



SCUOLA DI DOTTORATO
UNIVERSITÀ DEGLI STUDI DI MILANO-BICOCCA

Department of **Physics**

PhD program in **Physics and Astronomy**, **XXXVIII** cycle
Curriculum in **Subnuclear Physics**

Paving the Way for CUPID: Pre-Production,
Validation, and Optimization of Li_2MoO_4
Scintillating Cryogenic Calorimeters

Massimo Girola
Registration number: 813288

Tutor: **Chiara Brofferio**

Co-tutor: **Luca Gironi**

Supervisor: **Matteo Biassoni**

Coordinator: **Stefano Ragazzi**

Academic Year **2024–2025**

*'Cause what about all these children? And what about
all their parents?*

[...]

*And what about all their wishes wrapped up in garland
roses?*

[...]

*What about all these children? And all their children's
children?*

And why am I even wonderin' that today?

— *LDR*

Contents

Abstract	iv
1 Implications of Observing a Matter-Creating Process	1
1.1 The Neutrino and the β Decay	1
1.2 The Standard Model Neutrino	3
1.2.1 The Standard Model	3
1.2.2 Parity, helicity and chirality	3
1.2.3 The masses of the fermions	5
1.2.4 Accidental global symmetries in the Standard Model	7
1.3 Massive Neutrinos	7
1.3.1 Neutrino oscillations	8
1.3.2 Experimental measurement of the oscillation parameters	9
1.3.3 Explaining a nonzero neutrino mass	11
1.3.4 Experimental probes to the neutrino mass scale	14
1.4 Neutrinoless Double-Beta Decay	16
1.4.1 Double-beta decay	16
1.4.2 $0\nu\beta\beta$	17
1.5 The Relevance of $0\nu\beta\beta$ Searches	18
1.5.1 What distinguishes matter from antimatter?	18
1.5.2 Matter-antimatter asymmetry and a matter-creating process	19
1.5.3 What can $0\nu\beta\beta$ tell us about neutrinos?	19
2 Searching for $0\nu\beta\beta$ with Cryogenic Calorimeters: from CUORE to CUPID	22
2.1 Experimental Sensitivity to the Effective Majorana Mass	22
2.2 Experimental Search for $0\nu\beta\beta$	24
2.2.1 Experimental sensitivity to $0\nu\beta\beta$	24
2.2.2 Current and future experimental approaches	25
2.3 Cryogenic Calorimeters for $0\nu\beta\beta$ Searches	26
2.3.1 Cryogenic calorimeters	27
2.3.2 Advantages and limitations of cryogenic calorimeters	37
2.3.3 The CUORE experiment	38
2.4 The CUPID Experiment	43
2.4.1 Scintillating cryogenic calorimeters	43
2.4.2 The CUPID experiment: detector, backgrounds and sensitivity	47
2.5 Conclusion	52

3	CUPID Li_2MoO_4 Pre-Production	53
3.1	$\text{Li}_2^{100}\text{MoO}_4$ Crystal Requirements	53
3.2	Overview of The Crystal Production Process	55
3.2.1	CUPID crystal pre-production	56
3.3	Crystal Production	57
3.3.1	Precursor powders and Li_2MoO_4 synthesis	57
3.3.2	Crystal growth	58
3.3.3	Post-growth manufacturing	61
3.3.4	Crystal validation strategy	63
3.4	Conclusion	63
4	Validation of Pre-Production Li_2MoO_4 Crystals	64
4.1	The CUPID Pre-CCVR Campaign	64
4.1.1	Crystals in the early-stage pre-production phase	65
4.1.2	Pre-CCVR measurement campaigns	65
4.2	The Hall C Cryogenic Infrastructure for CUPID CCVR	66
4.2.1	Low-background and low-noise Hall C cryostat	66
4.2.2	Readout and biasing electronics	67
4.3	Analysis Pipeline	68
4.3.1	Digital signal processing	68
4.3.2	Treating signals from a dual-channel heat-light detector	69
4.3.3	From raw waveforms to the energy spectrum	69
4.3.4	Heat-light coincidences: Light Yield and Discrimination Power	75
4.4	Pre-CCVR1	76
4.4.1	Crystals and goals	76
4.4.2	Detector design and assembly	78
4.4.3	Calibration sources	81
4.4.4	Data taking	81
4.4.5	Thermalizations	82
4.4.6	Working point and noise in the light detectors	83
4.4.7	Light detectors calibration and performances	85
4.4.8	Working point for the heat channels	86
4.4.9	Characterization of $\text{Li}_2^{100}\text{MoO}_4$ NIIC crystals	87
4.5	Pre-CCVR2	94
4.5.1	Goals of the pre-CCVR2 campaign	94
4.5.2	First batch of SICCAS natural Li_2MoO_4 crystals	94
4.5.3	Detector design and assembly	96
4.5.4	Wiring and different thermalizations	97
4.5.5	Pre-CCVR2 (v2) data taking	100
4.5.6	Working point selection and noise in the LDs	101
4.5.7	Li_2MoO_4 calorimetric performance and WP optimization	104
4.5.8	Light Yield	112
4.5.9	Background measurement and radiopurity	114
4.6	Conclusion	123

5	Systematic Scintillation and Thermoluminescence Characterization of Li_2MoO_4 Crystals	124
5.1	Radioluminescence and Thermoluminescence Characterization of Scintillating Li_2MoO_4 Crystals	124
5.2	Crystal Samples	125
5.3	Experimental Setup	127
5.4	Radioluminescence and Low Temperature Scintillation Properties	128
5.4.1	Methodology	128
5.4.2	Results	129
5.4.3	Light Yield and RL properties	130
5.5	Thermoluminescence	133
5.5.1	Methodology	133
5.5.2	Results	133
5.6	Transmittance	136
5.7	Summary	138
6	Thin Al Coatings for Li_2MoO_4 Scintillating Cryogenic Calorimeters	139
6.1	Improving Light Collection Efficiency With Al Coatings	139
6.2	Thin Aluminum Coatings for Li_2MoO_4 Crystals	140
6.2.1	Coating depositions	140
6.2.2	Low-temperature properties of Al coatings	140
6.2.3	Radiopurity of Al coatings	141
6.2.4	Reflectance of Al coating	142
6.3	Impact of Al Coating on Calorimetric Performance	143
6.3.1	Cryogenic facility	143
6.3.2	Detector geometry and data taking	143
6.3.3	Calorimetric performance	145
6.4	Outlook	149
	Conclusions and Outlook	151
A	Dilution Cryostats	153
A.1	$^3\text{He}/^4\text{He}$ Dilution Cryostat: Design and Operation	154
	Bibliography	157
	Acknowledgements	176

Abstract

Neutrinoless double-beta decay ($0\nu\beta\beta$) is today regarded as the only experimentally viable probe of the nature of neutrinos. This hypothetical second-order weak nuclear process violates lepton-number conservation through the net creation of matter particles (a pair of electrons). Observing it would shed light on the origin of the matter–antimatter asymmetry in the Universe while establishing the Majorana nature of neutrinos.

CUPID is a next-generation experiment that will search for $0\nu\beta\beta$ of ^{100}Mo with scintillating cryogenic calorimeters. Operating 1596 isotopically enriched $\text{Li}_2^{100}\text{MoO}_4$ crystals with dual heat–light readout, it aims to reach the sensitivity required to explore the inverted-hierarchy region. The high Q -value of ^{100}Mo places the region of interest in an energy range free of most β/γ backgrounds, while particle identification through the scintillation channel suppresses degraded α events. Achieving this sensitivity imposes stringent requirements on every detector element. The $\text{Li}_2^{100}\text{MoO}_4$ crystals are the core component: they must deliver excellent calorimetric performance, high light yield, and ultra-high bulk and surface radiopurity. CUPID has tasked the Shanghai Institute of Ceramics (SICCAS) with large-scale production; however, because Li_2MoO_4 lacks an established industrial growth chain, SICCAS must develop a reliable, reproducible process that minimizes costly isotope losses while preserving radiopurity at every stage—objectives that are inherently in tension with one another.

CUPID and SICCAS are addressing these challenges in a *pre-production* phase focused on optimizing crystal growth. In this phase, SICCAS produces both natural and enriched crystals to tune growth parameters and requires CUPID to measure them for feedback on the achieved quality. In parallel, CUPID must develop a crystal-validation protocol—*CUPID Crystal Validation Runs* (CCVR)—that will be required to verify batch-by-batch compliance once full-scale production begins. Measurements on the pre-production crystals provide the ideal setting: CUPID can build and refine the CCVR protocol while performing the very measurements that SICCAS requires. This phase is called *pre-CCVR*.

This dissertation is situated at the early stage of pre-production, detailing my contributions to CUPID in the areas of crystal procurement and validation. I collaborated with SICCAS to document the complete growth process, from raw materials to finished crystals, so CUPID can establish a shared technical framework and common vocabulary with its producer for correlating growth parameters with detector performance. This knowledge will also serve as the basis for CUPID to develop a radiopurity control strategy that SICCAS will apply across the entire production chain. I then conducted two pre-CCVR campaigns: the first established reference performance using crystals from a previous supplier (Nikolaev Institute of Inorganic Chemistry); the second characterized the first natural Li_2MoO_4 batch grown by SICCAS, focusing on calorimetric and scintillation performance. Radiopurity was also assessed, though no strict targets were applied for this growth.

In parallel, I developed a reproducible strategy to assess calorimetric performance, encompassing working-point optimization, signal-noise characterization, and resolution metrics. This procedure will guide future pre-CCVR runs and, ultimately, the CCVRs during large-scale production to ensure consistent detector quality.

In addition to developing crystal validation and production procedures, this dissertation investigates the scintillation properties of Li_2MoO_4 . Scintillation is the cornerstone of CUPID's particle-identification capability, yet Li_2MoO_4 emits only modest light. The second part of this work therefore examines how this limitation can be better understood and mitigated by studying the influence of defects and impurities in pre-production crystals on light yield, and by exploring the use of reflective Al coatings to enhance light collection without introducing passive materials, assessing their impact on calorimetric performance.

Chapter 1

Implications of Observing a Matter-Creating Process

Since its postulation, the neutrino has played a central role in particle physics, cosmology, and our understanding of the universe. Initially proposed to account for the shape of the β -decay spectrum, it has proven to be a particle with far-reaching implications, challenging physicists since its theorization. For decades, researchers have developed theories and conducted experiments to characterize its properties. Yet, despite extensive efforts, many aspects of the neutrino remain mysterious.

Oscillation experiments have shown that neutrinos possess mass, contradicting the Standard Model, which treats them as massless. This discovery raises the basic question of how that mass is generated. Direct kinematic probes, such as beta-decay endpoint measurements, can determine the absolute mass scale but provide only limited information. Among indirect approaches, the search for neutrinoless double-beta decay is especially powerful. This hypothetical second-order nuclear transition converts energy into matter without producing antimatter. Detecting it would confirm lepton-number violation and, within specific theoretical frameworks, reveal the Majorana nature of neutrinos while constraining their effective mass. Finding such a process would point to new physics capable of explaining the matter-antimatter asymmetry observed in the Universe.

This chapter briefly introduces massive-neutrino physics, reviews the experiments that built our current understanding, and highlights the importance of neutrinoless double-beta decay in current particle-physics research.

1.1 The Neutrino and the β Decay

Postulation

In the early 1900s, studies on radioactivity led many scientists to investigate the properties of γ and β -rays. As early as 1899, before the experimental discovery of the atomic nucleus in 1911, Rutherford noticed that β -rays exhibited the same characteristics as electrons (e^-) [1]. A significant milestone in the history of weak interactions was reached in 1914 when Chadwick observed a continuous energy spectrum for the emitted electron in β -rays [2] with several monoenergetic

lines superimposed on it. A few years later, the monochromatic peaks were correctly attributed to an internal-conversion process associated with γ transitions [3], leaving the open problem of interpreting the continuous component in the e^- spectrum. Indeed, using modern notation, if we think about β -decay as a two-body process of the following form:

$$(A,Z) \rightarrow (A,Z+1) + e^-$$

then conservation of energy and momentum implies that the two final-state particles should be emitted back-to-back with equal and opposite momenta. In the absence of energy losses outside the active volumes of the detector, this would result in a monochromatic energy line for the emitted electron. This prediction conflicted with the measurements of Ellis and Wooster, who in 1927 measured the total energy released in the decay of Radium E (^{210}Bi) using a calorimeter [4] and observed a continuous spectrum, adding strength to the Chadwick results of 1914.

Although some physicists initially proposed that the law of energy conservation might be violated in subatomic processes such as β decay [1], Pauli suggested in 1930 that an additional, highly penetrating neutral particle might be present in the final state. This would imply that the decay is a three-body process, thereby preserving energy conservation while explaining the continuous spectrum. Building on Pauli's hypothesis, Fermi proposed a theoretical framework for β decay in 1934. He suggested that a neutron within the nucleus could transform into a proton while emitting an electron and a spin-1/2 neutral particle, which he named the "neutrino" [5].

Discovery

We now know that in the case of a β^- decay the emitted particle is an electron anti-neutrino ($\bar{\nu}_e$) and the β decay can be written as:

$$(A,Z) \rightarrow (A,Z+1) + e^- + \bar{\nu}_e \tag{1.1}$$

Although Fermi's theory was highly successful in describing β decay, direct detection of the neutrino was necessary to confirm its validity. Using his theory, it is possible to compute the cross-section for a process of the form:

$$\bar{\nu}_e + p \rightarrow n + e^+, \tag{1.2}$$

known as inverse β decay. The predicted cross section for this reaction is extremely small ($\sigma \sim 10^{-44} \text{ cm}^2$), making experimental detection very challenging. To enhance the interaction rate, both an intense flux of antineutrinos and a detector with a large active volume are required.

In 1956, Cowan and Reines used the intense antineutrino flux from the Savannah River nuclear power plant in the United States to search for the reaction described in Eq. 1.2, which has a clear experimental signature [6]. They observed a rate of ~ 3 interactions per hour when the reactor was operating and a significantly lower rate when it was shut down. From these measurements, it was possible to estimate the interaction cross-section, which was found to be in agreement with the theoretical prediction based on Fermi's theory.

So, after nearly 40 years, the puzzle of the continuous energy spectrum in β decay was finally resolved, and the existence of the neutrino had to be included in the theoretical description of weak interactions, an advancement which contributed to the development of the Standard Model as we know it today.

1.2 The Standard Model Neutrino

In the Standard Model, the neutrino is a massless fermion described by Weyl spinors. This is in contrast with the solid experimental evidence for neutrino oscillations, which implies a nonzero mass (this will be discussed in Sec. 1.3). This section gives a brief overview of the Standard Model detailing the role of the neutrino and introducing its properties.

1.2.1 The Standard Model

The Standard Model of particle physics (SM) is a quantum field gauge theory based on the symmetry group $SU(3)_c \times SU(2)_L \times U(1)_Y$. It represents our most accurate and comprehensive description of elementary particles and their interactions [7]. The $SU(2)_L \times U(1)_Y$ subgroup unifies electromagnetism and weak interactions within the Glashow–Weinberg–Salam model, forming the electroweak sector of the SM [8]. The $SU(3)_c$ component corresponds to Quantum Chromodynamics (QCD) [9], the gauge theory describing the strong interaction.

It is possible to categorize the particle content of the SM into fermions (spin-1/2 spinorial fields) and bosons (scalar fields for spin-0 particles and vector fields for spin-1 particles). In the theory, particles are quantized excitations of their corresponding fields, with bosons mediating the interactions. At this stage, it is important to note that the SM is formulated as a quantum field theory consistent with quantum mechanics and special relativity, thus assuming a flat spacetime. As a consequence, gravity (general relativity) is not included, which is generally considered one of the main weaknesses of the theory [10].

The Higgs mechanism [11] explains how the electroweak symmetry group $SU(2)_L \times U(1)_Y$ is spontaneously broken into $U(1)_{em}$ [8]. One important consequence of the spontaneous symmetry breaking of $SU(2)_L \times U(1)_Y$ is that the spin-1 gauge bosons are reorganized into one massless photon (A), which mediates the electromagnetic interaction $U(1)_{em}$, and three massive bosons: W^+ , W^- , and Z^0 [8]. A significant consequence of this mechanism is that the electric charge, gauge coupling constants, and the masses of the W and Z bosons are related by the weak mixing angle θ_W in specific mathematical relations. These predictions had been tested with high accuracy at LEP [12, 13].

Another relevant prediction of the theory is that the only scalar (spin-0) field in the theory, the Higgs field, acquires a vacuum expectation value (VEV) through spontaneous symmetry breaking. Quantum fluctuations around this VEV appear as the physical, massive Higgs boson. In 2012, the ATLAS and CMS experiments at the LHC confirmed this prediction by observing the Higgs boson and measuring its mass [14, 15].

Couplings to the Higgs field not only generate mass for the gauge bosons but also for the fermions except the neutrinos, which remain massless in the SM. To understand why, it is necessary to introduce the concept of chirality.

1.2.2 Parity, helicity and chirality

Parity

In the first half of the last century, physicists were convinced that the laws of nature were invariant under a parity transformation. In 1956, Lee and Yang questioned this principle in weak interactions [16]. Shortly after, Wu designed an experiment to test parity violation in β -decays [17]. Thanks to the technological progress achieved in those years, Wu was able to cool down at the \sim mK level

some ^{60}Co nuclei using the adiabatic demagnetization technique. At such temperatures, the thermal agitation is so low that the ^{60}Co nuclei can keep their spin aligned to the magnetic field lines, fixing a preferential direction for the experiment. ^{60}Co undergoes β -decay transitioning into an excited state of Nickel ($^{60}\text{Ni}^{**}$) that has one unit less of nuclear spin. The conservation of total angular momentum implies that the spin of both produced (spin-1/2) fermions are aligned to the same direction of the ^{60}Co spin, which is fixed by the orientation of the magnetic field. The conservation of linear momentum imposes that the daughters e^- and $\bar{\nu}_e$ are emitted back-to-back. Under parity transformation, being a pseudo-vector, the spin does not change sign. On the contrary, the linear momenta of the products, being a vector, flips under parity, reversing their direction of emission with respect to the spin. Thus, for a given orientation of the ^{60}Co spin, if the weak interaction conserves parity, then the two cases (e^- emitted in the same direction of the spin or in the opposite direction) should be equally likely. What Wu observed using a β counter was that electrons were emitted preferentially in the direction opposite to the orientation of the spin of the ^{60}Co , resulting in a β asymmetry. This means that β -decay violates parity, or, in other words, our universe has a preference for a defined handedness, a concept that can be better understood introducing helicity.

Helicity of the neutrino

If we define the helicity as the normalized projection of the spin \vec{S} along the momentum \vec{p} :

$$H = \frac{\vec{S} \cdot \vec{p}}{|\vec{p}|}$$

it is easy to see that H changes sign under parity, and particles with spin (anti-)parallel to the momentum have $H = +1/2$ ($H = -1/2$). Commonly, particles with $H = -1/2$ ($H = +1/2$) are called left-handed (right-handed), although this can be misleading as this terminology is also used for referring to the chiral states, introduced later. If the neutrino is massless, it is impossible to find a reference frame where \vec{p} is reversed or in which $\vec{p} = 0$ (rest frame); therefore, a massless neutrino has a defined helicity state. Landau [18], Lee and Yang [19], and Salam [20] independently concluded that such massless neutrinos would have a determined and opposite helicity with respect to anti-neutrinos [21]. Shortly after Wu experimentally verified the parity violation, another important step in the comprehension of the neutrino was made when Goldhaber came up with a clever idea to measure the helicity for ultra-relativistic neutrinos [22]. He used a decay chain involving ^{132}Eu , in which a neutrino and a photon having the same helicity are emitted. By measuring the helicity of the photon he demonstrated that only left-handed neutrinos (neutrinos with $H = -1/2$) are produced in the weak interaction, confirming that parity violation in weak interactions is maximal. Taking the words from Ref. [23]: *if neutrinos are supposed to be massless, helicity distinguishes them from antineutrinos.*

Chirality

All these experimental pieces of evidence helped to conclude that weak interactions have a vector-minus-axial (V-A) structure. Using the Dirac notation, it is possible to introduce the concept of chirality as the eigenvalue of the Dirac matrix γ^5 . Being a Lorentz-invariant operator, contrary to helicity, chirality is always defined: left-handed particles have an eigenvalue of -1, while right-handed particles have an eigenvalue of +1. This is often a source of confusion, as the terminology introduced earlier for helicity is the same, although the two concepts are different. Indeed, chirality

is always defined, and it coincides with helicity only for massless fermions. If the fermion is massive, a defined chirality state can be thought of as a superposition of left- and right-handed helicity states. In particular, if we consider an ultra-relativistic left chiral state, this will be mainly a left-handed helicity state, with a small right-handed helicity component $\mathcal{O}(m/E)$. The γ^5 matrix is used to introduce the left-chiral projector $P_L = (1 - \gamma^5)/2$ that, when applied to a generic fermion ψ , extracts the left handed component: $P_L\psi = \psi_L$. The V-A structure for weak interactions involves vertices in which a matrix $\gamma^\mu(1 - \gamma^5)/2$ is inserted between the spinor and the adjoint spinor [24]. Having introduced the left-chirality projector, it is easy to see that a direct consequence of the V-A structure is that, taking the words from Ref. [24] *the weak charged current acts only on the left-handed part of particle wavefunctions and the right-handed part of antiparticle wavefunctions*. If not specified otherwise, in the following discussion the handedness will refer to the chirality and not to the helicity.

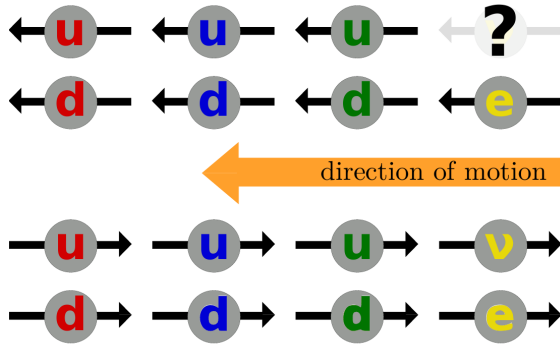
1.2.3 The masses of the fermions

Having introduced the concept of chirality, it is possible to express the fermion content of the SM as follows [26]:

$$\left\{ \begin{pmatrix} u_i \\ d_i \end{pmatrix}_{\text{rgb,L}}, (u_i)_{\text{rgb,R}}, (d_i)_{\text{rgb,R}}, \begin{pmatrix} \nu_i \\ \ell_i \end{pmatrix}_L, (\ell_i)_R \right\}_{i=1,2,3} \quad (1.3)$$

where $i = 1, 2, 3$ indicates the three generations of fermions observed in nature: up-type quarks $u_i = (u, c, t)$, down-type quarks $d_i = (d, s, b)$, charged leptons $\ell_i = (e, \mu, \tau)$, and neutral leptons $\nu_i = (\nu_e, \nu_\mu, \nu_\tau)$, the latter more commonly called neutrino flavors. {r, g, b} are the charges of the $SU(3)_C$ symmetry, and L and R are used to indicate defined handedness chirality states.

In this framework, the fermions are organized into a left-handed doublet under $SU(2)_L$, participating in weak interactions, and a right-handed singlet under $SU(2)_L$, that does not take part to the charged-current weak interactions, as shown in Table 1.1. This not only explains the results of the Wu and Goldhaber experiments but is also the deeper theoretical reason why parity is



Symbol	Multiplicity		Y	L	B
	$SU(3)_c$	$SU(2)_L$			
q_L doublet	3	2	+1/6	0	$\frac{1}{3}$
u_R singlet	3	1	+2/3	0	$\frac{1}{3}$
d_R singlet	3	1	-1/3	0	$\frac{1}{3}$
ℓ_L doublet	1	2	-1/2	1	0
e_R singlet	1	1	-1	1	0

Figure 1.1: Helicity of the 15 matter particles (in the ultra-relativistic limit) in each SM family. The arrow indicates the momentum direction. Reprinted from Ref. [25].

Table 1.1: SM fermions. Hypercharge is assigned according to $Q = T_{3L} + Y$. L and B are the lepton and the baryon numbers. Adapted from Ref. [25].

maximally violated in weak interactions. If we now focus on the leptons, we notice that there are no right-handed neutral fields in the model, as illustrated schematically in Fig. 1.1. This captures a key property of the neutrinos: after decades of experimental research *we have never detected a neutrino with right-handed chirality (or an anti-neutrino with left-handed chirality)*. The absence of right-handed neutral leptons in the SM leaves neutrinos without any gauge-invariant mechanism for acquiring mass, a point that will be clarified in the following discussion. By contrast, right-handed singlets *do* exist for the charged leptons, enabling a gauge-invariant term in the SM Lagrangian that generates their masses.

Before continuing, it's worth to notice that electromagnetism couples with the same electric charge to both chiralities of a four-component charged Dirac fermion, $\psi = \psi_L + \psi_R$; hence, the electromagnetic interaction is vector-like and conserves parity, showing no preference for handedness.

Returning to the matter at hand, the presence of right-handed charged leptons allows the SM to include a $SU(2)_L \times U(1)_Y$ gauge-invariant term that couples the Higgs doublet $\phi^T = (\phi_0, \phi_-)$ to the charged-lepton fields [27]:

$$-\mathcal{L}_{\text{Yukawa}} = g_i \begin{pmatrix} \nu_i \\ \ell_i \end{pmatrix}_L \phi^* (\ell_i)_R + h.c. \quad (1.4)$$

where $i = e, \mu, \tau$ labels the three generations and g_i are the Yukawa couplings, free parameters of the theory (analogous Yukawa terms can be written for the quarks, but the present discussion, being focused on the neutrino, is restricted to leptons).

After spontaneous symmetry breaking the neutral Higgs component acquires the vacuum expectation value $\langle \phi_0 \rangle = v/\sqrt{2}$, with $v = (\sqrt{2}G_F)^{-1/2} \simeq 246$ GeV, where G_F is the Fermi Constant. Equation (1.4) then becomes

$$-\mathcal{L}_{\text{Yukawa}} = g_i \frac{v}{\sqrt{2}} \bar{\ell}_i \ell_i$$

or, defining $m_i = g_i v/\sqrt{2}$,

$$-\mathcal{L}_{\text{Yukawa}} = m_i \bar{\ell}_i \ell_i = m_i (\bar{\ell}_{i,L} \ell_{i,R} + \bar{\ell}_{i,R} \ell_{i,L}) \quad (1.5)$$

For a more concrete example, we can focus on the electron. Its Yukawa term reads [28]:

$$-\mathcal{L}_{\text{Yukawa}}^e = g_e \frac{v}{\sqrt{2}} (\bar{e}_L e_R + \bar{e}_R e_L) = m_e (\bar{e}_L e_R + \bar{e}_R e_L)$$

where we can identify $m_e = g_e v/\sqrt{2}$ as the mass of the electron.

Equation (1.5) explicitly reveals that fermion masses arise as Dirac terms of the form $-m\bar{\psi}\psi$. Such terms *cannot* be written directly in the unbroken SM Lagrangian because they would violate $SU(2)_L$ invariance. By contrast, the Yukawa operator in Eq. (1.4) is invariant under $SU(2)_L$ and only generates mass after symmetry breaking. Equation (1.5) also shows that mass terms always couple fields of opposite chirality [27]; since there are no right-handed neutral leptons in the theory, SM neutrinos are massless. Observations of nonzero neutrino masses, therefore, point unambiguously to physics beyond the SM.

The neutrino mass puzzle

Neutrino oscillations, which are briefly introduced in Sec. 1.3.1, prove unequivocally that at least two neutrino states have nonzero mass, contradicting the SM. This observation compels us to extend the SM to accommodate neutrino masses.

One possibility would be to simply introduce one right-handed neutrino singlet $N_{i,R}$ per generation, augmenting the fermion content in Eq. (1.3). Since $N_{i,R}$ is an $SU(2)_L$ singlet, it is possible to introduce a gauge-invariant Yukawa interaction analogous to Eq. (1.4); after electroweak symmetry breaking, this term generates Dirac masses for the neutrinos just like it happens for the other charged fermions [29]. In this case the Weyl spinors $N_{i,R}$ constitute genuinely new degrees of freedom that would pair to the left-handed neutrino to form a four-component Dirac fermion.

A somehow more intriguing option is that the right-handed component is *not* an independent field but the charge-conjugate of the left-handed neutrino itself. In that case, the neutrino is a Majorana fermion, and the corresponding mass term is called *Majorana* mass.

These two possibilities (Dirac or Majorana fermions), together with neutrino oscillations, will be the topic of Sec. 1.3, which discusses the physics of massive neutrinos.

1.2.4 Accidental global symmetries in the Standard Model

One interesting feature of the Standard Model is the presence of several *accidental* global (non-local) symmetries. These are symmetries that emerge automatically in the SM, without being imposed *a priori* [21]. The conservation laws associated with these symmetries are: the baryon number B , the three lepton numbers L_e, L_μ, L_τ , and any linear combination such as the total lepton number $L = L_e + L_\mu + L_\tau$. Some of these symmetries are *anomalous*, meaning that they are respected only at the leading (perturbative) order. This means that we can find processes described by higher-order perturbative expansions that violate their associated conservation laws, such as it happens for $B+L$ [30, 31, 32].

After accounting for the anomalies, the SM retains the following *exact* (anomaly-free) global symmetries [21]:

$$B - L \quad L_e - L_\mu \quad L_\mu - L_\tau \quad (1.6)$$

and any linear combination of these.

Neutrino oscillations, briefly reviewed in Sec. 1.3.1, break $L_e - L_\mu$ and $L_\mu - L_\tau$ by mixing the flavor eigenstates, thereby exposing a limitation of the SM and leaving $B - L$ as the sole residual global symmetry. To date, no laboratory process violating $B - L$ has been observed; the only indirect evidence is the cosmological excess of matter over antimatter (see Section 1.5.2). As the authors of Ref. [21] emphasize, if $B - L$ is conserved, one can unambiguously distinguish matter from antimatter, in agreement with the SM predictions; if it is violated, we would expect to see matter/anti-matter transitions. Being the only neutral fermion, a transition of the type $\nu \rightarrow \bar{\nu}$ is particularly appealing, placing the neutrino at center stage in any test of this symmetry, as better explained in Sec. 1.5.

1.3 Massive Neutrinos

As introduced in Sec. 1.2, the neutrino played a central role in formulating the Standard Model. Yet, it is the same neutrino that now exposes the theory's limitations. Over time, a series of experimental anomalies converged toward a coherent picture: neutrinos oscillate between flavours, proving that they have nonzero mass and violating two of the global accidental exact symmetries of the SM. This section focuses on the physics of massive neutrinos mentioning both the experimental observables and the theoretical implications of a non-vanishing ν mass.

1.3.1 Neutrino oscillations

I oscillate...

The first evidence for neutrino oscillations came from measurements of the solar neutrino flux at the Homestake mine (now part of the Sanford Underground Research Facility). The detector observed solar neutrinos through the $^{37}\text{Cl}(\nu_e, e^-)^{37}\text{Ar}$ reaction and detected only about $\sim \frac{1}{3}$ [33] of the expected flux [34]. Other radiochemical [35, 36] and water-Cherenkov [37] experiments later confirmed the same anomaly. Bruno Pontecorvo had already suggested in the late 1950s that neutrinos might oscillate between flavors in vacuum [38, 39], which could be a possible solution to this “solar neutrino puzzle”. Initial skepticism was laid to rest in 2001 by the results from the SNO heavy-water Cherenkov detector, which was sensitive to all neutrino flavors and, therefore, was able to confirm that the total neutrino flux was meeting the expectations, with only around $\sim \frac{1}{3}$ of the total flux to be attributed to the electron flavor ν_e [40]. Several other experiments confirmed neutrino oscillations from non-solar sources, such as KamLAND for reactor anti-neutrinos [41] and Super-Kamiokande for atmospheric neutrinos [42], providing further support for the existence of neutrino oscillations.

Today, neutrino oscillation experiments play a central role in the effort to understand and characterize neutrino physics, as outlined in the next section.

...thus I have a mass

The theory of neutrino oscillations in vacuum was formulated by Maki, Nakagawa, and Sakata [43] and by Pontecorvo [44] in the 1960s, concurrently with the development of the SM. It describes a process in which, during propagation, a neutrino oscillates between its generational counterparts. This phenomenon arises naturally from basic quantum mechanical considerations [45], provided that neutrinos possess nonzero mass and that the mass eigenstates (propagation eigenstates) are not equivalent to their flavor eigenstates (weak interaction “production” or “detection” eigenstates). In other words, there is mixing between the flavor eigenstates (ν_e, ν_μ, ν_τ) and the mass eigenstates (ν_1, ν_2, ν_3), with each neutrino mass state being a superposition of all three flavors. This mixing is described by the 3×3 complex Pontecorvo-Maki-Nakagawa-Sakata matrix (\mathbf{U}_{PMNS}), which unitarily transforms between the two bases via the matrix product:

$$|\nu_\alpha\rangle = \sum_{i=1,2,3} U_{\text{PMNS}}^{\alpha i} |\nu_i\rangle, \quad |\nu_i\rangle = \sum_{\alpha=e,\mu,\tau} U_{\text{PMNS}}^{*\alpha i} |\nu_\alpha\rangle$$

In general, 3×3 unitary matrices have nine free parameters: three of these can be interpreted as rotation angles (commonly denoted as $\theta_{12}, \theta_{23}, \theta_{13}$), while the remaining six correspond to complex phases. One of these is the CP-violating phase (δ_{CP}), which is responsible for CP violation in the leptonic sector. The other five are relative phases that do not affect neutrino oscillations. If neutrinos are Dirac particles, all of these five phases can be absorbed into the definitions of the neutrino states. If, on the other hand, neutrinos are Majorana particles, two of these phases, α_1 and α_2 —known as Majorana phases—cannot be “rotated away”. The PMNS matrix can be parametrized using the following conventional form [46]:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \cdot \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \cdot \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} e^{i\alpha_1} & 0 & 0 \\ 0 & e^{i\alpha_2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (1.7)$$

where $s_{ij} = \sin(\theta_{ij})$ and $c_{ij} = \cos(\theta_{ij})$. For now, the matrix containing the Majorana phases can be ignored, as these phases do not affect neutrino oscillations [47, 48, 49]. However, these Majorana phases are highly relevant for understanding the fundamental nature of neutrinos and arise because, for a Majorana particle, the right-handed component is not independent of the left-handed one; this reduced freedom to re-phase the neutrino fields prevents those two phases from being rotated away. This possibility and its implications will be explored in more detail in Sec. 1.3.3 and Sec. 1.5.3.

Having introduced U_{PMNS} , using natural units ($\hbar = c = 1$), we can express the time evolution of the neutrino of a given flavor α as [50]:

$$|\nu_\alpha(t)\rangle = \sum_{i=1,2,3} U_{\alpha i}^* e^{-iE_i t} |\nu_i\rangle$$

where $E_i = \sqrt{p_i^2 + m_i^2}$. The probability of finding neutrino of flavor α having oscillated into another flavor β after a given time can then be calculated by taking the modulus square of the projection:

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \left| \sum_i U_{\beta i} U_{\alpha i}^* e^{-iE_i t} \right|^2$$

Assuming that the neutrino is ultra-relativistic and using natural units ($t = L$), this can be expanded as $E_i \sim p + \frac{m_i^2}{2E}$ (using $E \simeq p$) obtaining [50]:

$$\begin{aligned} P_{\alpha \rightarrow \beta}(L, E) = & \delta_{\alpha\beta} - 4 \sum_{i < j} \Re [U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin^2 \left(\frac{\Delta m_{ji}^2 L}{4E} \right) \\ & + 2 \sum_{i < j} \Im [U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}] \sin \left(\frac{\Delta m_{ji}^2 L}{2E} \right) \end{aligned} \quad (1.8)$$

What oscillations reveal about the SM and the neutrinos

Equation (1.8) provides the transition probability between any two flavours α, β . The cases where $\alpha = \beta$ are referred to as ‘‘disappearance’’, while those with $\alpha \neq \beta$ are called ‘‘appearance’’. Notably, the case $\alpha = \mu$ and $\beta = e$ describes the neutrino oscillation $\nu_\mu \rightarrow \nu_e$, observed for example by T2K [51] and NOvA [52]. This transition violates by two units the exact global accidental symmetry of the SM, $L_e - L_\mu$, introduced in Eq. (1.6): $\Delta(L_e - L_\mu) = 2$ [23]. Similarly, the appearance of tau neutrinos ($\nu_\mu \rightarrow \nu_\tau$), observed for example by OPERA [53] and Super-Kamiokande [54], violates $L_\mu - L_\tau$ by two units: $\Delta(L_\mu - L_\tau) = 2$ [23].

One extremely relevant consequence of Eq. (1.8) is that the transition probability depends only on the *differences* of the squares of the masses of the mass eigenstates, $\Delta m_{ij} := m_i^2 - m_j^2$. This means that oscillation experiments cannot determine the absolute mass scale of neutrinos and are sensitive only to the mass splittings. Moreover, the Majorana phases do not affect the transition probability, and oscillation experiments cannot provide any information about them [47, 48, 49].

1.3.2 Experimental measurement of the oscillation parameters

Oscillation experiment data can be used to perform global fits that allow the oscillation parameters to be measured with unprecedented accuracy [46], demonstrating that we have now entered the precision era of neutrino physics.




Figure 1.2: The two possible neutrino mass orderings. In the normal ordering (NO), the mass hierarchy is $m_1 < m_2 < m_3$, while in the inverted ordering (IO), it is $m_3 < m_1 < m_2$. The colors represent the probability of observing a specific mass eigenstate within a given flavor eigenstate. The curved portions of the colored regions show how these probabilities vary as δ_{CP} changes from 0 to 2π . Reprinted from Ref. [55].

From Eq. (1.8), we can deduce that the sensitivity of a given experiment depends on the probability “angle” $\propto \frac{\Delta m_{ij}^2 L}{E}$. This has profound implications for the parameters to which an oscillation experiment is sensitive. In fact, for a given Δm_{ij}^2 , if L/E becomes too small, the appearance (disappearance) probability approaches zero (one), making it impossible to observe any oscillation effect. Conversely, if L/E becomes too large, the oscillations occur so rapidly that an experiment cannot effectively resolve the differences between the three mass states. For this reason, different experiments will be maximally sensitive to certain oscillation parameters and less sensitive to others, depending on the energy spectrum and the distance L from their neutrino source(s). For example, solar neutrino experiments detect electron neutrinos produced in fusion reactions within the Sun and are maximally sensitive to Δm_{12}^2 and θ_{12} . By exploiting matter effects inside the Sun (MSW effect [56, 57]), it is possible to determine that $\Delta m_{21}^2 = m_2^2 - m_1^2$ is positive and much smaller than $|\Delta m_{31}^2|$. The sign of Δm_{31}^2 is still unknown. This implies that there are two non-equivalent orderings for the neutrino mass: the normal ordering (NO), with $m_1 < m_2 < m_3$, and the inverted ordering (IO), with $m_3 < m_1 < m_2$. The two orderings are illustrated in Fig. 1.2. In particular, adopting the same conventions as in Ref. [46], we can define:

$$\Delta m_{3\ell}^2 \quad \text{with} \quad \begin{cases} \ell = 1 & \text{for } \Delta m_{3\ell}^2 > 0: \text{ normal ordering (NO),} \\ \ell = 2 & \text{for } \Delta m_{3\ell}^2 < 0: \text{ inverted ordering (IO).} \end{cases}$$

Adopting this convention, we can express the three mass states as a function of the mass of the lightest neutrino state m_{lightest} and the two mass splittings:

$$\begin{cases} \text{NO:} & m_1 = m_{\text{lightest}}, & m_2 = \sqrt{m_{\text{lightest}}^2 + \Delta m_{21}^2}, & m_3 = \sqrt{m_{\text{lightest}}^2 + \Delta m_{3\ell}^2}, \\ \text{IO:} & m_3 = m_{\text{lightest}}, & m_2 = \sqrt{m_{\text{lightest}}^2 - \Delta m_{3\ell}^2}, & m_1 = \sqrt{m_{\text{lightest}}^2 - \Delta m_{3\ell}^2 - \Delta m_{21}^2}, \end{cases} \quad (1.9)$$

The most recent global fit has been performed in two versions: a baseline analysis that uses all public data with independently reproducible likelihoods, and an extended analysis that also incorporates the most recent atmospheric-neutrino likelihood released by Super-Kamiokande [46]. The results of

Table 1.2: Best-fit points (bfp), 1σ errors, and 3σ ranges for the oscillation parameters determined by the global fit on the most recent oscillation data excluding Super-Kamiokande atmospheric data. Note the opposite sign of $\Delta m_{3\ell}^2$ in the two hierarchies and the contrasting behaviour of the CP phase: for normal ordering δ_{CP} is close to the CP-conserving value 180° , whereas for inverted ordering it moves toward the maximally CP-violating point 270° . Also observe that the 3σ interval for θ_{23} still spans the maximal-mixing value 45° . Table adapted from [46].

Parameter	Normal Ordering		Inverted Ordering	
	bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range
θ_{12} [$^\circ$]	$33.68^{+0.73}_{-0.70}$	31.63 \rightarrow 35.95	$33.68^{+0.73}_{-0.70}$	31.63 \rightarrow 35.95
θ_{23} [$^\circ$]	$48.5^{+0.7}_{-0.9}$	41.0 \rightarrow 50.5	$48.6^{+0.7}_{-0.9}$	41.4 \rightarrow 50.6
θ_{13} [$^\circ$]	$8.52^{+0.11}_{-0.11}$	8.18 \rightarrow 8.87	$8.58^{+0.11}_{-0.11}$	8.24 \rightarrow 8.91
δ_{CP} [$^\circ$]	177^{+19}_{-20}	96 \rightarrow 422	285^{+25}_{-28}	201 \rightarrow 348
$\Delta m_{21}^2/10^{-5}$ [eV^2]	$7.49^{+0.19}_{-0.19}$	6.92 \rightarrow 8.05	$7.49^{+0.19}_{-0.19}$	6.92 \rightarrow 8.05
$\Delta m_{3\ell}^2/10^{-3}$ [eV^2]	$+2.534^{+0.025}_{-0.023}$	+2.463 \rightarrow +2.606	$-2.510^{+0.024}_{-0.025}$	-2.584 \rightarrow -2.438

the fit for the first case are shown in Table 1.2. The values of the oscillation parameters for each mass ordering still present some ambiguities:

- The **mixing angle** θ_{23} suffers from octant ambiguity: both octants are compatible with the data, and the 3σ precision remains about 20%; values close to 45° are not excluded.
- **Leptonic phase** δ_{CP} : For NO, the preferred values lie close to the CP-conserving point $\delta_{\text{CP}} \sim 180^\circ$; for IO, the preference shifts to $\delta_{\text{CP}} \sim 270^\circ$, corresponding to nearly maximal CP violation and making CP-conserving values less compatible with the data.
- **Mass ordering**: In the baseline fit, tensions between appearance and disappearance data leave normal and inverted orderings equally plausible. Adding the Super-Kamiokande data introduces a mild preference for NO, but the evidence remains inconclusive.

Forthcoming experiments, including the recently-launched JUNO (which just began data-taking) [58] and the planned Hyper-Kamiokande [59] and DUNE [60], should achieve the precision needed to resolve the mass-ordering ambiguity and tighten the remaining gaps in our oscillation picture. However, oscillation measurements alone cannot determine the absolute neutrino mass scale or reveal whether neutrinos are Dirac or Majorana particles; addressing these questions requires dedicated non-oscillation experiments.

1.3.3 Explaining a nonzero neutrino mass

Section 1.2.3 introduced how the mass of the fermions is generated through Yukawa couplings in the SM after spontaneous symmetry breaking. Lacking a right-handed neutral lepton, the SM does not allow the addition of $\text{SU}(2)_L$ invariant Yukawa terms for the neutrino in the unbroken Lagrangian. This leaves the neutrino without a mass, in contrast with the experimental evidence from neutrino oscillations.

Dirac mass terms

A straightforward way to give neutrinos a mass is to introduce one right-handed neutrino for each generation and couple them through a Yukawa interaction similar to the one in Eq. (1.4) that gives mass to the other fermions. After symmetry breaking, such a term would generate a neutrino mass term of the form [29]:

$$-\mathcal{L}_m = (m_D)_{\alpha\beta} \overline{\nu_{\alpha L}} \nu_{\beta R} + \text{h.c.} \quad (1.10)$$

where $m_D = h^\nu \langle H^0 \rangle$ and $\langle H^0 \rangle = v/\sqrt{2}$, with v having the value introduced in Sec. 1.2.3. One can diagonalize the matrix $(m_D)_{\alpha\beta}$ and unitarily transform the flavor eigenstates to the mass eigenstates. In the mass basis, one can introduce the Dirac neutrino field:

$$\nu_i := \nu_{iL} + \nu_{iR}$$

where i is an index representing a particular mass eigenstate. By doing so, Eq. (1.10) can be expressed as:

$$-\mathcal{L}_m = \sum_i m_i \bar{\nu}_i \nu_i$$

In this way, each neutrino mass eigenstate ν_i has four degrees of freedom: left- and right-handed neutrino and antineutrino chiralities, similarly to the other charged fermions. Such a Dirac mass term preserves the global lepton number symmetry corresponding to a U(1) phase transformation $\nu_j \rightarrow e^{i\theta} \nu_j$ [29].

With the current bounds on the absolute neutrino mass scale (which will be discussed in Section 1.3.4), we would obtain Yukawa couplings of order $\sim 10^{-13}$ – 10^{-12} . These are extremely small compared to those of the other fermions, which range from that of the electron $\sim 10^{-6}$ to that of the top quark ~ 1 . One option is that these coupling constants simply span across thirteen orders of magnitude. However, the community generally agrees that such fine-tuning of the Yukawa couplings would be rather unnatural.

Majorana mass terms

Given a spinor ψ , one can introduce the conjugate spinor ψ^C using the charge conjugation matrix C : $\psi^C = C\bar{\psi}^\top$. If this spinor is that of a charged fermion obeying the Dirac equation in the presence of an electromagnetic field, the spinor ψ^C obeys the Dirac equation with opposite charge and the same mass. For this reason, ψ^C is called the *antiparticle field*. Recalling the definition of the chiral projector given in Sec. 1.2.2, one can show that $(\psi_L)^c = (P_L \psi)^c = P_R \psi^c = (\psi^c)_R$ [29]. The conjugation operation, therefore, also flips chirality for a given spinor component. The field ψ_L and its conjugate $(\psi_L)^c$ contain the same two degrees of freedom, corresponding to the left-handed particle and the right-handed antiparticle. With this formalism, it is possible to introduce the Majorana field:

$$\chi = \psi_L + e^{i\theta} (\psi_L)^c$$

If we take the conjugate field χ^c , we see that a *Majorana field coincides with the original field χ up to an overall phase factor*, called Majorana phase¹: $\chi^c = (\psi_L)^c + e^{-i\theta} \psi_L = e^{-i\theta} \chi$ [27]. One can show that a generic fermion with four degrees of freedom (such as the electron) can be described by its two chiral components (ψ_L, ψ_R) but also by two independent Majorana fields (χ_1, χ_2) with

¹The Majorana phases do not affect neutrino oscillation probabilities. However, they play a significant role in the effective Majorana mass introduced in Sec. 1.5.3.

$\chi_1 = \psi_L + (\psi_L)^c$ and $\chi_2 = \psi_R + (\psi_R)^c$ [29]. This four-component description would also apply to neutrinos if their masses were of the Dirac type. On the other hand, if neutrinos were Dirac *and* massless, the right- and left-handed components would be completely independent (due to the absence of mass terms coupling opposite chiralities). Given the chiral (V-A) nature of weak interactions (which only couple to left-handed components), one logical possibility would be to simply omit the right-handed component from the theory, as mentioned in Sec. 1.2.3, since it would not participate in any interaction.

With the formalism introduced so far, for a neutral fermion, one could consider the possibility of using just two degrees of freedom to write a mass term. This idea was suggested initially by Majorana² in 1937 [61]. Considering the antiparticle field for the neutrino $(\nu_L)^c$, which is a right-handed field for the reasons mentioned earlier, we can introduce by hand a mass term of the form [29]:

$$-\mathcal{L}_M = \frac{1}{2}m \left[\overline{\nu_L}(\nu_L)^c + \overline{(\nu_L)^c}\nu_L \right].$$

Such a mass term couples fields of opposite chirality, making the left-handed and right-handed components no longer independent of one another. In contrast to the Dirac mass term introduced earlier, this term *does not preserve any U(1) phase symmetry and explicitly violates lepton number conservation* [29]. This violation has profound and far-reaching implications, as it would enable matter-creating processes such as $\frac{A}{Z}X \rightarrow \frac{A}{Z+2}X' + 2e^-$ (commonly known as neutrinoless double-beta decay), which will be discussed in Sec. 1.4.

One can introduce the Majorana neutrino field:

$$\nu := \nu_L + (\nu_L)^c$$

which under conjugation transforms into itself: $\nu^c = \nu$. With this field, the mass term becomes [29]:

$$-\mathcal{L}_M = \frac{m}{2}\nu\bar{\nu} \tag{1.11}$$

In this way, a mass term is introduced into the theory by construction. This is quite different from the case of Dirac terms, where the masses arise naturally via the Higgs mechanism after spontaneous symmetry breaking. Various extensions of the SM contain mechanisms that predict Majorana mass terms. One simple model is the so-called “see-saw” mechanism, which can naturally explain the smallness of neutrino masses compared to those of the other fermions.

Neutrino mass mechanisms

Some beyond Standard Model frameworks foresee an additional SU(2) singlet, referred to as N_R . Using this right-handed singlet, the most general mass term that can be added to the Lagrangian involving ν_L and N_R is of the form [29]:

$$-\mathcal{L}_M = \frac{1}{2} \begin{pmatrix} \overline{\nu_L} & \overline{(N_R)^c} \end{pmatrix} \begin{pmatrix} m_T & m_D \\ m_D^\dagger & M \end{pmatrix} \begin{pmatrix} (\nu_L)^c \\ N_R \end{pmatrix} + \text{h.c.} \tag{1.12}$$

The matrix introduced in this Lagrangian is referred to as the mass matrix, and different neutrino mass mechanisms emerge depending on the choice of this matrix. For example, choosing

²Majorana proposed this in a paper titled *Symmetric theory of the electron and positron* [61], which aimed to eliminate the negative energy states arising from the Dirac equation.

$m_T = M = 0$ corresponds to Dirac neutrinos, while nonzero entries for m_T or M lead to Majorana mass terms (making the neutrinos Majorana particles). One common choice is the so-called *Type-I see-saw model*, where $m_T = 0$ and $M \gg m_D$. In this case, the matrix in Eq. (1.12) may be block-diagonalized³ to yield effective Majorana masses. After diagonalization, one obtains the mass eigenstates, which are two Majorana fields with masses $m_{\text{light}} \simeq m_D^2/M$ and $m_{\text{heavy}} \simeq M$ [29]. This means that the light neutrino masses are suppressed by the large Majorana mass scale M . Hence, the smallness of the known neutrino masses is elegantly explained by the largeness of M , a mechanism after which the model is called “see-saw”. In particular, if we consider a Yukawa coupling of the neutrinos such that m_D is at the electroweak scale $\mathcal{O}(100 \text{ GeV})$, then a 10 meV light neutrino corresponds to a heavy mass scale of order 10^{15} GeV . Such energy scales are close to those of Grand Unification Theories (GUTs), such as $\text{SO}(10)$ [29]. Other mass mechanisms, such as Type-II see-saw and double see-saw, are also possible, as detailed in Ref. [62].

1.3.4 Experimental probes to the neutrino mass scale

As explained in Sec. 1.3.1, oscillation data fix only the mass splittings and leave the absolute neutrino mass scale yet to be determined. To complete the picture, we need complementary approaches, which can be categorized into two main groups. *Direct* methods aim to determine the mass kinematically and are model-independent, while *indirect* methods extract it from model-dependent observables. Direct measurements are the gold standard but are technically challenging because of the smallness of the neutrino mass.

Direct measurements

Using our knowledge of β -decay or other equivalently well-understood processes such as electron capture or inverse β -decay, it is possible to measure the energy of the charged particles in the final state of these reactions to infer the mass of the neutrino. In fact, if we consider, for example, β -decay, expressed in Eq. (1.1), in a simplified⁴ model, we can express the differential electron spectrum as proportional to [27]:

$$\frac{dN}{dE_e} \propto p_e E_e (Q - E_e) \sqrt{(Q - E_e)^2 - m_{\bar{\nu}_e}^2}, \quad (1.13)$$

where p_e , E_e , and m_e are, respectively, the momentum, energy, and mass of the emitted electron, and Q -value is the total available energy in the decay. From Eq. (1.7), we know that $\bar{\nu}_e$ is a superposition of three different mass eigenstates, so if we measure the energy spectrum, we would expect to observe three kinks near its endpoint. However, with the current technological capabilities, present-generation experiments do not have sufficient sensitivity to resolve such fine structure, and only a single physical observable, called m_β , which captures the overall effect, is considered:

$$m_\beta := \sqrt{\sum_{i=1,2,3} |U_{\text{PMNS}}^{ei}|^2 m_i^2} \quad (1.14)$$

As one can see, the presence of a nonzero m_β substituted for $m_{\bar{\nu}_e}$ in Eq. (1.13) has the effect of deforming the spectrum near its endpoint. In particular, the endpoint in the case of massless

³If we introduce one singlet neutrino per flavor, the mass matrix is 6×6 , with the entries in Eq. (1.12) being 3×3 .

⁴There are several effects not taken into account in this formula, such as the Coulomb interaction between the emitted electron and the daughter nucleus. These effects are included in experiments aiming to measure the neutrino mass but are irrelevant to this discussion.

neutrinos will be at Q with a finite slope, while in the case of $m_\beta > 0$ it will be at $(Q - m_\beta)$ with an infinite slope. It should be noted that, given the smallness of m_β , this is a very small effect, making this measurement extremely challenging. However, this technique has the advantage that it relies only on the conservation of energy and momentum, making it model-independent and highly reliable.

KATRIN is the world-leading experiment in direct mass measurement and performs precision spectroscopy of the β -decay spectrum of tritium to measure m_β . Its latest analysis, based on a statistics of 36 million electrons collected in 259 measurement days, resulted in an upper limit $m_\beta < 0.45$ eV at 90% confidence level [63]. It is worth noting that from oscillation experiments, we know that at least two neutrino mass states are nonzero due to the splittings, and from global fits, we have estimates for the mixing angles. It is, therefore, possible to set lower limits on m_β from oscillation data. Combining the results of global fits on oscillation data with KATRIN yields, at 95% confidence level [46]:

$$0.00085 \text{ eV} \leq m_\beta \leq 0.4 \text{ eV for NO, } 0.048 \text{ eV} \leq m_\beta \leq 0.4 \text{ eV for IO} \quad (1.15)$$

EcHO [64] and HOLMES [65] use electron capture of ^{163}Ho instead of β -decay, an idea originally proposed in 1982 [66]. The products of the process are an electron neutrino and ^{163}Dy in an excited state. Using cryogenic microcalorimeters, one measures the total de-excitation energy of ^{163}Dy , whose endpoint is at $(Q - m_\beta)$. This technique is still in the R&D phase and is not yet competitive with KATRIN.

Indirect measurements

Limits from cosmology Cosmological observations offer an indirect but powerful way to probe neutrino masses. In fact, every particle species influenced how galaxies and large-scale structures formed and evolved over time. In particular, neutrinos with nonzero masses slightly suppress the clustering of matter on small scales, because massive neutrinos move quickly enough to escape gravitational wells, smoothing out density fluctuations.

By carefully comparing the predicted distributions of galaxies and large-scale structure to observations from surveys such as DESI [67], it is possible to search for these subtle effects to set limits on the sum of neutrino masses:

$$\sum m_\nu := m_1 + m_2 + m_3. \quad (1.16)$$

Currently, no definitive evidence of neutrino mass effects has been observed, resulting in upper limits on $\sum m_\nu$ typically ranging from about 0.04 to 0.3 eV [46, 68, 69]. However, these limits rely on assumptions built into cosmological models, emphasizing the importance of direct model-independent laboratory measurements.

It is worth mentioning that neutrino oscillation data and limits from KATRIN can also be used to infer lower and upper limits for the variable $\sum m_\nu$. Combining the most recent results yields, at 95% confidence level [46]:

$$0.058 \text{ eV} \leq \sum m_\nu \leq 1.2 \text{ eV for NO, } 0.098 \text{ eV} \leq \sum m_\nu \leq 1.2 \text{ eV for IO} \quad (1.17)$$

This makes it clear that setting limits on $\sum m_\nu$ using direct detection strategies has the additional advantage that, if such results were to come into tension with predictions from cosmological limits, this could help to discriminate between different cosmological models.

Neutrinoless double-beta decay A complementary model-dependent probe is *neutrinoless double-beta decay* ($0\nu\beta\beta$), a second-order hypothetical weak nuclear process whose rate depends indirectly on the absolute neutrino mass scale. The following sections introduce the underlying physics of neutrinoless double-beta decay and explores its profound implications.

1.4 Neutrinoless Double-Beta Decay

This section introduces neutrinoless double-beta decay while Sec. 1.5 addresses the far-reaching implications of its experimental observation.

1.4.1 Double-beta decay

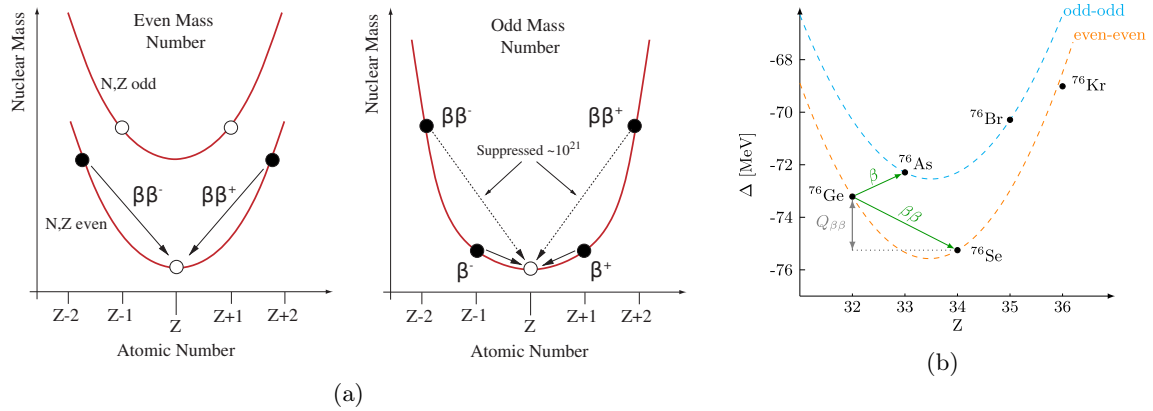


Figure 1.3: (a) Average nucleon binding energy as a function of atomic number. Odd nuclei lie on a single parabola, and due to the weak interaction structure, second-order processes like $2\nu\beta\beta$ are strongly suppressed compared to single β -decay. In contrast, spin-coupling effects split even nuclei into two parabolas: lower mass for even-even nuclei and higher for odd-odd nuclei. In such cases, single β -decay may be energetically forbidden, making the second-order transition the dominant decay mode. Reprinted from Ref. [25]; (b) Mass excess $\Delta = (m_A - A)u$ for isobars with mass m_A and mass number $A = 76$, where u is the atomic mass unit. For ^{76}Ge the β -decay is forbidden, leaving $2\nu\beta\beta$ as the leading decay channel. Reprinted from Ref. [21].

Double-beta decay ($2\nu\beta\beta$) is a second-order, weakly mediated nuclear transition of the form:



which is allowed within the Standard Model. This decay was first proposed by Maria Goeppert-Mayer in 1935, who provided an early estimate of its half-life as $T_{\frac{1}{2}}^{2\nu} \gtrsim 10^{17}$ years [70], using the Fermi theory of β -decay introduced in Sec. 1.1. Being a second-order weak process, this decay channel is highly suppressed compared to ordinary β -decay, making its observation extremely challenging. However, certain nuclei exist for which β -decay (or electron capture) is energetically forbidden. For these nuclei $2\nu\beta\beta$ becomes the dominant decay channel. This process has been observed for many

Isotope	Daughter	$Q_{\beta\beta}$ (keV)	f_{nat} (%)	Isotope	Daughter	$Q_{\beta\beta}$ (keV)	f_{nat} (%)
^{48}Ca	^{48}Ti	4267.98(32)	0.187(21)	^{116}Cd	^{116}Sn	2813.50(13)	7.512(54)
^{76}Ge	^{76}Se	2039.061(7)	7.75(12)	^{130}Te	^{130}Xe	2527.518(13)	34.08(62)
^{82}Se	^{82}Kr	2997.9(3)	8.82(15)	^{136}Xe	^{136}Ba	2457.83(37)	8.857(72)
^{96}Zr	^{96}Mo	3356.097(86)	2.80(2)	^{150}Nd	^{150}Sm	3371.38(20)	5.638(28)
^{100}Mo	^{100}Ru	3034.40(17)	9.744(65)				

Table 1.3: $Q_{\beta\beta}$ and natural abundance f_{nat} for isotopes currently being pursued by leading $0\nu\beta\beta$ -decay experiments. Isotopic abundances from [73]. Table adapted from Ref. [21].

isotopes with half-lives in the range of 10^{18} to 10^{21} yr [71].

In particular, this situation arises in even-even nuclei, like the ones listed in Table 1.3. In such nuclei, the pairing among nucleons increases the nuclear binding energy and thus reduces the nuclear mass. According to Eq. (1.1), a β -decay would convert a neutron into a proton, transforming the even-even nucleus into one with an odd number of protons and neutrons. This would reduce pairing, lower the binding energy, and raise the mass, making the β -decay energetically forbidden. In contrast, $2\nu\beta\beta$ transforms the nucleus into another even-even one, preserving pairing effects and lowering the mass, thus becoming the only allowed decay mode. This is shown in Fig. 1.3, with a numeric example for ^{76}Ge , which is one of the most studied isotope in this context.

It is important to emphasize that $2\nu\beta\beta$ is a *nuclear transition*, in which the *entire* nucleus participates, strongly influencing the process’s half-life. Theoretically predicting this half-life requires nuclear many-body models to compute the nuclear matrix elements (NME), which encode the complex nuclear dynamics and incorporate models of the nucleus and inter-nucleon correlations.

1.4.2 $0\nu\beta\beta$

Furry, in 1939, combined Majorana’s theory of neutrinos with the double-beta decay process introduced by Goeppert-Mayer and conceived an alternative version of double-beta decay in which no neutrinos are emitted in the final state:

$${}^A_Z X \longrightarrow {}^A_{Z+2} X' + 2e^- \quad (1.19)$$

This hypothetical nuclear transition became known as neutrinoless double-beta decay ($0\nu\beta\beta$) and violates both lepton number conservation and $B-L$ symmetry by two units. In 1960, the first explicit discussion of $0\nu\beta\beta$ as a lepton-number-violating process connected to the mass and Majorana nature of the neutrino was published [72]. Today, a large international experimental effort is dedicated to the search for this process, as its observation would have far-reaching implications, which are discussed in Sec. 1.5. A simplified Feynman diagram of neutrinoless double-beta decay is shown in Fig. 1.4a. As in $2\nu\beta\beta$, this is a nuclear transition involving the entire nucleus, with its half-life strongly dependent on nuclear matrix elements (NMEs). The role of NMEs, particularly in the search for $0\nu\beta\beta$, will be discussed further in Sec. 2.1.

It should be noted that if this process exists, its half-life is expected to exceed that of $2\nu\beta\beta$ by several orders of magnitude, making its detection extremely challenging. Despite its extreme rarity, $0\nu\beta\beta$ has a distinctive experimental signature: by conservation of energy and momentum, the total energy of the two emitted electrons equals the Q -value of the transition ($Q_{\beta\beta}$), producing a sharp peak in their summed energy spectrum, as illustrated in Fig. 1.4b. The experimental strategies and challenges involved in detecting this signature are the focus of Chapter 2.




Figure 1.4: (a) One possible simple diagram representing $0\nu\beta\beta$ which schematically highlights the role of nuclear physics in the transition. Reprinted from Ref. [45]; (b) Theoretical spectra of $2\nu\beta\beta$ and $0\nu\beta\beta$ decays with 1.5% energy resolution (FWHM) and arbitrary normalization. Reprinted from Ref. [21].

1.5 The Relevance of $0\nu\beta\beta$ Searches

The observation of neutrinoless double-beta decay provides a unique research opportunity, as it would simultaneously address numerous fundamental questions.

1.5.1 What distinguishes matter from antimatter?

From the theoretical formalism introduced in Sec. 1.2.3 and in Sec. 1.3.3 several considerations naturally arise.

Within the SM, the chiral (V-A) structure of weak interactions indicates that in processes such as β^+ -decay, neutrinos in the final state are strictly left-handed. For a massless neutrino (as it is in the SM) chirality and helicity coincide. Thus, SM neutrinos would always have LH helicity, clearly distinguishing it from an antineutrino emitted in β^- -decay, which would have RH helicity. Hence, in the SM helicity is as a good quantum number to distinguish neutrinos from antineutrinos.

However, the scenario changes dramatically when considering the results from neutrino oscillation experiments, which prove that neutrinos have nonzero mass: not only is the SM unable to explain these masses as it does for other fermions, but the distinction between neutrinos and antineutrinos also becomes problematic after spontaneous symmetry breaking⁵. In fact, for a massive neutrino, it is always possible to find a reference frame in which it is at rest. In this rest frame, neutrinos and antineutrinos have zero momentum, and consequently, zero helicity. In the absence of other quantum numbers, neutrinos and antineutrinos would differ solely by their spin orientation. However, rotational symmetry (or, equivalently, the freedom to arbitrarily choose the coordinate axes in the rest frame) implies that neutrinos and antineutrinos are indistinguishable. Taking the words from Ref. [21]: *the structure of the Standard Model, together with the hypothesis that neutrinos have mass, suggests that the neutrino and the antineutrino are the same particle in the rest frame*. To preserve the possibility of distinguishing neutrinos from antineutrinos, one could introduce an additional quantum number, such as lepton number L . In the rest frame, this would

⁵Hypercharge, which neutrinos possess, is broken spontaneously.

imply the existence of another distinct neutral particle characterized by the opposite value of this quantum number, thus enabling a Dirac mass term as the one in Eq. (1.10). However, this quantum number would have to be added coercively to the theory. On the other hand, one may accept that, for massive neutrinos, only a single neutrino state exists in the rest frame, making neutrinos absolutely neutral particles (with respect to all quantum numbers). Consequently, lepton number L would not be conserved. Since baryon number B is not involved in this discussion, the violation of L would imply a corresponding violation of $B - L$.

This discussion clearly demonstrates how $B - L$ is fundamentally linked to the nature of the neutrino, making experimental tests of this symmetry a top priority in contemporary particle physics. $0\nu\beta\beta$ is a process which has $\Delta B = 0$ and $\Delta L = 2$ and, having a clear experimental signature, it is a precious tool to test $B - L$.

In this respect, another important consideration must be made. In fact, a direct consequence of $B - L$ violation, is that $0\nu\beta\beta$ is a process in which energy is converted only into matter particles, without the creation of any antimatter particle. This is radically different from β or $2\nu\beta\beta$ decays, in which the net amount of particles and antiparticles remains unchanged.

1.5.2 Matter-antimatter asymmetry and a matter-creating process

Cosmological observations provide robust experimental evidence that the universe consists almost exclusively of baryons, with only an insignificant amount of antibaryons [74, 75]. This phenomenon, known as “matter-antimatter asymmetry”, has been debated for many years. In 1967 Sakharov formulated three conditions that a model must satisfy to account for this baryon excess [76]:

- violation of B;
- violation of C and CP;
- interactions taking place out of thermal equilibrium.

Given the success of the SM in describing particle physics, it is natural to examine whether it fulfills these Sakharov conditions. Such an examination is nontrivial because some of the SM’s global symmetries are anomalous (see Sec. 1.2.4). However, these anomalies are far too suppressed to generate the observed asymmetry [77], so the SM remains unable to explain the baryon excess.

Observation of a $B - L$ violating process such as $0\nu\beta\beta$ would signal physics beyond the SM capable of explaining the matter-antimatter asymmetry observed in the universe [78].

1.5.3 What can $0\nu\beta\beta$ tell us about neutrinos?

The search for $0\nu\beta\beta$ addresses the question introduced in Sec. 1.3.3: are the right-handed components of neutrinos new degrees of freedom (Dirac fermions), or are they related to the degrees of freedom of antineutrinos (Majorana fermions)? In the case of Majorana fermions, the mass term presented in Eq. (1.11) breaks the phase symmetry and consequently the number symmetry, leading to lepton-number violation and enabling processes like $0\nu\beta\beta$ to occur. Extrapolating from the argument presented in Sec. 1.5.1, in a simplified picture one could imagine that the lepton-number-violating process enabling $0\nu\beta\beta$ arises due to the nonzero neutrino mass, which implies the neutrino is not in a pure helicity state. As discussed in Sec. 1.2.2, the amount of “wrong” helicity in the neutrino field is $\mathcal{O}(m/E)$, strongly suppressing the rate of this process.

Although this is a very naïve way of looking at the process described by the Feynman diagram shown in Fig. 1.4a, this argument clearly highlights the connection among (i) the neutrino mass, (ii) the rate of $0\nu\beta\beta$, and (iii) the Majorana nature of neutrinos. Although this topic is a matter of speculation [25, 79, 80, 81, 82], the community tends to agree that the observation of a non-null rate of $0\nu\beta\beta$ would imply that neutrinos are Majorana particles. However, given the model-dependent nature of these arguments, a few points have to be made clear. Firstly, non-observation of $0\nu\beta\beta$ would not imply that neutrinos are Dirac fermions. Secondly, an observation of $0\nu\beta\beta$ could be related to the nonzero neutrino masses very indirectly. Indeed, even if an observation is made, it can't be assumed that it comes *solely* from the Majorana mass term as other processes (or even a combination of them) could be responsible for the decay. A detection of $0\nu\beta\beta$ would therefore open the much more challenging problem of experimentally testing what is the underlying mechanism mediating the process [82, 83]. After all, there are no neutrinos in the final state, and the diagram shown in Fig. 1.4a is just one of the many possible realizations of it.

With these caveats in mind, one can adopt several working assumptions, the most relevant being that the dominant mechanism for $0\nu\beta\beta$ is the exchange of light Majorana neutrinos, whose diagram is shown in Fig. 1.4a [23] (sometimes called light Majorana-neutrino exchange). Under these hypothesis one can show that the decay amplitude, and therefore the expected rate, is proportional to a quantity known as the *effective Majorana mass*, defined as [46, 21]:

$$m_{\beta\beta} := \left| \sum_{i=1}^3 |U_{ei}^2| e^{i\phi_i} m_i \right| = \left| m_1 c_{13}^2 c_{12}^2 e^{2i(\alpha_1 - \delta_{\text{CP}})} + m_2 c_{13}^2 s_{12}^2 e^{2i(\alpha_2 - \delta_{\text{CP}})} + m_3 s_{13}^2 \right| \quad (1.20)$$

Here U_{ei} are the elements of the PMNS matrix, ϕ_i are the Majorana phases, and the right hand term is written explicitly in terms of the PMNS parameters defined in Eq. (1.7). The quantities m_i are the three neutrino mass eigenvalues; this model assumes only the three standard neutrino states. Since this formula depends on the Majorana phases, a precise measurement of the $0\nu\beta\beta$ rate, combined with independent information on the absolute neutrino mass scale, could in some situations constrain or even fix one or both CP-violating Majorana phases [84].

By inserting into Eq. (1.20) the expressions for m_1 , m_2 , and m_3 from Eq. (1.9), one can rewrite $m_{\beta\beta}$ in terms of m_{lightest} and the two mass splittings Δm_{21} and $\Delta m_{3\ell}$. Doing this puts into even more clear evidence that, within this framework, the parameters space of neutrino oscillations is connected to the one of neutrinoless double-beta decay. In particular, under these assumptions, a measurement of the rate of $0\nu\beta\beta$ would fix the value of m_{lightest} , providing an indirect measurement of the neutrino mass scale, as anticipated in Sec. 1.3.4.

Quantitatively, the global analysis on the oscillation data can be combined with the bounds from KATRIN to give, at 95% CL [46]:

$$0 \leq m_{\beta\beta} \leq 0.41 \text{ eV} \quad \text{for NO}, \quad 0.015 \text{ eV} \leq m_{\beta\beta} \leq 0.41 \text{ eV} \quad \text{for IO}. \quad (1.21)$$

The limits expressed in Eq. (1.15) and Eq.(1.17), together with those in Eq. (1.21), define the allowed regions in the parameter space spanned by m_{lightest} , m_β , $m_{\beta\beta}$, and $\sum m_\nu$. The regions obtained from global fits to oscillation data are displayed in Fig. 1.5. Neutrinoless double-beta decay experiments, which set model-dependent indirect limits on $m_{\beta\beta}$, can further shrink these regions. To date, no experiment has observed $0\nu\beta\beta$. Nonetheless, measurements have set lower limits on its half-life, which, in this framework, translate into upper limits on $m_{\beta\beta}$ as shown in Fig. 2.1. This will be better discussed in Sec. 2.1.

Figure 1.5 shows that, for IO, a nonzero lower bound exists for $m_{\beta\beta}$, whereas in the NO case no such bound appears. Consequently, if a future $0\nu\beta\beta$ experiment were to establish an upper




Figure 1.5: Allowed regions obtained by projecting the results of the global analysis of oscillation data over $m_{\beta\beta}$ and (a) m_β as in Eq. (1.14); (b) Σm_ν as in Eq. (1.16); (c) m_{lightest} as in Eq. (1.9). Colored areas (grey lines) are obtained with (without) the inclusion of the KATRIN bound. Upper bounds on $m_{\beta\beta}$ shrink the allowed regions horizontally from the top of these plot. Upper bounds on Σm_ν from cosmology can also shrink the allowed regions. Adapted from [46].

limit on $m_{\beta\beta}$ below the IO-allowed band, it would (indirectly) demonstrate normal ordering even in the absence of a detection of $0\nu\beta\beta$. It is worth noting, however, that forthcoming oscillation experiments like JUNO are expected to determine the mass ordering independently in the coming years [58]. In addition, if neutrinos follow NO, specific combinations of the Majorana phases could suppress $m_{\beta\beta}$ to extremely small values, potentially making $0\nu\beta\beta$ practically impossible to detect. Certain theoretical frameworks predict values for these parameters [85]; therefore, systematically excluding larger portions of this space provides valuable guidance for model building even prior to the experimental observation of $0\nu\beta\beta$. Taken together, these considerations demonstrate the exceptional informative power provided by the experimental searches for this process.

Given the extremely small values possible for $m_{\beta\beta}$, the amplitude of $0\nu\beta\beta$ is expected to be tiny. For this reason, one might look at other processes to test the nature of the neutrino, such as interactions involving non-relativistic neutrinos. However, the community largely agrees that, for now, only searches for $0\nu\beta\beta$ are experimentally practical. This has also been shown in very recent studies [86].

In summary, detecting a matter-creating process such as $0\nu\beta\beta$ would greatly deepen our understanding of neutrinos, physics beyond the Standard Model, and why the Universe is formed by matter. Its observation would prove that $B - L$ is violated and, under some assumptions, offer information on whether neutrinos are Majorana particles, on the absolute neutrino mass scale and ordering, and on the CP-violating Majorana phases. $0\nu\beta\beta$ is the only practical laboratory probe capable of addressing these foundational questions simultaneously. For these reasons, the search for $0\nu\beta\beta$ remains one of the highest worldwide priorities in particle physics today [87, 88, 89].

Chapter 2

Searching for $0\nu\beta\beta$ with Cryogenic Calorimeters: from CUORE to CUPID

The far-reaching implications of discovering neutrinoless double-beta decay ($0\nu\beta\beta$) have motivated a broad international experimental programme. Present-generation experiments, based on several detector technologies, already set stringent lower limits on the half-life of many candidate isotopes, and the next-generation ones are designed to improve those limits by at least an order of magnitude.

Among the available techniques, the one based on cryogenic calorimeters has emerged as especially promising. The CUORE experiment currently holds the most stringent half-life limit for ^{130}Te . Its multi-year operation of a tonne-scale array at ~ 15 mK has demonstrated the long-term stability, excellent energy resolution, and low backgrounds attainable with this method.

CUPID will extend this approach. It targets the $0\nu\beta\beta$ decay of ^{100}Mo using scintillating cryogenic Li_2MoO_4 crystals with a dual-channel readout that will allow it to have CUORE-level energy resolution while providing event-by-event particle identification, reducing backgrounds in the region of interest. In addition, ^{100}Mo offers a Q -value above the natural radioactivity end-point and benefits from favourable nuclear phase space factor, enhancing sensitivity to the effective Majorana mass.

This chapter reviews the current status of $0\nu\beta\beta$ searches, explains the advantages and challenges of scintillating cryogenic calorimeters, and contextualizes CUPID within a broader experimental program.

2.1 Experimental Sensitivity to the Effective Majorana Mass

Assuming that $0\nu\beta\beta$ is driven by the light–Majorana-neutrino exchange mechanism introduced in Sec. 1.5.3, the decay rate is proportional to the effective Majorana mass $m_{\beta\beta}$ defined in Eq. (1.20) [25]. Consequently, for a given half-life one can compute the corresponding value of $m_{\beta\beta}$. Although this conversion is model-dependent, it enables a direct comparison of the discovery reach of exper-

iments using different isotopes. As recalled in Sec. 1.4, neutrinoless double-beta decay is a nuclear transition, so the entire nucleus influences the rate, which can be written as follows [25]:

$$\frac{1}{T_{1/2}^{0\nu}} = G^{0\nu} |\mathcal{M}|^2 \frac{m_{\beta\beta}^2}{m_e^2} \quad (2.1)$$

where $G^{0\nu}$ is the phase-space factor, \mathcal{M} the nuclear matrix element (NME), and m_e the electron mass. Hence, a measurement of $T_{1/2}^{0\nu}$ can be converted (indirectly) into a value of $m_{\beta\beta}$. Conversely, if $0\nu\beta\beta$ is not observed, the experimental lower limit on the half-life translates into an upper bound on $m_{\beta\beta}$ according to [25]:

$$m_{\beta\beta} \leq \frac{m_e}{\mathcal{M} \sqrt{G^{0\nu} T_{1/2}^{0\nu}}} \quad (2.2)$$

Equation (2.2) shows that both the NME and the phase-space factor are crucial. Although $G^{0\nu}$ can vary a lot among the candidate isotopes (listed in Tab. 1.3) [90], its uncertainty is only about 7% [91]. NMEs, on the other hand, are difficult to calculate accurately and represent the largest source of uncertainty when deriving bounds on $m_{\beta\beta}$.

Nuclear matrix elements

Because $0\nu\beta\beta$ and $2\nu\beta\beta$ occur in nuclei with many protons and neutrons, calculating NMEs requires solving a demanding many-body problem. Current theoretical approaches differ by up to a factor of three, indicating that important underlying physics is still missing in these models [95]. Moreover, reliably estimating the associated uncertainties is itself challenging. A common practice is, therefore, to evaluate Eq. (2.2) using the largest and smallest \mathcal{M} values available, thereby giving

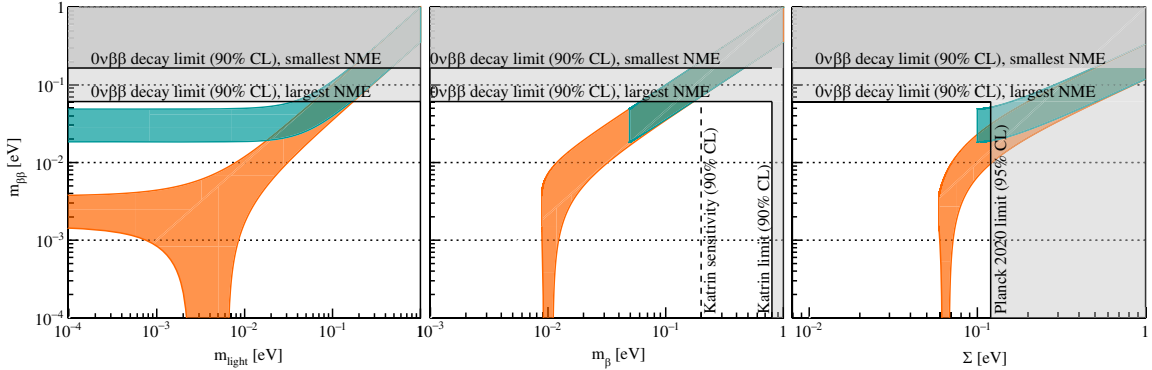


Figure 2.1: Colored bands show the maximally allowed parameter space for $m_{\beta\beta}$ as a function of m_{lightest} as defined in Eq. (1.9), m_{β} as in Eq. (1.14), and Σm_{ν} as in Eq. (1.16) assuming the central value of the neutrino oscillation parameters in IO (teal) and NO (orange). Vertical lines in the middle panel mark the regions excluded by the KATRIN limit (gray area) [92] and sensitivity [93], while the vertical gray band in the right panel show cosmological constraints [74]. The horizontal gray band denotes the region ruled out by the 2016 KamLAND-Zen $0\nu\beta\beta$ data assuming the smallest and the largest NME [94]. Reprinted from Ref. [21].

a plausible range for $m_{\beta\beta}$.

With this procedure, one can map existing half-life limits onto the allowed regions shown in Fig. 1.5 to further shrink the permitted values by introducing exclusion regions as shown in Fig. 2.1. Although model-dependent, these bounds already surpass those derived from global fits on oscillation data combined with KATRIN results (reported in Eq. (1.21)). For instance, the current best limit, $T_{1/2}^{0\nu} > 3.8 \times 10^{26}$ yr (90% C.L.) from KamLAND-Zen [96], corresponds to $m_{\beta\beta} < 28\text{--}122$ meV, even tighter than the range shown in Fig. 2.1, which was obtained with the 2016 data [94]. Even for the worse (smallest) NME, the resulting upper limit of 0.122 eV is more than three times stronger than the global-fit bound of 0.41 eV quoted in Eq. (1.21).

2.2 Experimental Search for $0\nu\beta\beta$

Equation (2.1) provides a model-dependent connection between the half-life of neutrinoless double-beta decay, $T_{1/2}^{0\nu}$, and the effective Majorana mass $m_{\beta\beta}$. The real experimental challenge is to measure $T_{1/2}^{0\nu}$ directly or, in the absence of a signal, to set a lower bound on it. Although the decay is expected to be exceptionally rare, its signature is very clear: the sum of energies of the two emitted electrons must equal the Q -value of the transition $Q_{\beta\beta}$, as shown in Fig. 1.4b. In practice, every experimental approach that one could think of boils down to looking for a narrow, monochromatic peak at $Q_{\beta\beta}$ in the total electron-energy spectrum.

Many detector technologies have been developed to achieve this goal, yet their sensitivities can all be expressed through the same core experimental quantities. This section introduces these shared parameters and shows how they determine the achievable limits (or, eventually, the first measurement) of $T_{1/2}^{0\nu}$.

2.2.1 Experimental sensitivity to $0\nu\beta\beta$

From the law of radioactive decay it is easy to show that, if $N_{\beta\beta}$ nuclei are observed for a total live time T with a detector whose total efficiency for the two-electron signal is ϵ , then, in the lucky case $n_{\beta\beta}$ events are identified as $0\nu\beta\beta$, the half-life is [25]

$$T_{1/2}^{0\nu} = \ln 2 \cdot T \cdot \epsilon \cdot \frac{N_{\beta\beta}}{n_{\beta\beta}} \quad (2.3)$$

If every quantity except $n_{\beta\beta}$ is known with negligible uncertainty, the only error on $T_{1/2}^{0\nu}$ arises from the Poisson fluctuations of the event count.

When no $0\nu\beta\beta$ candidates are observed, it is convenient to define a *detector factor of merit*, or *sensitivity* $S^{0\nu}$. This is the half-life corresponding to the largest signal that could be hidden by background fluctuations at a chosen confidence level (CL). Despite several simple approximations, the most important being that the background counts grows linearly with time and detector mass, the sensitivity can be used to compare present and future experiments on the basis of common experimental parameters [90]. At 68% C.L. the sensitivity is [97]:

$$S^{0\nu} = T_{1/2}^{\text{Back.Fluct.}} = \ln 2 N_{\beta\beta} \epsilon \frac{T}{n_B} = \ln 2 \times \frac{x \eta \epsilon N_A}{A} \sqrt{\frac{MT}{B \Delta}} \quad (n_B = \sqrt{MTB \Delta} \text{ at 68\% C.L.}) \quad (2.4)$$

Where, for an experiment that studies a $\beta\beta$ isotope embedded in a material of molecular mass A containing x atoms of that isotope per molecule,

- $N_{\beta\beta}$, ϵ , and T are as defined in Eq. (2.3); N_A is Avogadro's number;
- B is the *background index* (BI), i.e. the background rate per unit energy and detector mass in the region near $Q_{\beta\beta}$, commonly called the *region of interest* (ROI);
- M is the detector mass;
- Δ is the FWHM energy resolution in the ROI;
- η is the isotopic abundance of the $\beta\beta$ nuclei under study, equal to the natural abundance f_{nat} of Table 1.3 or higher if isotopic enrichment strategies are used.

Equation (2.4) shows that, for a fixed mass and observation time, a detector with lower background, better energy resolution, or higher enrichment reaches an enhanced sensitivity. Defining the *exposure* MT , the sensitivity scales as the square root of the exposure in the finite-background regime.

If an experiment operates *background-free* in the ROI, namely $MT\Delta B \lesssim 1$, then, at 68% C.L., its sensitivity becomes [97]:

$$S_{\text{ZB}}^{0\nu} = \ln 2 N_{\beta\beta} \epsilon \frac{T}{n_B} = \ln 2 \times \frac{x \eta \epsilon N_A}{A} \frac{MT}{1.14} \quad (n_B = 1.14 \text{ at } 68\% \text{ C.L.}) \quad (2.5)$$

In this *zero-background* (ZB) regime the sensitivity increases linearly with exposure and no longer depends explicitly on Δ , although maintaining good resolution remains essential, since the experiment re-enters the finite-background regime once $MT\Delta B$ exceeds unity.

Finally, combining Eqs. (2.4) and (2.5) with Eq. (2.2) gives the dependence of the effective Majorana-mass reach on exposure:

$$m_{\beta\beta} \leq \begin{cases} \propto (MT)^{-1/2} & \text{zero background,} \\ \propto (MT)^{-1/4} & \text{finite background.} \end{cases}$$

Thus, in the finite-background case, improving the $m_{\beta\beta}$ sensitivity by a factor of two requires increasing the exposure by a factor of 16. Minimizing backgrounds is therefore crucial, as every reduction moves an experiment closer to the zero-background regime, where sensitivity grows much faster with exposure and discovery potential is maximized. Not only it is crucial to keep the backgrounds as low as possible, but also to measure it with high precision, ensuring that background-related uncertainties are completely negligible, as they can translate into an uncertainty on the half-life.

2.2.2 Current and future experimental approaches

Current- and next-generation experiments employ multiple complementary technologies to search for $0\nu\beta\beta$ decay.

Experiments using different isotopes and detector technologies are actively pursued to maximize the discovery potential for $0\nu\beta\beta$ decay. Prominent technologies include semiconductor detectors (HPGe), liquid xenon TPCs, high-pressure xenon gas TPCs, large-scale liquid-scintillator detectors, and cryogenic calorimeters [21]. HPGe diodes typically provide excellent energy resolution and proven scalability but limit the search of the process to ^{76}Ge . Xenon TPCs benefit from high-purity operation and effective event reconstruction capabilities, whereas large-scale liquid scintillator

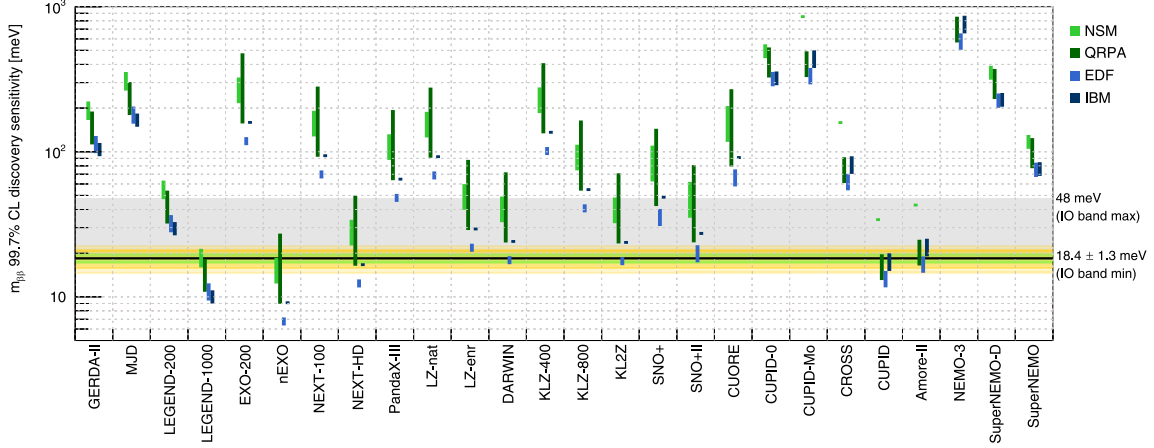


Figure 2.2: Discovery sensitivities of current- and next-generation $0\nu\beta\beta$ -decay experiments under the light-neutrino-exchange hypothesis. The coloured vertical bars illustrate the spread that arises when different NMEs are used, as discussed in Sec. 2.1. Values of $m_{\beta\beta}$ above the indicated thresholds are probed at higher confidence levels. The grey band marks the range of $m_{\beta\beta}$ expected for an inverted mass ordering with a vanishing lightest mass, while the minimum IO value and its 1σ , 2σ , and 3σ uncertainty intervals are shown by the black, green, orange, and yellow bands, respectively. Adapted from Ref. [21].

detectors leverage their significant mass but pay the price of a lower energy resolution. Cryogenic calorimeters, especially when coupled with scintillation readouts (scintillating cryogenic calorimeters), offer excellent energy resolution and particle identification capabilities, enabling highly effective background suppression. Each of these technologies involves distinct challenges and opportunities regarding isotope enrichment, radiopurity, detector scalability, and operational complexity [21]. Figure 2.2 shows the discovery sensitivities of current- and next-generation $0\nu\beta\beta$ -decay experiments. The vertical colored bars illustrate the significant spread resulting from different nuclear matrix element calculations (discussed in Sec. 2.1). Scintillating cryogenic calorimeters, in particular, aim at achieving sensitivity deep into the inverted-ordering neutrino-mass regime, highlighting their prominent role within the future landscape of $0\nu\beta\beta$ -decay searches.

The most stringent constraint to date is set by KamLAND-Zen, which reports a limit of $T_{1/2}^{0\nu} > 3.8 \times 10^{26}$ yr (90% C.L.) [96]. This bound translates into an effective Majorana mass constraint of $m_{\beta\beta} < 28\text{--}122$ meV, depending on the adopted nuclear matrix element.

2.3 Cryogenic Calorimeters for $0\nu\beta\beta$ Searches

Cryogenic calorimeters are devices that provide all the requirements expressed by Eq. (2.4). This section outlines the working principle of this technology explaining how it can be effectively used to set competitive limits on $0\nu\beta\beta$ and $m_{\beta\beta}$ as it is done by the CUORE experiment. Scintillating cryogenic calorimeters are devices that benefit of all the same advantages demonstrated by CUORE with the added benefit of being able to further reduce backgrounds in the ROI by exploiting

scintillation light to perform particle identification. These devices will be introduced in the following section.

2.3.1 Cryogenic calorimeters

Over a century ago, Marie and Pierre Curie, together with Albert Laborde, showed that radium’s radioactive decay releases measurable heat [98, 99], establishing that ionizing radiation inevitably warms any material it traverses. At ordinary temperatures, however, the temperature increase from a single particle’s interaction is far too small to register. In 1984, E. Fiorini and T. O. Niinikoski recognized¹ that this difficulty could be overcome by cooling the absorber until its heat capacity became exceptionally low; under such conditions, even the minuscule heat released by one single particle could be detected [102]. Their insight launched the field of low-temperature detectors and opened a new avenue for excellent energy resolution spectrometry.

Working principle

A cryogenic calorimeter consists of a dielectric and diamagnetic absorber crystal instrumented with a thermal (phonon) sensor. Figure 2.3a presents an idealised schematic of such device. In practice the crystal is mechanically supported, thermally linked to the bath, and equipped with sensors and heaters, elements whose purpose will be detailed later; Fig. 2.3b shows the CUORE implementation, typical of cryogenic setups yet not the only possible realisation. Ideally the system behaves as a perfect calorimeter: when a particle deposits an energy E , the absorber temperature increases by

$$\Delta T = \frac{E}{C}, \quad (2.6)$$

where C is the absorber’s heat capacity. An accurate measurement of ΔT with the thermal sensor therefore yields the deposited energy E . For the temperature rise to be measurable with high precision, ΔT must be sufficiently large; hence the absorber should have a very small heat capacity. This requirement is met by operating dielectric or diamagnetic crystals at temperatures of 10–100 mK. In the Debye regime [105], the heat capacity of such materials is

$$C(T) \simeq \frac{12\pi^4}{5} N k_B \left(\frac{T}{\theta_D} \right)^3 = 1944 \frac{m}{M} \left(\frac{T}{\theta_D} \right)^3 \text{ J K}^{-1} \quad (T \ll \theta_D), \quad (2.7)$$

where θ_D is the Debye temperature, N the number of atoms, m the absorber mass, M the molar mass, and k_B Boltzmann’s constant. Cooling the crystal to ~ 10 mK reduces C to a level at which energy deposits from a few hundreds of eV up to several MeV yield detectable temperature pulses.

Equation (2.7) shows a key strength of cryogenic calorimeters: as long as the absorber possesses a sufficiently high θ_D , its exact chemical composition is not critical. This gives wide freedom in choosing the absorber crystal, making these detector extremely versatile in many fields of particle physics². For $0\nu\beta\beta$ searches, any isotope in Table 1.3 can, in principle, be used if a stable, good-quality radiopure crystal containing that isotope can be produced. The same freedom allows to

¹Early hints that cryogenic calorimetry could be used as a particle detector appeared in 1974, even earlier than the idea published by Fiorini, when T. O. Niinikoski and F. Udo suggested measuring the minute heating produced by a neutrino beam traversing matter [100]. A second milestone was a 1982 communication by G. V. Mikelmacher, B. S. Neganov, and V. N. Trofimov, which outlined how emerging cryogenic techniques might register extremely small energy releases through the resulting temperature rise [101].

²For example, scintillating cryogenic calorimeters have been proposed to detect astrophysical neutrinos [106] and to search for dark matter [107, 108]. A review of suitable materials and their uses is given in Ref. [109].

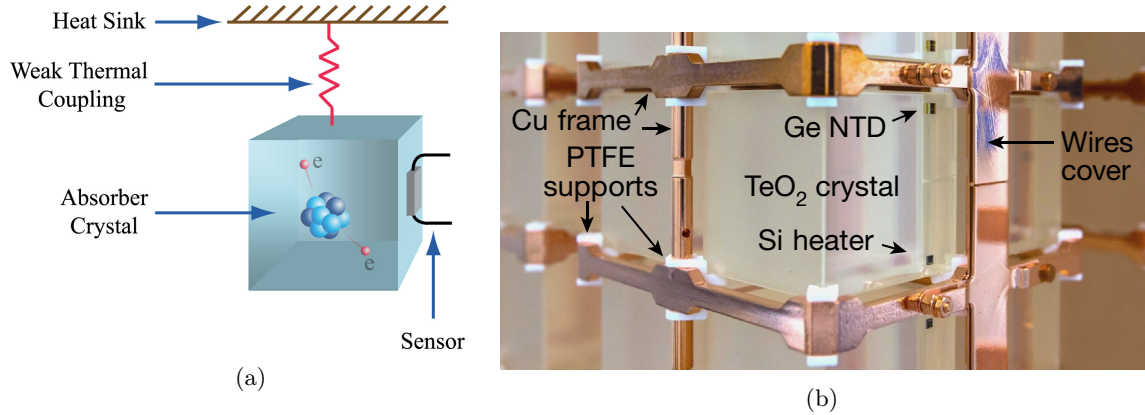


Figure 2.3: (a) Idealised sketch of a cryogenic calorimeter: a $0\nu\beta\beta$ decay occurs within the absorber crystal, and the two emitted electrons deposit their energy as heat, which is sensed by the thermal sensor. Reprinted from Ref. [103]. (b) Real-world implementation: the crystal (TeO_2) is mechanically supported (Cu frames and PTFE supports) and instrumented with thermistors (Ge NTD), heaters (Si heater), and read-out wiring; the photograph shows a CUORE detector module. Reprinted from Ref. [104].

pick compounds that provide secondary detection channels; for example, if a scintillating material is chosen, the light emitted by the absorber crystal can be used to perform particle identification and reduce backgrounds in the ROI. This idea will be introduced in Sec. 2.4.1 and discussed further in later chapters. Table 2.1 lists Debye temperatures for the absorber materials considered in this thesis, together with their corresponding heat capacities and expected ΔT values.

Material	Θ_D [K]	M $\left[\frac{\text{g}}{\text{mol}}\right]$	C $\left[\frac{\text{MeV}}{\text{K}}\right]$	ΔT $\left[\frac{\text{mK}}{\text{MeV}}\right]$
TeO_2	232 ± 7 [110]	159.6	1217.7	~ 0.8
Li_2MoO_4	316.15 [111]	173.82	441.8	~ 2.2

Table 2.1: Values for Debye temperature, heat capacity and temperature increase for 1 MeV of deposited energy in a 200g crystal typically used for $0\nu\beta\beta$ searches kept at 10 mK. The values of the thermal capacities reported in this table are evaluated using Eq. (2.7) under the approximation of considering only the lattice contributions (Debye law).

Converting the energy deposited by a single particle into a measurable heat pulse

Equation (2.6) describes an ideal, perfectly isolated calorimeter subjected to an instantaneous energy deposit. In a practical detector, such as the one shown in Fig. 2.3b, the absorber is weakly coupled to a thermal bath maintained at the base temperature $T_b \sim 10$ mK. After a particle interaction, the absorber warms up according to Eq. (2.6) and then relaxes back to T_b on a time scale dependent on its heat capacity C and the thermal conductance G of the link to the bath. The resulting signal is a transient temperature pulse (see Fig. 2.4a) whose amplitude is proportional to




Figure 2.4: (a) Block diagram of a typical thermal model for a detector configuration similar to that in Fig. 2.3b. E_0 is the energy deposited by particle interactions, G_i are thermal conductances, C is the thermal capacity of the crystal and T_b is the heat sink temperature. This simplified network specifically describes CUORE and is reprinted from Ref. [112]. (b) Typical pulse from a cryogenic calorimeter employing a Ge-NTD sensor; the example shown is taken from a CUORE TeO_2 crystal, and its time profile can be naïvely modeled by the thermal network in panel (a); the pulse height is proportional to the energy E_0 deposited in the crystal.

the deposited energy and inversely proportional to C , while its decay time scales as C/G . Thus a small heat capacity and a weak thermal link guarantee that the full energy release is converted into the amplitude of the pulse. A small heat capacity also maximises the pulse amplitude for a given energy release. An accurate measurement of the amplitudes of the pulses provides excellent energy resolution spectroscopy.

The discussion so far has been deliberately naïve, treating the detector as a single heat-capacity element weakly coupled to a heat sink. A more realistic picture must include every physical component of the detector and the conductances linking them: absorber crystal, thermistor, copper frames, polytetrafluoroethylene (PTFE) supports, and their interfaces, readout wiring, and so on. Figure 2.4b shows such a network for the CUORE module shown in Fig. 2.3b; although still idealised, it already captures the dominant heat-flow paths. Nonetheless, even this extended network still remains a simplified representation. A complete model must track all significant heat capacities and conductances, embed the thermistor bias and read-out circuitry, and incorporate the temperature dependence of both $C(T)$ and the $G_i(T)$, whose variation introduces non-linearities into the time evolution of the pulse and, possibly, on its amplitude [112].

A read-out electronic chain is also required to transform the thermistor’s temperature variation into the voltage pulse shown in Fig. 2.4a. This electronics not only enables signal acquisition but also couples back to the sensor, so the detector must be described with a more comprehensive electro-thermal model [113, 114] that treats the detector (absorber and thermistor) and the read-out electronics as a whole.

A typical CUORE thermal pulse height is about $100 \mu\text{K}/\text{MeV}$ [115] (close to the order-of-magnitude value estimate reported in Table 2.1, naïvely derived from Eq. (2.7)). The thermistor

converts this thermal pulse into a voltage pulse that has a typical amplitude of $100 \mu\text{V}/\text{MeV}$ [116]. This signal is then amplified by the front-end electronics [117] to give a pulse similar to the one shown in Fig. 2.4a. Typical CUORE pulses have rise and decay times of roughly 0.1 s and 1 s, respectively [115] and a frequency bandwidth extending up to 20 Hz [116].

Read-out electronics: converting a $\mathcal{O}(100 \mu\text{K})$ heat pulse into a measurable signal

The previous sections illustrated, using a CUORE module as a reference, how a thermal pulse arises in a cryogenic *macro*calorimeter. The same principles applies to other cryogenic detector classes like microcalorimeters, which chiefly differ in absorber mass, thermal coupling, and in the thermal sensor chosen to read out the pulse. The discussion that follows remains focused on macrocalorimeters, which are suitable detectors to perform competitive $0\nu\beta\beta$ searches.

Thermal sensor The sensor type largely sets the read-out scheme and pulse characteristics. The most common types of sensors include: Transition-Edge Sensors (TES) [118], Metallic Magnetic Calorimeters (MMC) [119], Kinetic Inductance Detectors (KID) [120], and Neutron-Transmutation-Doped Si or Ge wafers (NTD). These sensors present different features that makes them suitable for different applications. The choice of the sensor therefore depends on the experiment’s energy range, event rate, and construction constraints. For example, TES sensors require a complex micro-fabrication sequence and are normally deposited directly on the absorber; if the crystal cannot withstand those steps, the TES must be built on a separate support that is then thermally coupled to the absorber [121], adding complexity to both the thermal model (and the understanding of the detector’s response) and the assembly of experiments based on large detector arrays. TES and MMC are low-impedance ($\leq 0.1 \Omega$) devices and therefore need dedicated cryogenic based read-out electronics. NTD thermistors, by contrast, have high impedance, in the range $\mathcal{O}(0.1\text{--}1 \text{ G}\Omega)$ at $T_b \sim 15 \text{ mK}$, and can be read with simpler room-temperature electronics described in the following. The trade-off is slower response and a slightly poorer energy threshold. All the measurements discussed in this work are performed using Ge-NTDs, which are the same type of sensor used for the CUORE and CUPID experiments (introduced in Sec. 2.3.3 and Sec. 2.4). For this reason, in the following discussion, only NTD sensors are considered.

Neutron-Transmutation-Doped Si or Ge wafers (NTD) A Si or Ge crystal behaves as an insulator at millikelvin temperatures. Neutron-transmutation doping (performed in nuclear reactors) introduces a very uniform distribution impurities. There exists a critical dopant concentration N_c for which a transition from insulator to conductor occurs. For dopant concentration $N \simeq N_c$, i.e. near the transition regime, the material resistivity strongly depends on the temperature. At low temperatures conduction occurs through electrons transitioning between impurity sites without passing through the conduction band via the *hopping mechanism* [122]. When temperatures are extremely low and the doping concentration is slightly below a critical density N_c the device is described by variable-range hopping [123], and its resistivity can be described by the following law:

$$\rho(T) = \rho_0 e^{\sqrt{\tau_0/T}}$$

For a sensor of length l and cross section A

$$R(T) = \rho(T) \frac{l}{A} = R_0 e^{\sqrt{\tau_0/T}} \quad (2.8)$$

where R_0 is a factor that depends on the dopant concentration and the device geometry. Eq. (2.8) shows that the device's resistance is a very steep function of the temperature, nearly exponential, and is able to convert a $\mathcal{O}(10\text{--}100\ \mu\text{K})$ pulse into a big resistance change, which can be measured as a voltage drop if the sensor is biased with a constant current. If compared to MMC and TES, the NTD is intrinsically slower, with response times of the order of $\mathcal{O}(\text{ms})$.

Biasing and read-out circuit When a constant current I flows through the NTD, any variation in its resistance produces a corresponding change in the voltage across the device:

$$V_{\text{NTD}}(T) = IR_{\text{NTD}}(T) \quad (2.9)$$

with T denoting the temperature and $R_{\text{NTD}}(T)$ is given by Eq. (2.8) [124]. A convenient way to exploit this relation is to employ a two-wire circuit that simultaneously biases and reads out the thermistor. Driving the series combination of the NTD and a *load resistor* $R_L \gg R_{\text{NTD}}$ with a low-noise voltage power supply fixes the current I (limited by R_L), keeping it nearly constant for small variations of R_{NTD} . The corresponding scheme is depicted in Fig. 2.5. A temperature variation,




Figure 2.5: Biasing and readout circuit for an NTD thermistor. The load resistor R_L is selected so that $R_L \gg R_{\text{NTD}}$, keeping the current nearly constant at $I \simeq V_{\text{bias}}/R_L$ even when R_{NTD} varies. Reprinted from Ref. [124].

such as the thermal pulse generated by a particle interaction, first alters the resistance via Eq. (2.8) and then is converted into a time-dependent voltage signal measured across the thermistor according to Eq. (2.9). The time evolution of the waveform $V(t)$ is shaped not only by this conversion but also by the transfer functions of the readout circuitry and by several noise sources acting in different frequency bands. These noise contributions are examined in the next section.

Electro-thermal feedback and detector's working point Equation (2.9) implies that, for a given temperature change, the voltage signal scales with the bias current. Raising I therefore looks like a trivial way to amplify amplitude and, if there are noise sources that do not grow concurrently, the signal-to-noise ratio (SNR) also increases, potentially improving the energy resolution. This strategy, however, works only up to a certain biasing current because the NTD always dissipates Joule power according to:

$$P(T) = I^2 R_{\text{NTD}}(T) \quad (2.10)$$

which feeds back on the detector response. If the dissipated power becomes comparable with the thermal conductance G linking the thermistor to the heat sink, the NTD warms above the bath temperature T_b [124]:

$$T = T_b + \frac{P}{G}$$

A higher T reduces the NTD resistance and, consequently, the signal amplitude.

This phenomenon is called *electro-thermal feedback*: for $P \ll G$, the NTD temperature tracks the bath closely ($T \approx T_0$), and the signal amplitude initially grows linearly with I through Eq. (2.9). Once the current exceeds a critical value (called *inversion*), self-heating raises T , the resistance R_{NTD} decreases, and, since according to Eq. (2.9) the signal amplitude is proportional to R_{NTD} , the height of the voltage pulses also begins to decline. At the same time, Eq. (2.10) shows that the dissipated power decreases with R_{NTD} , giving rise to a non-trivial interplay between current, self-heating, and resistance.

The effects of electro-thermal feedback can be seen in the $I-V$ characteristic as well as in the $R-P$ curve. Figure 2.6a displays a representative $I-V$ load curve for an NTD thermistor mounted on an absorber, together with the amplitude of a monochromatic thermal pulses injected in the absorber plotted as a function of V_{NTD} . The derivative of the inverse of the $I-V$ curve is called dynamic resistance ($R_{\text{dyn}} = dV/dI$). For small bias currents the response is linear: as I increases, V_{NTD} rises proportionally and the device behaves like a conventional resistor ($R_{\text{NTD}} = R_{\text{dyn}}$); this operational domain is referred to as the *ohmic region*. When the current becomes large enough to heat the




Figure 2.6: (a) Blue line: the $I-V$ characteristic of the same NTD. At low bias the device behaves linearly (ohmic); the point where the slope diverges is the inversion. Red points: amplitude of monochromatic pulses injected in the calorimeter as a function of V_{NTD} ; after the inversion, the signal drops as consequence of the reduction of R_{NTD} . Both curves clearly show the effects caused by electro-thermal feedback. (b) $R-P$ curve for CUORE NTDs at heat sink temperatures T_b from 11 mK to 27 mK; when the current increases, the power dissipated by the NTD (horizontal axis) start to have a non-negligible self heating effect, reducing R_{NTD} (vertical axis). The value of the resistance in the low power limit is called R_{base} . Both adapted from Ref. [124].

thermistor appreciably, its resistance drops and V_{NTD} starts to decrease. The point where the slope of the I - V curve diverges ($R_{\text{dyn}} \rightarrow 0$) is called the *inversion*. If the bias current is increased further, the detector moves into the *overbias* regime. Beyond the inversion, the concurrent reductions in V_{NTD} and R_{NTD} cause the signal amplitude to decline, clearly showing the influence of electro-thermal feedback on detector performance.

The value of the biasing voltage V_{bias} and the load resistor R_L fix a biasing current I which sets the NTD in one of the operating regions described above. The choice of V_{bias} and R_L for a given heat sink temperature T_b is called Working Point (WP).

Figure 2.6b shows the resistance R_{NTD} as a function of the dissipated power P for a CUORE NTD thermistor ($R_0 \approx 1.0 - 1.5 \Omega$, $T_0 \approx 4.0 - 5.0 \text{ K}$) measured at several bath temperatures T_b [124]. In the low-power limit (ohmic region) the thermistor is in thermal equilibrium with the bath, so $T \simeq T_b$. For a given NTD (with given R_0 , T_0), the corresponding resistance, called the *base resistance* R_{base} , depends only on T_b through Eq. (2.8). R_{base} can therefore serve as a proxy for the heat sink temperature. Increasing the bias current raises the power, which heats the thermistor and drives its resistance downward. Each working point (which could be chosen in the ohmic region, near the inversion, or in overbias), has a corresponding NTD resistance. The value of R_{NTD} at the WP is called *working resistance* and is indicated with R_{WP} .

The preceding discussion makes it evident that the choice of the WP strongly affects the signal amplitude, with Fig. 2.6a indicating that the highest gain occurs close to the inversion. The figure of merit that ultimately governs both energy resolution and threshold, however, is the Signal-to-Noise Ratio (SNR), and not just the signal amplitude. Because the frequency response of the read-out chain changes in a nontrivial way across different WPs, each noise source is filtered differently, and their combined impact on the SNR cannot be inferred from the signal amplitude alone. For this reason, it is always important to perform dedicated working point optimization measurements when operating cryogenic calorimeters.

Signal amplification and digitization Typical pulses obtained with the circuit of Fig. 2.5 exhibit amplitudes of order $\mathcal{O}(10\text{-}100 \mu\text{V}/\text{keV})$. This quantity, which expresses the pulse amplitude per unit deposited energy and is usually reported in $\mu\text{V}/\text{keV}$ or nV/keV , will be referred to as *intrinsic gain* or *sensitivity*³ in the following sections. Signals this small are readily obscured by noise; the main noise sources will be discussed in the next section. To keep their impact negligible, the signal must be amplified before it meets any subsequent noise-producing stage.

Figure 2.7 shows a simplified front-end chain for an NTD-based cryogenic calorimeter [117]. The thermistor voltage is carried to room temperature through cryostat wiring and enters a low-noise differential voltage pre-amplifier with a selected pair of Junction Field Effect Transistors (JFETs) at the inputs. A secondary stage, realised with a programmable-gain amplifier (PGA), provides additional adjustable gain. After amplification, pulse amplitudes of order $\mathcal{O}(0.1\text{-}1 \text{ mV}/\text{keV})$ are obtained. The amplified signal then passes through an anti-aliasing low-pass Bessel filter and is digitised by the analog-to-digital converter (ADC) boards of the data-acquisition (DAQ) system, which record the waveforms for offline analysis. The ADCs sampling frequency is typically chosen between 1 and 4 kHz. All the measurements described in Chapter 4 have been acquired with an ADC having 18-bit resolution while the ones described in Chapter 6 have been acquired with 24-bit ADC resolution. The PGA gain is chosen so that the quantisation noise is negligible compared with the intrinsic detector noise. Except for the anti-aliasing filter, all further shaping and filtering

³Not to be confused with the sensitivity $S^{0\nu}$ introduced in Sec. 2.2.1.

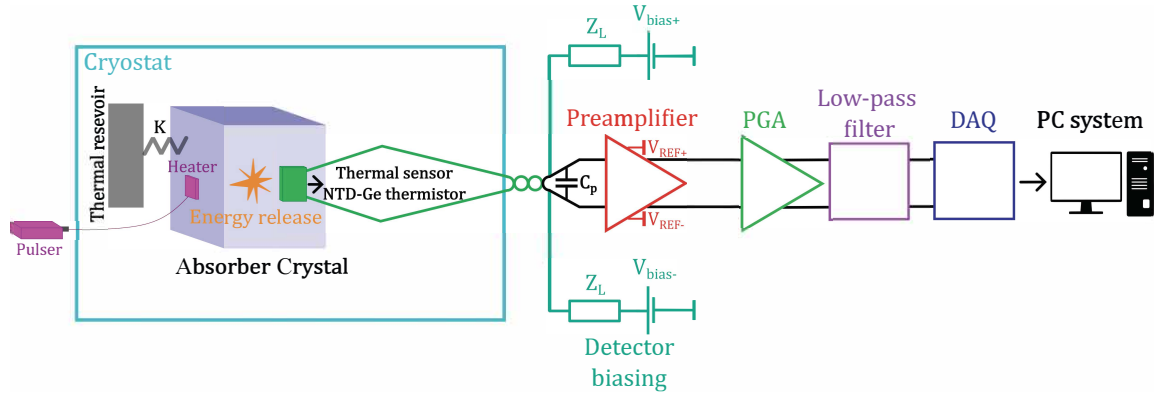


Figure 2.7: Typical front-end (FE) and back-end electronics of NTD equipped cryogenic calorimeters. Z_L is $\approx R_L/2$. C_p denotes the parasitic capacitance of the leads that connect the room-temperature front-end back-plane to the NTD thermistor inside the crystal; this capacitance often has non-negligible effects and modifies the overall transfer function, effectively introducing an initial shaping stage in the signal chain. The pulser and the heater may be used to inject monochromatic pulses that emulate energy release of particles and can be used for many purposes, including building plots like the one shown in Fig. 2.6a. Adapted from Ref. [125].

are performed digitally on the stored data. This front-end electronics is employed for all studies of NTD-equipped cryogenic calorimeters presented in the following chapters.

Noise sources

Detectors that rely on ionisation or scintillation convert the deposited energy into a countable number of charge carriers or photons. Because the creation of a single carrier requires a few eV, the Poissonian fluctuations in that number are relatively large and set an intrinsic limit on the energy resolution. Cryogenic calorimeters, which measure the temperature rise of the absorber instead of counting carriers, are not subject to this specific limitation. In practice, though, macrocalorimeters instrumented with NTD thermistors must contend with several additional noise mechanisms, which are detailed in the following.

Intrinsic noise When energy is deposited in a thermal-equilibrium calorimeter, it first produces a population of *athermal* phonons whose energies are much larger than $k_B T$. These phonons lose energy by scattering from lattice sites and electrons, eventually converting into a larger number of lower-energy phonons. Once this thermalisation is complete, the absorber reaches a new equilibrium at a slightly higher temperature and the phonon spectrum becomes thermal again.

During the entire process the thermometer contains, on average, only a fraction β of the total energy, and this fraction can fluctuate. If the thermal link between thermometer and absorber is strong⁴, the thermometer effectively samples the phonon distribution many times [126]. Under these conditions the dominant limitation on energy resolution arises from random heat flow through

⁴That is, the time required for the whole system to reach equilibrium is shorter than the signal-formation time.

the weak link to the bath. One can regard these fluctuations as \sqrt{N} variations in the number of thermal phonons, each with energy $k_B T$.

The corresponding intrinsic limit, known as *thermodynamic fluctuation noise* (TFN), is [127]

$$\Delta E = \sqrt{k_B T^2 C} \quad (2.11)$$

where T is the absorber temperature and C its heat capacity. Thermalization effects, the detector's transfer function bandwidth, and other non-idealities can already augment the intrinsic energy resolution, which can be expressed as [128]:

$$\Delta E = \xi \sqrt{k_B T_b^2 C} \quad (2.12)$$

where T_b is the temperature of the heat sink, $C = C(T_b)$ is the capacity of the detector at T_b , and ξ is an adimensional factor, typically between 1.5 and 2.0, which contains the information about the sensor, the thermal link, and the temperature dependence of $C(T)$. ξ can be made smaller in electrothermal-feedback devices [126, 128].

The intrinsic TFN limit is extremely small. For example, a 1 mg Si crystal operated at 50 mK would, in principle, achieve a resolution below 1 eV, two orders of magnitude better than state-of-the-art Si(Li) detectors [129]. In practice, macro-calorimeters fall short of this ideal. CUORE, whose TeO_2 absorbers weigh 750 g each, attains a full-width-half-maximum (FWHM) energy resolution of (7.310 ± 0.024) keV at $Q_{\beta\beta}$ [116], roughly one order of magnitude above the TFN estimate of $\mathcal{O}(100)$ eV. This degradation has both intrinsic origins, and extrinsic contributions (which usually dominate) such as Johnson–Nyquist noise and microphonic vibrations; the latter are discussed in this section.

- **Radiative losses** [127]. Ionisation produced by the interaction can recombine emitting photons that escape the crystal, removing energy and adding \sqrt{N} fluctuations. This is especially true in scintillating cryogenic calorimeters (introduced in Sec. 2.4.1), in which the absorber crystal emits scintillation light.
- **Metastable states and thermalisation noise** [127, 130]. The energy released by the interaction can create long-lived excitations (e.g., permanent deformations in the crystal's lattice) reducing the measurable amount of energy. Fluctuations in the energy that goes into long-lived metastable states can impact the energy resolution.
- **Early-stage phonon effects** [127, 130]. Athermal phonons may leak out through mechanical supports broadening the observed spectrum.

Extrinsic noise Cryogenic calorimeters are usually subject to noise from various electrical and mechanical noise sources.

- **Johnson–Nyquist noise.** Thermal agitation of charge carriers in a conductor produces white noise with spectral density $\overline{e_n^2} = 4k_B T R$, where k_B is Boltzmann's constant, T is the temperature, and R is the conductor's resistance. Examining the read-out circuit in Fig. 2.5, two distinct Johnson–Nyquist contributions are evident:

1. *Series noise* generated by the NTD thermistor itself.

2. *Parallel noise* from the load resistors, which operate at room temperature and therefore dominate. The current fluctuations they introduce are $i_n^2 = 4k_B T/R_L$, which naturally translate into voltage fluctuations across the NTD: $e_n^2 = R_{\text{NTD}}^2 i_n^2$.

The total Johnson–Nyquist voltage noise is the quadratic sum of these two terms:

$$\overline{V_{\text{Johnson}}^2} = 4k_B T_b R_{\text{NTD}} + 4k_B T \frac{R_{\text{NTD}}^2}{R_L} \approx 4k_B T \frac{R_{\text{NTD}}^2}{R_L} \quad (2.13)$$

Because the NTD operates at $T_b \simeq 15$ mK, the first term is usually negligible and it has always been ignored in the measurements presented in the following chapters. The load resistors R_L , on the other hand, are at room-temperature ($T \sim 300$ K) and their Johnson–Nyquist noise cannot be ignored.

- **Preamplifier noise.** The preamplifier contributes with both parallel and series noise. All the measurements described in the following chapters were acquired with the front-end electronics described in Refs. [131, 132, 117], which employs ultra-low-noise JFET amplifiers. Their white series noise is approximately $3.2\text{--}3.3$ nV/ $\sqrt{\text{Hz}}$, while the $1/f$ component reaches about 7.8 nV/ $\sqrt{\text{Hz}}$ at 1 Hz.

The combined contributions of preamplifier noise and Johnson–Nyquist noise define an irreducible lower bound on the system’s white noise. In the following chapters, this bound will be referred to as the *noise floor*, given by

$$\text{Noise Floor} = \sqrt{\overline{V_{\text{Johnson}}^2} + \overline{V_{\text{pre}}^2}} \quad (2.14)$$

where $\overline{V_{\text{pre}}^2}$ denotes the white component of the preamplifier noise, referred to its input. The noise floor is typically reported as a voltage noise spectral density with units V/ $\sqrt{\text{Hz}}$.

- **Electromagnetic pick-up noise:** Stray fields and poor grounding around the sensor can inject sharp spikes or momentarily shift its bias, producing brief jumps in the output voltage. Encasing the front-end electronics in a Faraday shield largely blocks these external fields. Using tightly twisted pairs of wires can also reduce this effect. Some residual electrical interference noise might still be present even when all these mitigation strategies are used.
- **Mechanical vibrations** of the cryostat can cause microphonic noise in the cables. The wires change their mutual and ground capacitances, so the charge sensed by the read-out stage fluctuates. Using very short leads, tightly twisted pairs, and clamping them firmly to the cryostat structure helps to suppress this problem. Another consequence of vibrations is in the mechanical coupling with the absorber crystals: the vibrations reaching them dissipate energy as heat and mimic genuine thermal pulses whose bandwidth is overlapped to the one of the signal. The exact level depends on the cryostat design and detector mounting, but it can be lessened by mechanically isolating the detector–cryostat assembly from external supports. Even with all these mitigation strategies in place, mechanical vibrations and microphonic noise can still have an impact on the detector’s energy resolution.

All of the noise sources listed above act before the first amplifier, so they directly degrade the SNR and, by extension, the energy resolution and threshold. Noise generated after the amplification stage is negligible when referred back to the preamplifier input and can therefore be disregarded. Finally, the chosen WP sets the system’s transfer function; adjusting it can selectively enhance or suppress particular noise contributions.

2.3.2 Advantages and limitations of cryogenic calorimeters

Section 2.3.1 presents the operating principles of NTD-equipped cryogenic macrocalorimeters; the features described therein demonstrate that the intrinsic properties of these detectors render them particularly suitable for $0\nu\beta\beta$ searches. The following list shows that many of the requirements dictated by the detector factor of merit $S^{0\nu}$ introduced in Eq. 2.4 are met:

- **Excellent energy resolution:** Even when all the noise sources described in Sec. 2.3.1 are taken into account, various experiments and demonstrators have proven that these devices can have relative FWHM energy resolutions $\Delta E/E$ of the order of $\mathcal{O}(0.1\%)$ at $Q_{\beta\beta}$ [109], comparable to HPGe diodes.
- **Source = detector configuration:** The $\beta\beta$ isotope is embedded in the absorber, giving typically high detection efficiencies.
- **Scalability and exposure:** Long-term stable operation of an array of approximately 1000 crystals cooled by a single dilution refrigerator has been demonstrated by the CUORE experiment [116, 104], achieving raw TeO_2 exposures at the tonne scale.
- **Broad choice of absorber materials:** Any compound that can be grown as a good quality, stable, radiopure, dielectric and diamagnetic crystal, such as TeO_2 , Li_2MoO_4 , can in principle serve as an absorber. This versatility offers several advantages:
 - A material containing atoms whose corresponding $0\nu\beta\beta$ candidate isotope have large natural abundance, for example ^{130}Te (see Table 1.3), can be chosen, reducing enrichment requirements.
 - Crystals containing isotopes endowed with favourable nuclear matrix elements can be selected to maximise sensitivity to $m_{\beta\beta}$.
 - If a scintillating compound is employed, the different light yield for α and β/γ interactions⁵ permits efficient α -background discrimination in the ROI if a secondary detection channel to measure scintillation light is included in the detector’s design.
- **Low detection threshold and linear response:** The CUORE experiment has shown that a large array of cryogenic calorimeters can be used to perform analysis in a wide energy range, from keV to MeV, thereby demonstrating the versatility of this technique for rare-event searches [134]. A $0\nu\beta\beta$ cryogenic calorimeter apparatus can thus operate as a multi-purpose detector.

Experiments based on other technologies can benefit from topology reconstruction to reduce backgrounds. Experiments based on external sources can be designed to have tracking-capabilities that will be essential to study the properties and the physics of the $0\nu\beta\beta$ in case of discovery. With respect to these experiments, cryogenic calorimeters show the following limitations:

- **No event-topology reconstruction:** Large arrays of cryogenic calorimeters register just a single energy deposit per crystal, in contrast with gaseous time-projection chambers that reconstruct event topology. As a result, multi-site γ -induced backgrounds can be reduced only through inter-crystal anti-coincidence, a technique that affords limited rejection power.

⁵The light quenching of α particles is described by Birks’ empirical law [133]. It arises from local saturation of excitations along the dense α track due to its larger stopping power relative to β/γ .

- **Slow signals and pile-up:** As Fig. 2.4a shows, pulses have rise times of order 0.1 s. This slow characteristic time implies that pile-up may be a concern. In particular, isotopes with comparatively short $T_{1/2}^{2\nu}$, such as ^{100}Mo , could suffer from $2\nu\beta\beta$ pile-up unless dedicated rejection techniques are implemented.
- **Crystal-growth challenges:** Germanium lends itself to ultra-high-purity single-crystal growth because several intrinsic material properties and mature processing steps act in concert; HPGe detectors therefore rely on a well-industrialised technology. For cryogenic calorimeters the challenge is greater: each absorber compound must be developed with a bespoke growth protocol, and stringent radiopurity targets must be met at every chemical and thermal step, factors that lengthen schedules and inflate costs when large arrays are planned. The high price of isotopically enriched raw materials implies that the production yield must be high, raising complexity and reducing manufacturing speed. If the absorber crystal is hygroscopic or mechanically fragile it must be handled in inert-atmosphere and surface treatment may require customized steps that can add further complexity.

The tonne-scale CUORE experiment demonstrates that cryogenic calorimeters constitute an effective technology for $0\nu\beta\beta$ searches taking advantage of many of the features described above.

2.3.3 The CUORE experiment

The Cryogenic Underground Observatory for Rare Events (CUORE) is an experiment located at the *Laboratori Nazionali del Gran Sasso* (LNGS), an underground laboratory in central Italy. Its main goal is the search for $0\nu\beta\beta$ of ^{130}Te [104]. ^{130}Te has a relatively high natural isotopic abundance, around 34% (see Table 1.3) and $Q_{\beta\beta} = 2526.97(23)$ keV [136], which falls between the ^{208}Tl 2615 keV γ line (the end-point of natural γ radiation) and its compton shoulder, in an energy region which is relatively free from most environmental γ backgrounds.

The detector is constituted by 988 TeO_2 cubic crystals, each with dimensions $5 \times 5 \times 5$ cm³ arranged into 19 independent towers each having 13 floors with 4 crystals per floor. The absorber mass totals 742 kg, corresponding to 206 kg of ^{130}Te . Each crystal is instrumented with a Ge-NTD thermistor and constitutes an independent cryogenic calorimeter; its operating principle and the associated electronic readout system are described in detail in Sec. 2.3.1. A silicon heater on each absorber injects periodic, monochromatic pulses that emulate the energy released by particle interactions and can be used to stabilise for intrinsic gain changes due to temperature drifts throughout data-taking. A photograph of a CUORE Ge-NTD equipped TeO_2 calorimeter is shown in Fig. 2.3b. Each crystal is supported by polytetrafluoroethylene (PTFE) holders that keep it in place in the copper structure. The towers are anchored to the mixing chamber (MC) of a custom cryogen-free low-noise dilution cryostat that maintains a base temperature of ~ 15 mK [137]. The copper structure holding the crystals is thermalized to the MC and also serves as a heat sink; the TeO_2 crystals are thermally coupled to the bath according to the thermal model depicted in Fig. 2.4b. A rendering of the CUORE experiment is shown in Fig. 2.8.

An overburden of ~ 1300 m of rock provided by the *Gran Sasso* mountain suppresses the surface muon flux by ~ 6 orders of magnitude [138]. Additional passive shielding encloses the cryostat: multiple external and internal shields with an almost 4π coverage attenuate environmental γ backgrounds and neutrons [139, 137]. The inner shield is mostly made of ancient Roman lead, which is depleted in ^{210}Pb [140]. The main background in the ROI originates from degraded α particles emitted by trace contaminants in copper surfaces facing the calorimeters or on the surface of




Figure 2.8: The CUORE experiment and its infrastructure. (a) Photorealistic rendering of the whole experiment, including the support structure and external shields (Lightmap Creative S.r.l.). (b) Rendering of the cryostat (top) and one tower (bottom) as implemented in the Monte Carlo simulation used to build the background model. Only the cryostat rendering (top) is drawn to the indicated length scale. Reprinted from Ref. [135].

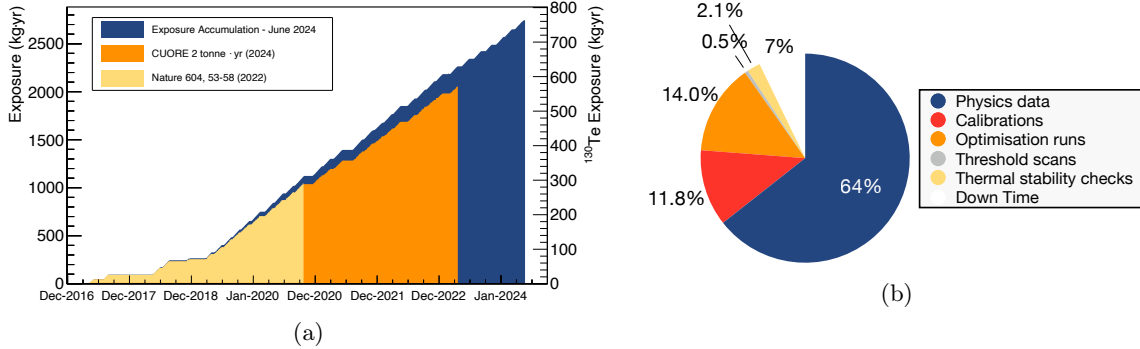


Figure 2.9: (a) CUORE exposure up to June 2024. The analysis of the first tonne-yr data set is reported in Ref. [103], whereas the results obtained with the complete 2 tonne-yr exposure are presented in Ref. [116]. The difference between the raw exposure accumulation (blue) and the published data in the corresponding time period (yellow,orange) arises from data-quality cuts that reject occasional thermal instabilities of the detectors. (b) Distribution of CUORE running time among physics data, calibration runs, and maintenance over the same period.

the crystals themselves. This contribution is mitigated by dedicated cleaning procedures and strict radiopurity controls on the crystals and nearby structures [141, 142, 143].

Current results from CUORE

CUORE is the first tonne-scale experiment based on cryogenic calorimeters able to set competitive limits on $T_{1/2}^{0\nu}$ and $m_{\beta\beta}$. With over 5 years of stable operation it has already collected and analyzed a raw TeO_2 exposure of $MT = 2039 \text{ kg}\cdot\text{yr}$. CUORE is nearing the end of its data taking and plans to continue until reaching $3 \text{ tonne}\cdot\text{yr}$ which corresponds to around $1 \text{ tonne}\cdot\text{yr}$ of ^{130}Te exposure.

Data taking The CUORE data taking intervals background and calibration measurements with routine maintenance operations on the cryostat. Each dataset contains roughly one week of calibration using external ^{232}Th and ^{60}Co sources, followed by ~ 1.5 months of background measurement. The background data contribute to the total exposure and form the sample employed in the search for $0\nu\beta\beta$ decay. Since 2019 CUORE has maintained a duty cycle above 90 %, confirming that a large-scale array of cryogenic calorimeters operated at $\sim 15 \text{ mK}$ can run stably over multi-year periods [116]. The CUORE accumulated exposure is shown in Figure 2.9a and its duty cycle in Figure 2.9b.

Analysis pipeline The data collected by each calorimeter is digitized with the readout scheme shown in Fig. 2.7. After triggering the data [144], the analysis proceeds by denoising [145] and optimally filtering the waveforms [146], after which various data-quality selections are applied to reject spurious events [147]. Pulse shape discrimination (PSD) removes pulses with irregular temporal profiles that may correspond to misreconstructed or spurious events, improving energy resolution and sensitivity. In addition, an anti-coincidence (AC) cut rejects events depositing more than 40 keV in more than one crystal within a 5 ms window, suppressing background from multi-Compton γ




Figure 2.10: CUORE total energy spectrum after successive stages of event selection obtained from an exposure of 2039.0 kg · yr of TeO₂. Spectra after the basic quality (base), anti-coincidence (AC), and pulse-shape discrimination (PSD) cuts are shown separately. A red vertical line marks $Q_{\beta\beta}$ for ¹³⁰Te, while the shaded band indicates the ROI. Reprinted from Ref. [116].

interactions and surface radioactivity. The total spectrum obtained after each of these cuts is shown in Fig. 2.10.

Results The most recent results are obtained from a total exposure of 2039.0 kg · yr of TeO₂, corresponding to an effective exposure of 567.0 kg · yr of ¹³⁰Te used in the $0\nu\beta\beta$ analysis [116]. A region of interest (ROI) with sufficient statistics to fit backgrounds is chosen around $Q_{\beta\beta}$ in the range [2465, 2575] keV. The main background contributions in the ROI consists of a nearly flat component from degraded surface α particles, a 2505.7 keV sum peak originating from the coincident absorption of the two γ rays emitted by ⁶⁰Co, and a sub-dominant contribution from multi-Compton scattering of the 2615 keV ²⁰⁸Tl γ rays. The fitting model therefore includes a Gaussian line at $Q_{\beta\beta}$, a dataset dependent flat background, and a time-dependent Gaussian for the ⁶⁰Co peak. The detector response is modeled using the 2615 keV ²⁰⁸Tl γ line. The knowledge of the detector response allows to define priors for the signal amplitude and resolution. The FWHM energy resolution of the detector, obtained with an exposure-weighted harmonic mean, is (7.540 ± 0.024) keV at the 2615 keV calibration peak, corresponding to $\sim 0.3\%$ relative resolution. Several strategies are used to compute the overall signal efficiency, given by the product of containment and cuts efficiencies. Its value and its associated systematic uncertainties are treated as nuisance parameters. The number of events above background is used to search for $0\nu\beta\beta$. No excess compatible with $0\nu\beta\beta$ decay is observed. The Bayesian analysis yields a lower limit on the half-life [116]:

$$T_{1/2}^{0\nu} > 3.5 \times 10^{25} \text{ yr} \quad (2.15)$$

at 90% confidence interval. The best-fit result is plotted in Fig. 2.11a.

Under the assumptions made in Sec. 2.1 the half-life limit can be converted into an upper bound on the effective Majorana mass. Consequently an upper limit of:

$$m_{\beta\beta} < 70\text{--}250 \text{ meV} \quad (2.16)$$

is obtained, where the interval reflects the spread of current nuclear matrix-element calculations. Figure 2.11b presents this result on the $(m_{\beta\beta}, m_{\text{lightest}})$ plane defined in Sec. 1.5.3.

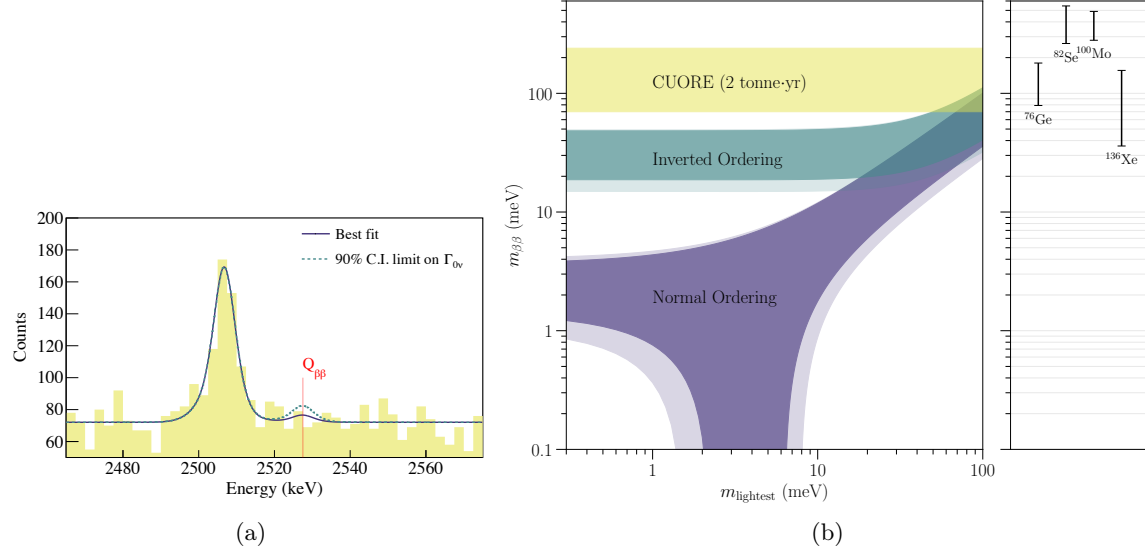


Figure 2.11: (a) ROI spectrum after all selection cuts (yellow), overlaid with the best-fit model (purple) and with the same model constrained to the 90 % confidence-interval upper limit on the $0\nu\beta\beta$ decay rate (green dashed). (b) Limits on $m_{\beta\beta}$ plotted against m_{lightest} . The area permitted by oscillation data for inverted (normal) mass ordering appear in green (purple), with lighter shades indicating the 3σ uncertainties on the mixing parameters. The yellow band represents the latest constraint from CUORE, while exclusion regions obtained for ^{76}Ge [148], ^{82}Se [149], ^{100}Mo [150], and ^{136}Xe [151] are overlaid for comparison.

CUORE experience and the road to CUPID

With an excellent relative energy resolution of $\simeq 0.3\%$ at the $Q_{\beta\beta}$ of ^{130}Te and an exposure exceeding $2 \text{ tonne} \cdot \text{yr}$ of TeO_2 accumulated in more than 5 years of uninterrupted data taking (see Fig. 2.9a), CUORE fulfils two of the key requirements that enter the factor of merit introduced in Eq. (2.4), demonstrating the potentials of cryogenic calorimeters for next-generation $0\nu\beta\beta$ searches.

CUORE backgrounds The background index (BI) is computed by performing a fit in the ROI in the background-only hypothesis. The exposure-weighted mean value obtained on an exposure of $2039.0 \text{ kg} \cdot \text{yr}$ is [116]:

$$\text{BI} = (1.42_{-0.02}^{+0.03}) \times 10^{-2} \text{ counts keV}^{-1} \text{ kg}^{-1} \text{ yr}^{-1} \quad (2.17)$$

The main contribution to this BI is due to surface α decays. Recently, a data-driven background model was presented in Ref. [135]. This analysis, based on an exposure of $1038.4 \text{ kg} \cdot \text{yr}$ (whose corresponding $0\nu\beta\beta$ result is published in Ref. [103]), indicates that degraded α particles account for approximately 75 % of the BI. The dominant contribution originates from the so-called detector's

close parts (defined in Ref. [135]), whose most relevant components are the copper pieces facing the crystals (tower frames, supports for readout wires and tiles covering the inside of the MC shield). In addition, α decays occurring within the outermost micrometres of the TeO_2 crystals may contribute when either the α particle or the recoiling daughter nucleus escapes the active volume after depositing a degraded fraction of its energy and subsequently interacts with adjacent passive material, thereby leaking into the ROI. Fig. 2.14a shows a breakdown of the contributions to the BI from other volumes according to the results of the data-driven background model. If the residual α component can be further suppressed through dedicated mitigation strategies, the background in the ROI may fall into the zero-background regime, where the sensitivity given by Eq. (2.5) grows linearly with exposure.

The reduction of the contribution of α events in the ROI is precisely the goal of CUPID, a next-generation $0\nu\beta\beta$ experiment introduced in Sec. 2.4, which will use active tagging strategies to remove counts from degraded α particles.

Radiopurity of the absorber crystals The CUORE experience underlines the need for intrinsically clean absorber material. High-purity growth protocols [152], followed by systematic underground validation runs (CUORE Crystal Validation Runs) [153], proved effective in keeping the bulk radio-contamination of TeO_2 crystals negligible in the ROI, as demonstrated by the background model [135]. The same philosophy now guides the development of any next-generation $0\nu\beta\beta$ cryogenic calorimeter experiment, where the benchmarks and know-how inherited from CUORE provide the reference for crystal production and validation.

2.4 The CUPID Experiment

CUPID (CUORE Upgrade with Particle IDentification) [154] is a next-generation experiment that will deploy 1596 isotopically enriched scintillating Li_2MoO_4 crystals, currently in pre-production phase, to search for neutrinoless double-beta decay of ^{100}Mo . If no signal will be observed, CUPID will constrain the half-life $T_{1/2}^{0\nu}$ and hence set competitive limits on $m_{\beta\beta}$.

CUPID will operate each enriched Li_2MoO_4 crystal as an independent scintillating cryogenic calorimeter, measuring both heat and scintillation light. The high $Q_{\beta\beta}$ of ^{100}Mo and the ability to tag α events (via their reduced light yield compared to β/γ interactions) will lower the BI in the ROI by two orders of magnitude relative to CUORE. To achieve this result, each Li_2MoO_4 crystal must satisfy stringent requirements that combine ultra-high radiopurity with excellent calorimetric and scintillation performance.

2.4.1 Scintillating cryogenic calorimeters

Scintillating cryogenic calorimeters are a specific subclass of cryogenic calorimeters that employ a dual-channel readout system:

- **Heat (main) channel:** An absorber crystal, instrumented with a thermal sensor and maintained at ~ 15 mK, records the energy deposited by each particle interaction through temperature pulses registered by the thermometer. If operated as detailed in Sec. 2.3.1, this channel can be used to perform excellent energy-resolution spectroscopy.
- **Light (side) channel(s):** One (or multiple) secondary detector(s), facing the absorber crystal, measures scintillation or Cherenkov light. This detector, known as the *light detector*

(LD), is itself a cryogenic calorimeter operated at ~ 15 mK optimized for low thresholds and converts incident photons into heat pulses. Typically, the LD consists of a Ge or Si wafer equipped with a thermal sensor, which, similarly to the main channel, is operated as detailed in Sec. 2.3.1. By correlating signals in the light channel with those in the main channel, a light yield can be assigned to each corresponding heat pulse, on an event-by-event basis. Although it is called “light” detector, the LD also detects ionizing radiation from other types of impinging particles, such as X rays or nuclear recoils as well as γ/β rays and α particles.

Fig. 2.12a presents an idealized sketch of the typical design of a scintillating cryogenic calorimeter while Fig. 2.12b shows a real-world implementation. This dual (heat + light) readout system enables active, event-by-event particle identification (PID). In particular, when a scintillating crystal is used in the main channel, the quenching described by Birks’ empirical law [133] results in a lower light yield (LY) for α particles (or nuclear recoils) than for electrons. Thanks to the combination of a wide selection of absorber materials and PID capabilities, scintillating cryogenic calorimeters constitute highly versatile devices for diverse applications in particle




Figure 2.12: (a) Idealised sketch of a cryogenic calorimeter: a particle interaction deposits energy in the absorber crystal, producing both a thermal and a light signal. The thermal sensor attached to the absorber measures the heat, while the light detector converts the scintillation light into a heat pulse that is read out by its thermometer. Courtesy of M. Mauri. (b) Example of a real-world implementation: two absorber crystals (in this case Li_2MoO_4) are secured in copper frames with PTFE supports and instrumented with Ge NTD thermistors and read-out wiring. Two Ge LDs, each equipped with Ge NTD sensors and mounted on PTFE elements, face the top and bottom surfaces of each crystal, for a total of four LDs. This photograph shows the CUPID design described in Ref. [155].

physics. For instance, they have been proposed for the detection of astrophysical neutrinos [106] and for searches of dark matter [107, 108]. Ref. [109] provides a review of suitable materials and their applications. When applied to a $0\nu\beta\beta$ experiment, this dual-channel technique suppresses the dominant α -induced background component.

Light Yield and Discrimination Power

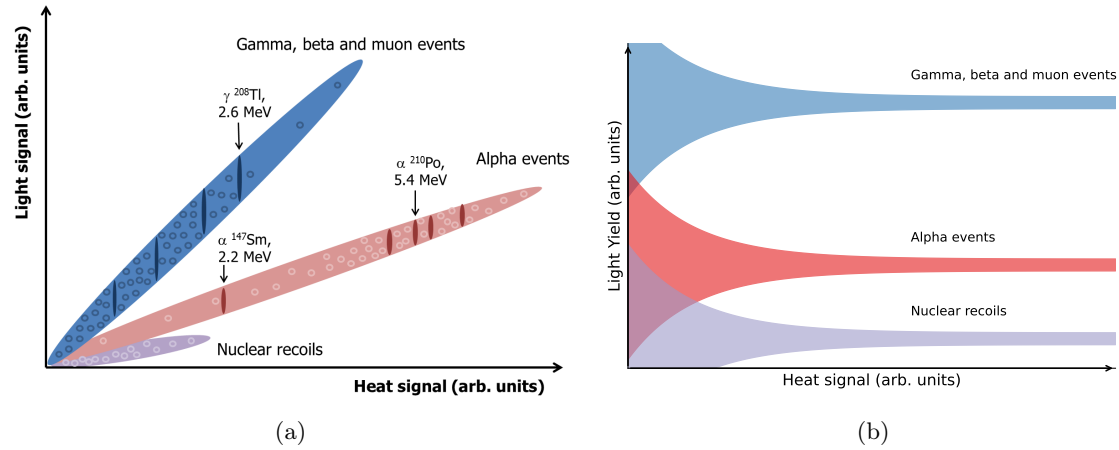


Figure 2.13: (a) Sketch showing the energy registered in the LD as a function of the coincident energy measured in the main channel. Each circle marks an event, which aggregate into peaks (dark ovals) along the $\beta/\gamma/\mu$ and α bands. Highly ionizing particles produce less scintillation than electrons, γ -rays or muons of the same energy. As a result, different ionization properties lead to clear separation in a light-vs-heat scatter plot, a distinction that grows stronger at higher energies when more scintillation light is emitted. Reprinted from Ref. [109]. (b) Light-to-heat ratio (i.e. the y-over-x coordinate in the left-hand scatter plot) as a function of the energy measured in the main channel. Under the assumptions discussed in the text, this ratio is directly proportional to each particle species' LY in the crystal. The separation between these bands (relative to their widths) enables PID.

For each particle interaction, two thermal pulses are recorded simultaneously: a signal in the absorber and a corresponding scintillation light pulse in the LD. Plotting the light-signal amplitude versus the heat-signal amplitude yields a scatter plot similar to that sketched in Fig. 2.13a, in which different particle species form distinct bands. Assuming that the energy converted into scintillation light is negligible compared to that converted into heat in the absorber⁶, the light yield of each event can be estimated by the L/H ratio, defined simply as the ratio between the energy registered by the light detector and that registered by the heat channel (i.e., the ratio of the y- and x-coordinates in the plot of Fig. 2.13a). It should be noted that, for designs such as that shown in Fig. 2.12, the

⁶For energy conservation, the light-to-heat (L/H) ratio exhibits anti-correlation for particles with associated light emission [156, 157]. Although it is often negligible, this effect may be visible in crystals having a higher scintillation yield. In these cases, the LY can still be defined by the L/H ratio, but several effects arise that non-negligibly affect both the reconstructed energy and the energy resolution in the heat channel; these must be addressed with a dedicated analysis which is typically required only for crystals having higher light yield than Li_2MoO_4 .

geometric light-collection efficiency is less than unity. However, under the above assumption, the L/H ratio is directly proportional to the intrinsic LY of the absorber crystal. In the measurements presented in the following chapters, “LY” will refer to this L/H ratio.

Both plots in Fig. 2.13 demonstrate that scintillating cryogenic calorimeters separate different types of interacting particles into distinct bands. In the LY plot, β and γ interactions exhibit the same LY, whereas α particles (and neutrons) locally saturate the radiative recombination centers and thus produce a reduced LY. Consequently, the LY provides an effective parameter for event-by-event particle identification. However, several effects broaden the LY distributions, including Poissonian fluctuations in the number of photons collected by the LD and the intrinsic energy resolution of the LD. The energy-dependence of these broadenings is discussed in Sec. 5.4.3. Typically, in Li_2MoO_4 detectors, given the small amount of light emitted by highly ionizing particles, the α and nuclear-recoil bands overlap. This does not impair the signature of $0\nu\beta\beta$, since, for a neutrinoless double-beta decay experiment, the critical requirement is to discriminate β/γ interactions from α events in the ROI around $Q_{\beta\beta}$. To quantify the separation between two bands (e.g., α versus β/γ), the discrimination confidence level, or discrimination power (DP), is defined as a function of the heat-channel energy E as

$$\text{DP}_{\alpha/\gamma(\beta)}(E) = \frac{|\mu_{\gamma(\beta)} - \mu_{\alpha}|}{\sqrt{\sigma_{\gamma(\beta)}^2(E) + \sigma_{\alpha}^2(E)}} \quad (2.18)$$

where μ_j and $\sigma_j(E)$ are the mean and standard deviation of the light-signal distribution for species j .

Light detection

An excellent discrimination power, DP, is indispensable for effective PID capabilities and, in turn, for suppressing backgrounds in the ROI. In order to achieve this, two conditions must be met simultaneously. First, the absorber crystal must have a sufficiently high LY so that Poisson fluctuations in the number of collected photons do not have a drastic impact on the relative width of the distribution. The LY is governed by the compound itself and, for a given compound, by its microscopic quality: point defects or trapping centres can compete with radiative recombination and quench the light output. This issue is especially relevant for Li_2MoO_4 , whose intrinsic LY is modest and can be further reduced by growth defects [158, 159, 160]. Second, the LD must act as a cryogenic calorimeter with good calorimetric performances and excellent energy resolution so as not to have a drastic impact on the width of the LY distribution. From these premises, three experimental optimization opportunities to improve the DP in the ROI arise:

- (a) The number of photons reaching the LD could be maximised by applying reflective coatings on lateral faces that do not view the LD. Any solution adopted must leave both the radiopurity and the heat-channel calorimetric performance unchanged. Dedicated measurements to assess the impact of reflective aluminum coatings on the calorimetric performances of the heat channel are presented in Chapter 6.
- (b) The LD must have excellent calorimetric performances. Firstly, the LD must be operated at a WP that maximises its SNR ratio, thereby minimising its contribution to the denominator of Eq. (2.18). Secondly, additional ways to increase the SNR may be used. In particular, Neganov–Trofimov–Luke (NTL) amplification [161, 162], implemented by biasing metal electrodes fabricated on the Ge absorber, can linearly increase the heat signal and, consequently,

the SNR. With this technology it is possible to get an effective amplification factor up to 60 [163, 164].

- (c) The absorber crystal must be of excellent quality and free of defects that could degrade its scintillation yield. Dedicated measurements to assess the presence of defects in newly produced Li_2MoO_4 crystals are reported in Chapter 5.

Properly optimized scintillating cryogenic calorimeters are thus an ideal technology to search for $0\nu\beta\beta$ decay.

2.4.2 The CUPID experiment: detector, backgrounds and sensitivity

When applied to $0\nu\beta\beta$ searches, PID capabilities enabled by scintillating cryogenic calorimeters can be exploited to reduce backgrounds in the region of interest. The effectiveness of this technique was demonstrated by CUPID-0 which used 26 scintillating ZnSe cryogenic calorimeters to search for $0\nu\beta\beta$ of ^{82}Se [149] and by CUPID-Mo which operated 20 scintillating Li_2MoO_4 crystals to search for $0\nu\beta\beta$ of ^{100}Mo [150]. The PID capabilities enabled both demonstrators to reduce the α contribution to the BI in the ROI by at least one order of magnitude [154].

Isotope choice

Figure 2.14 illustrates that employing PID significantly mitigates the dominant background contribution from degraded α particles. As outlined in Section 2.3.3, this background constitutes approximately 75 % of the BI in CUORE near the $Q_{\beta\beta}$ of ^{130}Te . Furthermore, Fig. 2.14b demonstrates the advantage of selecting an isotope with a $Q_{\beta\beta}$ value above 2615 keV, as this places the ROI beyond the energy range of most natural γ -ray backgrounds. In this regard, ^{100}Mo stands out as a particularly promising candidate. As shown in Fig. 2.14b, its $Q_{\beta\beta}$ value of 3034 keV locates

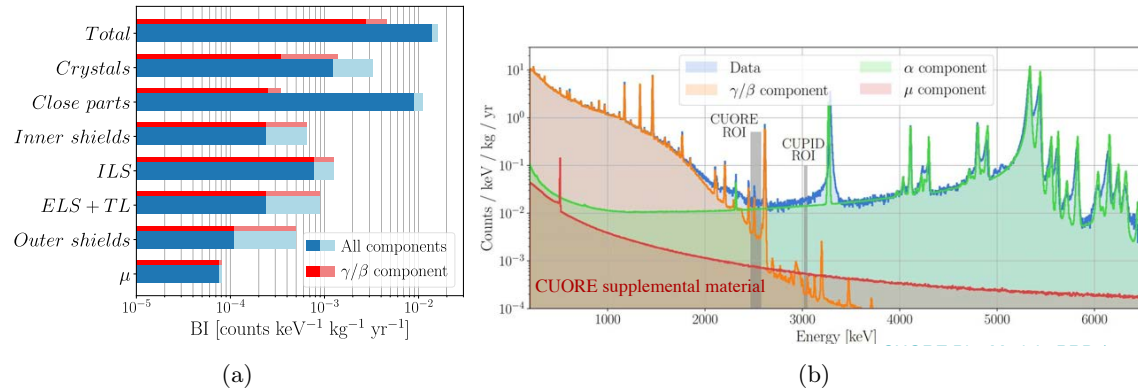


Figure 2.14: (a) Breakdown of the contributions to the BI in CUORE, evaluated using a data-driven background model built with an exposure of 1 ton·year. The light and dark blue bands indicate the impact of systematic uncertainties on the BI distribution. The red bands represent the corresponding contribution from β/γ events. The volumes shown correspond to those defined in Ref. [135] and displayed in Fig. 2.8. Reprinted from Ref. [135]. (b) Breakdown of the CUORE background components as determined from the background model study [135].

the ROI in a region substantially cleaner of γ/β backgrounds compared to isotopes with lower $Q_{\beta\beta}$ values, such as ^{130}Te .

The high $Q_{\beta\beta}$ of ^{100}Mo not only helps in background suppression but also enhances the phase space factor $G_{0\nu}$, whose leading term scales with $Q_{\beta\beta}^5$. As it can be seen from Eq. (2.2), this also improves sensitivity to $m_{\beta\beta}$. The CUPID-Mo small-scale demonstrator has shown that Li_2MoO_4 crystals exhibit excellent performance when used as scintillating cryogenic calorimeters. This material offers outstanding calorimetric characteristics and intrinsic scintillation properties, with a LY sufficient for effective PID even without doping, thereby enabling the required DP capabilities [150]. Moreover, CUPID-Mo has demonstrated that it is possible to produce Li_2MoO_4 crystals enriched in ^{100}Mo with the radiopurity level required for performing $0\nu\beta\beta$ searches [165], confirming their viability for use in next-generation experiments. A full-scale experiment employing this material must ensure that strict radiopurity standards are guaranteed during each step of a large-scale production, with consistent and reproducible quality from the initial to the final batches. Chapter 4 discusses how these criteria can be monitored and validated starting from the pre-production phase.

The CUPID detector

The CUPID detector consists of 1596 closely packed scintillating Li_2MoO_4 cryogenic calorimeters, arranged in 57 towers of 14 floors each, with two crystals per floor [154]. The tower structure vertically alternates LDs with Li_2MoO_4 crystals, for a total of 1710 light detectors. The towers will be assembled using an innovative gravity-assisted design, in which each floor is stacked on top of the previous one solely by gravity [155]. A rendering showing the CUPID detector modules and

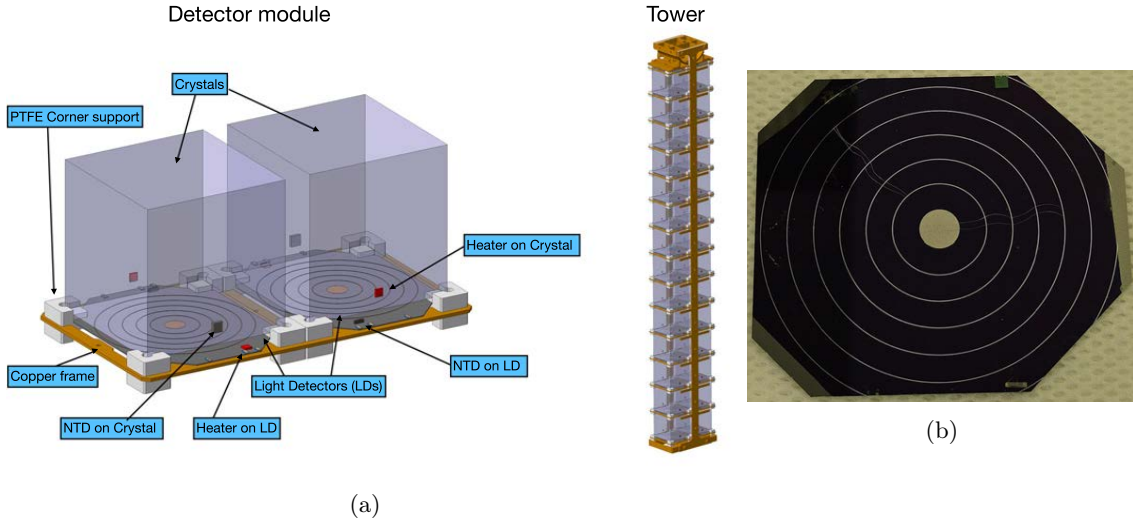


Figure 2.15: (a) Rendering of one tower with 14 floors and 28 detector modules (right), and of a single CUPID floor featuring two adjacent detector modules (left). Each module includes a $\text{Li}_2^{100}\text{MoO}_4$ absorber shaped as a cube with 45 mm sides and a Ge light detector. A photograph of this module is shown in Fig. 2.12b. (b) Example of an NTL Ge light detector equipped with two arrays of concentric, annular interleaved Al electrodes. Both images reprinted from Ref. [154].

one tower is displayed in Fig. 2.15a.

Absorber crystals Each crystal will be isotopically enriched to $\geq 95\%$ in ^{100}Mo ($\text{Li}_2^{100}\text{MoO}_4$), for a total ^{100}Mo mass of 240 kg. The crystals will be cubes with dimensions $4.5 \times 4.5 \times 4.5 \text{ cm}^3$; this size was chosen to optimize the $0\nu\beta\beta$ containment efficiency while keeping the pile-up from $2\nu\beta\beta$ events [166] under control. Chapter 3 provides a comprehensive description of the requirements for $\text{Li}_2^{100}\text{MoO}_4$ crystals, as well as details of the crystal growth process.

Light Detectors Each LD will consist of an octagonal 300 μm -thick HPGe wafer coated with a 70 nm-thick SiO anti-reflective layer that improves photon collection. The SNR of the LDs will be enhanced by NTL amplification: a set of aluminum electrodes evaporated on the wafer surface will be biased with a voltage of $\mathcal{O}(100 \text{ V})$, producing a large electric field that will amplify the signals. The improved SNR achievable with this technology will allow the contribution to the backgrounds in the ROI from $2\nu\beta\beta$ pile-up to be kept under control. The impact of this and other background sources on the BI in the ROI is discussed in the following. A picture of one Ge LD with NTL electrodes is displayed in Fig. 2.15b.

Thermal sensors and readout system Each $\text{Li}_2^{100}\text{MoO}_4$ crystal and LD will be instrumented with a Ge NTD thermistor for signal readout. Both channels will be operated according to the principles described in Sec. 2.3.1. In addition, each detector will be equipped with a Si heater, which can be used for stabilizing thermal drifts and for injecting custom pulses to study pile-up in the LDs. All the noise sources described in Sec. 2.3.1 also affect the LDs. With respect to what is described there, a few key differences must be noted: since pile-up rejection requirements demand LDs with optimal SNR, new JFET preamplifiers with a white noise contribution of 1.3–1.5 nV/ $\sqrt{\text{Hz}}$ and a $1/f$ component $\lesssim 5 \text{ nV}/\sqrt{\text{Hz}}$ at 1 Hz will be employed [125]. Moreover, a low-noise voltage power supply capable of providing a V_{bias} up to 100 V will be used in combination with $\mathcal{O}(10\text{--}100 \text{ G}\Omega)$ R_L to reduce Johnson-Nyquist parallel noise. It should also be noted that, being faster detectors, the LDs will have a signal bandwidth extending to higher frequencies, making them potentially more sensitive to high-frequency noise. To accommodate the increased LD bandwidth, anti-aliasing Bessel filters with higher cut-off frequencies will be used, along with ADCs operating at sampling frequencies up to 10 kHz. The ADCs employed will have a resolution of 24 bits.

Background budget

The CUPID detector will be installed in the same cryogenic facility currently housing CUORE at LNGS, shown in Figure 2.8. The cryogenic infrastructure for CUPID will remain largely unchanged from CUORE, allowing CUPID to inherit many of CUORE design optimizations. These include shielding from external radiation and vibrational noise mitigation systems [167, 168, 145]. The increased thermal load from having three times as many channels as CUORE will require only a few key upgrades to the cooling system.

A key advantage of reusing this infrastructure is the ability to make accurate projections of background contributions using the data-driven background model developed for CUORE [135]. As Fig. 2.8 shows, the CUORE infrastructure already has a massive shielding: a minimum of 25 cm of lead in all directions for attenuation of environmental γ radiation, and layers of 18 cm polyethylene and 2 cm boric acid powder to moderate and absorb environmental neutrons. The

plots in Fig. 2.14 show that these shieldings, combined with the high $Q_{\beta\beta}$ of ^{100}Mo , are effective in suppressing β/γ backgrounds, making degraded α particles the dominant component. CUPID PID capabilities will allow for lower background levels, resulting in the need for further suppression of neutron- and muon-induced events in order to keep them subdominant in the BI. For this reason, CUPID will add an extra 10 cm layer of polyethylene for enhanced neutron shielding and implement a compact muon veto system able to fit in the existing infrastructure. This veto will consist of plastic scintillator panels equipped with SiPMs, designed to tag approximately 99% of muons [154].

Simulations estimate a $0\nu\beta\beta$ containment efficiency of about 78% for CUPID. The region around the ^{100}Mo $Q_{\beta\beta}$ is affected by natural radioactivity, where γ interactions predominantly undergo Compton scattering, depositing energy across multiple crystals. This enables the application of an AC cut, such as the one shown in Fig. 2.10, to reduce this contribution to the BI. Combining PID, the high $Q_{\beta\beta}$, and the AC cut, CUPID can robustly suppress backgrounds in the ROI, defined as a 30 keV window around $Q_{\beta\beta} = 3034$ keV. Nevertheless, even with optimized shielding and analysis cuts, residual γ/β events from the cryogenic infrastructure with unfortunate topologies may persist, contributing to the BI. These include, for example, surface contaminants of ^{214}Bi or coincident ^{208}Tl γ rays that may be present in the innermost thermal shield [154].

Another relevant background source is $2\nu\beta\beta$ pile-up [166]. While ^{100}Mo high $Q_{\beta\beta}$ makes it an excellent candidate for $0\nu\beta\beta$ searches, it also results in a relatively short $2\nu\beta\beta$ half-life of 7.1×10^{18} yr [169], shorter than that of isotopes like ^{130}Te . In $\text{Li}_2^{100}\text{MoO}_4$ crystals of size $4.5 \times 4.5 \times 4.5$ cm³, this corresponds to a rate of 2.6 mHz per crystal [154]. If not properly mitigated, this high rate, combined with the slow detector response, can produce a flat background in the ROI from the accidental pile-up of two unresolved events. The scintillation of lithium molybdate is key to reduce this background: LDs, which have faster thermal pulses (lower rise times), can be used to reject pile-up by using the coincident light pulses from such events. Pile-up rejection depends not only on the signal's time evolution but also on its SNR. LDs with NTL amplification are employed to improve the SNR and thereby reduce pile-up backgrounds [170].

Crystal radiopurity also contributes to the background. The CUPID-Mo background model [165] is currently being used to provide a data-driven estimate of the contribution to the BI from $\text{Li}_2^{100}\text{MoO}_4$ crystals and highlights the importance of controlling crystal contamination. Pre-production measurements are essential to provide feedback to manufacturers and tailor a customized radiopurity protocol that covers all steps from production to delivery, storage and installation. To achieve this, strategies similar to those employed in CUORE [152, 153] should be adopted. A comprehensive understanding of the production process and a rigorous validation program, both of which should be established starting from the pre-production phase, are crucial. Chapter 4 will focus on these efforts.

To fully probe the IO region of $m_{\beta\beta}$, CUPID will have to achieve a BI of 10^{-4} counts keV⁻¹ kg⁻¹ yr⁻¹, which is a factor of 100 lower than the CUORE BI reported in Eq. (2.17). A breakdown of the CUPID background budget required to reach the target $m_{\beta\beta}$ sensitivity is shown in Fig. 2.16a.

Sensitivity

The sensitivity of the CUPID experiment to $m_{\beta\beta}$ depends on a few key experimental parameters, which coincide with those used to define the detector factor of merit introduced in Sec. 2.1.

One of the many advantages arising from the successful operation of demonstrators [150, 169], combined with the operational experience of CUORE, is that most of the experimental parameters influencing the $m_{\beta\beta}$ sensitivity are known with an uncertainty of the order of $\mathcal{O}(1\%)$. However,

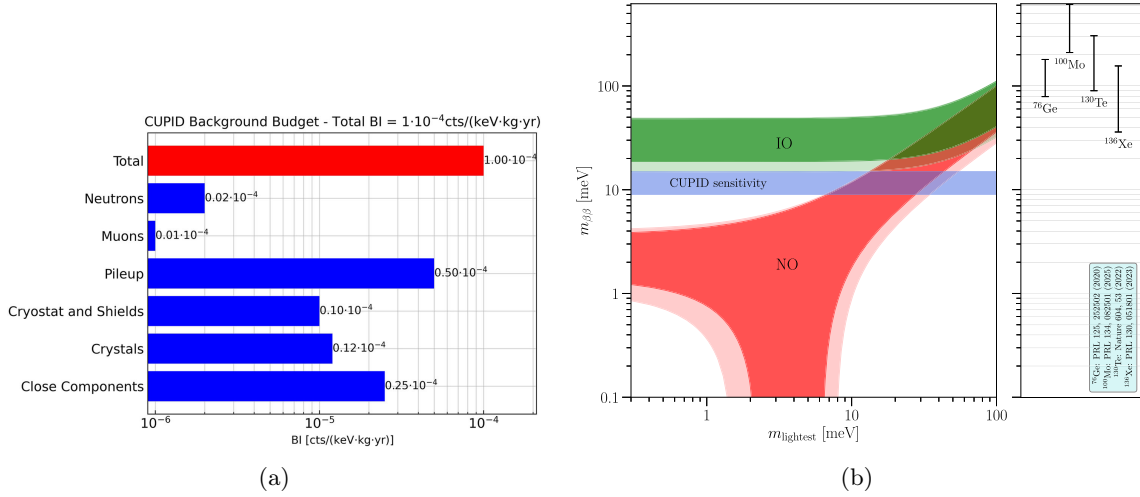


Figure 2.16: (a) Breakdown of the BI at the $Q_{\beta\beta}$ value of ^{100}Mo for the CUPID experiment. The crystal-related background contributions are inferred from the CUPID-Mo background model [165], while contributions from the CUORE infrastructure are derived from the CUORE background model [135] and those due to external radiation are obtained via dedicated Monte Carlo simulations. Reprinted from Ref. [154]. (b) Limits on $m_{\beta\beta}$ plotted against m_{lightest} . The area permitted by oscillation data for inverted (normal) mass ordering appear in (red), with lighter shades indicating the 3σ uncertainties on the mixing parameters. The horizontal shaded region indicates the expected CUPID sensitivity, which corresponds to a half-life of ^{100}Mo greater than 1.8×10^{27} yr, assuming a representative set of NMEs. The side panel displays the current limits on $m_{\beta\beta}$ derived from ^{76}Ge [148], ^{100}Mo [171], ^{130}Te [103], and ^{136}Xe [151]. Reprinted from Ref. [154].

the BI and the energy resolution are currently subject to larger uncertainties, due to an ongoing optimization campaign aimed at improving both the detector response and the radiopurity of several detector components, including the crystals themselves. Chapter 4 provides a detailed explanation of how these optimizations are being implemented, starting from the pre-production phase. Dedicated measurements to assess the radiopurity of preliminary batches of crystals from this phase

Detector parameter	Target value
Total crystal mass	450 kg
Enrichment fraction	95%
Isotope mass	240 kg
FWHM energy resolution	5 keV
Livetime	10 years
Containment efficiency	78%
Selection efficiency	90%
Background index	10^{-4} counts $\text{keV}^{-1} \text{kg}^{-1} \text{yr}^{-1}$

Table 2.2: Experimental parameters of the CUPID experiment. Adapted from Ref. [154].

will not only be useful for providing feedback to the producer, but will also serve to evaluate the calorimetric performance and contamination levels of the newly produced crystals. This knowledge can in turn be used to update and reduce the uncertainties on the parameters employed in the $m_{\beta\beta}$ sensitivity calculation.

In a recent publication [172], the sensitivity of the CUPID experiment was evaluated under various deployment scenarios of the 1596 $\text{Li}_2^{100}\text{MoO}_4$ crystals. In particular, that work computes the sensitivity achievable in both a full deployment scenario and a staged deployment scenario. With the parameters listed in Table 2.2, the median exclusion sensitivity is $T_{1/2}^{0\nu} = 1.8 \times 10^{27}$ yr, corresponding to exclusion limits for $m_{\beta\beta}$ in the range $m_{\beta\beta} = 9.0\text{--}15.2$ meV, fully covering the IO region, as shown in Fig. 2.16b. This figure also indicates that CUPID will probe the NO region for $m_{\text{lightest}} \gtrsim 10$ meV.

An alternative scenario considered in Ref. [172] involves deploying the CUPID detector in two stages. CUPID Stage-I would include one-third of the detectors, arranged in 19 towers. This staged approach offers the advantage of yielding early science data while the remaining $\text{Li}_2^{100}\text{MoO}_4$ crystals are produced and the additional towers are assembled. The full 1596 crystal configuration, corresponding to a total mass of 450 kg, would be deployed following CUPID Stage-I. The sensitivity of CUPID Stage-I can be estimated by assuming 3 years of data-taking with 150 kg of deployed crystals. Using the same parameters as in Table 2.2 and assuming a BI of 1.0×10^{-4} counts $\text{keV}^{-1} \text{kg}^{-1} \text{yr}^{-1}$, CUPID Stage-I would achieve a median discovery sensitivity of 0.2×10^{27} yr, corresponding to a range of $m_{\beta\beta} = 26.3\text{--}44.4$ meV.

It should be noted, however, that the reduced number of detectors in this initial phase has a less granular detector, limiting the informations coming from coincident events. This weakens the effectiveness of coincidence cuts used to reject muons and events from crystal contaminations. Nonetheless, current background studies suggest that the BI will remain below 1.5×10^{-4} counts $\text{keV}^{-1} \text{kg}^{-1} \text{yr}^{-1}$. For this reason, the analysis in Ref. [172] also includes sensitivity estimates for this higher BI scenario.

2.5 Conclusion

This chapter has presented the CUPID experiment and its strategy to achieve the sensitivity required for the search for $0\nu\beta\beta$ decay of ^{100}Mo , through the deployment of 1596 isotopically enriched scintillating cryogenic calorimeters based on Li_2MoO_4 crystals. CUPID builds upon the foundation established by CUORE, which demonstrated the potential of cryogenic calorimeters over more than five years of continuous operation, as well as on the experience gained from small-scale demonstrators such as CUPID-Mo.

The following chapters focus on the procurement, characterization and validation of $\text{Li}_2^{100}\text{MoO}_4$ crystals during the pre-production phase, focusing mainly on calorimetric and scintillation performance.

Chapter 3

CUPID Li_2MoO_4 Pre-Production

The realization of CUPID critically depends on the availability of 1596 high-quality and radiopure $\text{Li}_2^{100}\text{MoO}_4$ crystals. To this end, CUPID has tasked the Shanghai Institute of Ceramics (SICCAS), Chinese Academy of Sciences, with large-scale production. This is the same producer that successfully delivered the 988 radiopure TeO_2 crystals for the CUORE experiment.

Unlike TeO_2 , which has several non-academic applications in the industry, Li_2MoO_4 lacks an established production chain, making crystal growth intrinsically more challenging. In addition, the use of costly isotopically enriched molybdenum imposes the need to minimize material losses throughout all stages of production. These aspects introduce further constraints compared to CUORE, where only radiopurity had to be optimized and crystal quality was taken for granted. Despite these difficulties, the success of the CUPID-Mo demonstrator [150] demonstrated that radiopure, high-quality $\text{Li}_2^{100}\text{MoO}_4$ crystals can indeed be produced [173].

The success of the CUORE experiment relied on a close collaboration with SICCAS [152], demonstrating that sharing a technical framework with the producer is essential to enforce radiopurity requirements [153]. For CUPID this collaboration is even more critical, since crystal quality, in addition to radiopurity, directly affects both calorimetric and scintillation performance.

Accordingly, crystal production at SICCAS was documented during a *pre-production* phase that covered all stages, from raw powders to finished crystals, establishing a shared framework and common vocabulary between CUPID and the producer. This chapter provides a concise overview of this process.

3.1 $\text{Li}_2^{100}\text{MoO}_4$ Crystal Requirements

CUPID will be built with 1596 isotopically enriched $\text{Li}_2^{100}\text{MoO}_4$ crystals, with enrichment $\geq 95\%$ in ^{100}Mo . The crystals will be cubic with dimensions $4.5 \times 4.5 \times 4.5 \text{ cm}^3$, corresponding to a mass of $\sim 285 \text{ g}$ each, for a total ^{100}Mo mass of 240 kg. Table 3.1 lists the high-level specifications for the experiment. In order to achieve the target sensitivity (see Sec. 2.4.2), strict requirements are imposed on crystal properties.

Radiopurity CUPID's requirement is that crystal radioactivity contributes less than $(0.11 \pm 0.03) \times 10^{-4} \text{ cts keV}^{-1} \text{ kg}^{-1} \text{ yr}^{-1}$ to the BI, divided into a bulk component of at most $(1.4 \pm 0.6) \times 10^{-6} \text{ cts keV}^{-1} \text{ kg}^{-1} \text{ yr}^{-1}$ and a surface component of at most $(9 \pm 4) \times 10^{-6} \text{ cts keV}^{-1} \text{ kg}^{-1} \text{ yr}^{-1}$.

Requirement	Value
Total $\text{Li}_2^{100}\text{MoO}_4$ Mass	450 kg
Number of Crystals	1596
Dimensions	$4.5 \times 4.5 \times 4.5 \text{ cm}^3$
Enrichment in ^{100}Mo	$\geq 95\%$
Total ^{100}Mo mass	240 kg

Table 3.1: High-level specifications for $\text{Li}_2^{100}\text{MoO}_4$ crystals.

These numbers are derived from the radiopurity results demonstrated by the contaminations of the CUPID-Mo enriched crystals [165], which were grown at NIIC (Nikolaev Institute of Inorganic Chemistry, Novosibirsk, Russia) using material selected for radiopurity and adopting controlled growth protocols [174, 173]. Continuous interaction with SICCAS, the producer identified by CUPID, is aimed at reproducing and possibly improving these results for the full-scale production. The radiopurity requirements are summarized in Table 3.2.

	Isotope	Requirement
Bulk	^{238}U	$< 0.4 \mu\text{Bq/kg}$
	^{226}Ra	$< 0.4 \mu\text{Bq/kg}$
	^{232}Th	$< 0.4 \mu\text{Bq/kg}$
	^{228}Th	$< 0.4 \mu\text{Bq/kg}$
	^{40}K	$< 1 \text{ mBq/kg}$
	^{210}Pb	$< 0.01 \text{ mBq/kg}$
	^{60}Co	$< 0.01 \text{ mBq/kg}$
	^{137}Cs	$< 0.5 \text{ mBq/kg}$
Surface	^{226}Ra – ^{210}Pb	$< 2.0 \text{ nBq/cm}^2$
	^{228}Th – ^{208}Pb	$< 2.5 \text{ nBq/cm}^2$

Table 3.2: Radiopurity requirements for bulk and surface contamination of $\text{Li}_2^{100}\text{MoO}_4$ crystals.

Energy Resolution CUPID requirement is 5 keV FWHM at $Q_{\beta\beta}$. As discussed in Sec. 2.3.1, the energy resolution results from several intertwined detector parameters, which include not only intrinsic crystal properties but also thermal couplings, the working point configuration, and the overall noise conditions. For this reason, although the final requirement is expressed in terms of energy resolution, the crystal contribution is usually evaluated through its intrinsic gain (sensitivity). This quantity, however, must always be interpreted in the context of the working point chosen for the NTD sensor and the absorber temperature. A systematic procedure to compare crystals in terms of both energy resolution and intrinsic gain is one of the objectives of the present work and will be addressed in Sec. 4.5.7.

Light Yield The intrinsic light yield of the crystals at ~ 15 mK is difficult to measure because the geometry of scintillating cryogenic calorimeter modules has a light collection efficiency far from unity, and reflections from other elements or from the crystal surfaces are not straightforward to account for. For this reason, CUPID defines the requirement on LY in terms of the light detected on one of the two LDs (top/bottom) in the typical CUPID geometry, as illustrated in Fig. 2.12b. The requirement is $LY \sim 0.35$ keV/MeV [154], which corresponds to the amount of light needed to achieve a $\beta/\gamma-\alpha$ discrimination efficiency larger than 5 when using a properly performing LD. This requirement will also allow to use scintillation light to identify (and reject) events induced by two $2\nu\beta\beta$'s randomly producing a pile-up event. The methodology adopted to evaluate the LY in the geometry of typical CUPID modules is presented in Chapter 4.

Surface Quality The surface quality of the crystals may affect light collection efficiency. Surfaces must be homogeneous and free of halos. At present, the optimal level of surface roughness for maximizing light collection on the LDs is not yet established. A polished (optical quality) surface scatters less light but may increase internal reflections, while a rougher (opaque) surface may reduce reflections, though at the cost of partial absorption or scattering of the emitted light. To take advantage of these effects, one solution would be to polish only the two opposite faces facing the LDs and leave the other faces opaque, or vice versa, to optimize overall light collection. The best configuration is not clear yet, and it will be investigated during pre-production, where crystals with different surface finishes will be produced and compared. CUPID is also developing Monte Carlo (MC) simulations of light propagation in Li_2MoO_4 to clarify these aspects. In all cases, the chosen polishing strategy must comply with the surface contamination requirements listed in Table 3.2. Moreover, Li_2MoO_4 is hygroscopic. If handled in a humid atmosphere, the surface can degrade and become opaque or haloed. All handling at SICCAS must therefore be performed under low humidity to prevent moisture-induced surface damage.

Other Requirements Crystals must be $4.5 \times 4.5 \times 4.5$ cm³ Li_2MoO_4 cubes with phenakite structure, space group $R3$, lattice parameters $a = 9.590$ Å and $c = 16.659$ Å. They must be isotopically homogeneous with enrichment $\geq 95\%$ in ^{100}Mo . Dimensional tolerance is ± 0.2 mm, with 0.5 mm bevels and mean planarity < 0.2 mm. Crystals must be transparent, colorless, bubbleless, free of inclusions, and with absorption length > 50 cm in the 550–700 nm spectral range. The chemical purity must be $> 99.9\%$.

3.2 Overview of The Crystal Production Process

The following list provides a bird's-eye overview of the main steps of the growth process and highlights the most delicate aspects for CUPID.

1. **Precursor powders:** Li_2MoO_4 is synthesized from Li_2CO_3 and MoO_3 precursors. These powders must be procured with ultra-high radiochemical purity. Enriched crystals are produced using $^{100}\text{MoO}_3$ (95% enrichment), whose high price dominates the total cost of crystal procurement. This imposes minimizing material losses during growth and processing.
2. **Li_2MoO_4 synthesis:** Several synthesis strategies exist to produce Li_2MoO_4 powder for crystal growth. The reaction between raw powders must yield only Li_2MoO_4 , with no side products or residual unreacted precursors. Parasitic chemical species or unreacted powders

would compromise growth and crystal quality. Any ancillaries and consumables used at this stage must be selected to avoid introducing radiochemical contaminants.

3. **Li₂MoO₄ crystal growth:** Crystals are grown in platinum crucibles using the Bridgman (BG) technique. Crystal growth must yield ingots of high structural quality that meet the required specifications. Post-growth annealing is used to stabilize the crystals and prevent crack formation. Since crystallization purifies the material, multiple crystallizations will be performed to increase radiopurity. The Czochralski (CZ) technique may also be used as an intermediate crystallization stage. Each re-crystallization, however, introduces material losses, thereby increasing the crystal cost owing to the high price of enriched ¹⁰⁰MoO₃.
4. **Cutting:** The ingots are cut to the CUPID crystal geometry. This step must avoid damaging the crystal or introducing radioactive contaminants, while minimizing material loss. Leftover fragments are ground back to Li₂MoO₄ powder, mixed with freshly synthesized powder, and reused in new growths. These recycling strategies are designed to recover costly enriched material.
5. **Surface quality:** Owing to the hygroscopic nature of Li₂MoO₄, polishing must be carried out with water-free methods in low humidity atmosphere. All consumables, such as pads, powders, and lubricants, must have ultra-high radiopurity to prevent surface contamination.
6. **Packaging and delivery:** Crystals must be packed in humidity-tight, vacuum-sealed, radiopure containers that also guarantee safe transport.

3.2.1 CUPID crystal pre-production

Before large-scale Li₂¹⁰⁰MoO₄ production begins, SICCAS must optimize the growth parameters to address the delicate aspects of the stages outlined above. This optimization phase, referred to as *pre-production*, has several goals; the three primary ones are:

- **Recovery Efficiency:** One of the goals is to optimize and measure the ¹⁰⁰Mo recovery efficiency of the production process. CUPID requires SICCAS to achieve a recovery efficiency of at least 90% to minimize material losses and contain costs. To this end, SICCAS is developing recycling strategies to recover waste material. Since repeated recycling may introduce or concentrate contaminants, its impact on radiopurity and on calorimetric and scintillation performance must be evaluated by CUPID.
- **Production Yield:** Li₂MoO₄ crystal growth is intrinsically challenging, and defects such as cracks or inclusions may occur, particularly if parasitic chemical species remain after powder synthesis. Pre-production therefore focuses on refining the synthesis procedure and quantifying the yield. Defective ingots can be reprocessed for subsequent growths, although such recycling negatively affects recovery efficiency.
- **Radiopurity:** CUPID documents all stages of the production process to identify those that may critically affect radiopurity. This detailed knowledge will be used to define a radiopurity protocol containing precise guidelines, such as the use of ancillaries and consumables screened for radiopurity and the cleanroom operation of selected steps. SICCAS will implement this protocol in the final stages of pre-production, which will focus primarily on radiopurity.

During the period covered by this work, SICCAS was in the early pre-production phase, with the primary emphasis on optimizing crystal growth and evaluating recycling strategies. To contain costs, predominantly natural (non-enriched) crystals were grown. CUPID tests these crystals as scintillating cryogenic calorimeters to quantify calorimetric and scintillation performance and the achieved radiopurity. The resulting measurements provide feedback to SICCAS to correlate detector response with growth parameters and the implemented recycling steps.

SICCAS has now entered a more advanced pre-production stage. In the past few months, it has focused on enriched crystal growth, producing the first batches of $\text{Li}_2^{100}\text{MoO}_4$. At present, the main effort has shifted toward radiopurity optimization, and SICCAS is implementing the radiopurity protocol developed by CUPID.

The following section gives a brief overview of the documented growth process and highlights the improvements achieved during the early stages of pre-production.

3.3 Crystal Production

3.3.1 Precursor powders and Li_2MoO_4 synthesis

The precursor powders used for the synthesis of Li_2MoO_4 must possess ultra-high radiochemical purity. A summary of the requirements on these powders is presented below. Multiple strategies are employed to verify compliance with these requirements, both by the producer and independently by CUPID on test samples. The screening techniques used include but are not limited to: ICP-MS (inductively coupled plasma mass spectrometry), high-purity germanium (HPGe) γ -spectroscopy in low-background underground labs; neutron activation analysis (NAA); and X-ray diffraction (XRD).

Li_2CO_3 powder


The technical specifications for lithium carbonate powder are a chemical purity not less than 99.99% and a radiopurity that satisfies the requirements reported in Table 3.3. Commercially available material usually meets the chemical purity standards, but radiopurity certifications are typically not provided. CUPID, in collaboration with SICCAS, is considering several suppliers. ICP-MS and HPGe assays of samples bought from selected vendors have shown that it is possible to obtain precursor material meeting the specifications.

Isotope	^{232}Th	^{238}U	^{40}K	^{235}U	^{226}Ra	^{137}Cs
Requirement [mBq/kg]	< 5	< 5	< 15	< 4	< 5	< 0.5

Table 3.3: Radiopurity requirements for lithium carbonate precursor.

Enriched $^{100}\text{MoO}_3$ powder

^{100}Mo -enriched molybdenum is produced using gas centrifuges. Natural molybdenum is first fluorinated to MoF_6 gas:



and then enriched through a cascade of centrifuges. The enriched $^{100}\text{MoF}_6$ gas is subsequently converted into $^{100}\text{MoO}_3$ via an industrial multi-step wet-chemistry process.


CUPID requires an enrichment of about 95% in ^{100}Mo , a chemical purity of at least 99.9%, and compliance with the radiopurity limits reported in Table 3.4. SICCAS has been tasked with procuring the isotope from IPC, CNNC (Institute of Physics and Chemistry, China National Nuclear Corporation), which has already produced $\mathcal{O}(10\text{kg})$ of enriched MoO_3 . The material was independently screened by CUPID using ICP-MS and NAA to verify the ^{100}Mo enrichment and the radiochemical purity. The enriched $^{100}\text{MoO}_3$ powder was found to be compliant with CUPID requirements.

Isotope	^{232}Th	^{238}U	^{40}K
Requirement	< 0.8 mBq/kg (< 0.2 ppb)	< 2.5 mBq/kg (< 0.2 ppb)	< 50 mBq/kg

Table 3.4: Radiopurity requirements for enriched molybdenum oxide precursor.

Li_2MoO_4 synthesis

Li_2MoO_4 can be synthesized from precursors using either wet chemistry or solid-state sintering. Although both methods are under study, the first batches of pre-production crystals were grown from powder prepared via wet chemistry. In this process, the raw powders are dissolved in warm distilled water at controlled temperature to form Li_2MoO_4 in solution:



Afterwards, the solution is warmed and evaporated under controlled conditions until saturation, leading to the precipitation of Li_2MoO_4 as a fine powder.

During natural crystal pre-production, SICCAS obtained a lower-than-expected yield due to unstable crystal quality in part of the growths. Investigations showed that wet chemistry synthesis occasionally produced secondary lithium molybdate phases ($\text{Li}_2\text{Mo}_2\text{O}_7$, $\text{Li}_2\text{Mo}_3\text{O}_{10}$, $\text{Li}_2\text{Mo}_4\text{O}_{13}$, $\text{Li}_2\text{Mo}_5\text{O}_{16}$) together with unreacted Li_2CO_3 . Moreover residual Li_2CO_3 , can also attack the platinum crucibles forming Li_2PtO_3 . An example of crucible attacked by unreacted Li_2CO_3 can be seen in Figure 3.1b. These byproducts negatively interfere with crystal growth, resulting in unstable quality and low production yield. XRD analyses of both the synthesized Li_2MoO_4 powder and the platinum crucibles confirmed this hypothesis.

SICCAS is refining the synthesis procedure to suppress these side products and is rapidly converging toward a more stable process. In parallel, CUPID is collaborating with IPC to explore an alternative strategy. Since the industrial process that converts $^{100}\text{MoF}_6$ into $^{100}\text{MoO}_3$ already produces $\text{Li}_2^{100}\text{MoO}_4$ as an intermediate byproduct (though with additional contaminants), it may be possible to adapt this step to directly yield ultra-pure Li_2MoO_4 powders. IPC and SICCAS are therefore working together to develop an optimized chemical treatment capable of providing Li_2MoO_4 powder of sufficient purity to be directly used for crystal growth.

3.3.2 Crystal growth

Although the details of crystal growth are beyond the scope of this work, a brief overview is provided here, highlighting the most critical aspects affecting crystal quality and radiopurity.

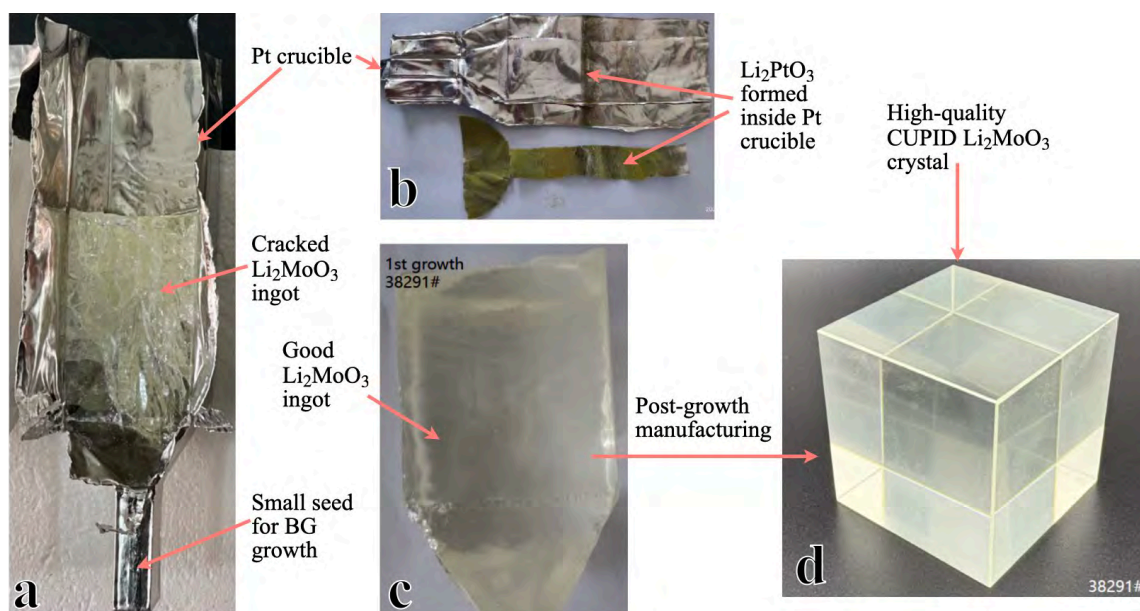


Figure 3.1: Progress in BG crystal growth during pre-production: (a) ingot grown using a small seed exhibiting cracks due to suboptimal furnace cooldown and annealing; (b) crucible where a Li_2PtO_3 layer formed during growth, caused by unreacted Li_2CO_3 in the Li_2MoO_4 powder; (c) and (d) show an ingot and the corresponding high-quality CUPID crystal obtained after resolving synthesis issues and adapting the cooldown and annealing program to the small BG seed size.

Growth techniques

SICCAS employs two standard techniques for growing Li_2MoO_4 crystals. The descriptions below are qualitative and are not intended as an exhaustive, quantitative treatment of these techniques.

- **Bridgman (BG):** a seed crystal is placed at the bottom of a high-purity platinum crucible and the Li_2MoO_4 powder is loaded above it. The crucible is then placed in a furnace with a controlled vertical temperature gradient: the upper zone is above the Li_2MoO_4 melting point and the lower zone is below it. The Li_2MoO_4 powder is fully melted while most of the seed remains solid, then the crucible is lowered slowly through the gradient. Directional solidification starts from the seed and progresses upward, yielding a single-crystal ingot having the crucible shape. The seed's material partially melts and get mixed with the rest of the crystal. Figure 3.1c shows a Li_2MoO_4 ingot resulting from a successful BG growth.
- **Czochralski (CZ):** Li_2MoO_4 powder is melted in a high-purity platinum crucible. A seed crystal is attached to a rotating platinum rod, dipped into the melt, and then slowly pulled upwards while rotating. By controlling pull rate, temperature, and rotation speed, the melt crystallizes on the seed and a cylindrical single-crystal ingot is obtained.

Both techniques are standard for oxide crystal growth and have yielded high-quality Li_2MoO_4 . To meet CUPID's throughput, BG is adopted as the baseline, since multiple crucibles can run in parallel in large furnaces. CZ is retained as an intermediate re-crystallization step. CZ growths

can also start without a seed, so they serve as a seed-generation tool. CZ-generated seeds can be used for BG growths, avoiding the mixing of legacy BG seeds with new Li_2MoO_4 powder. This is especially important when moving from natural to enriched powders, to avoid diluting the isotopic enrichment with a natural seed. For these reasons CZ is used only for intermediate re-crystallization and seed preparation.

Large-scale production thus proceeds with BG that, however, requires optimization:

1. **Annealing:** Thermal stress during furnace cooldown can induce cracks in the crystals, lowering the production yield. SICCAS is optimizing the post-growth annealing program and adjusting the cooldown profile to reduce internal stress and guarantee structural integrity.
2. **Crucible and seed size and shape:** The sizes and shapes of the crucible and the seed influence thermal stress, cutting strategy, and material loss. In BG, the seed size is non-negligible. To ensure uniform isotopic enrichment in the final ingot, the seeds must also be made of enriched material. Smaller seeds require less enriched materials, reducing loss of isotope. However, they increase stress and the probability of cracks, potentially lowering the production yield. Moreover, the crucible shape determines the cutting strategy and the resulting material losses. SICCAS is optimizing crucible geometry and seed dimensions to balance production yield and recovery efficiency.

An example of a failed BG growth due to a small seed size combined with a suboptimal furnace cooldown and annealing program is shown in Figure 3.1a. During pre-production, SICCAS is optimizing these parameters. This optimization allows to obtain good-quality BG ingots like that shown in Figure 3.1c.

Recycling and purification by multiple re-crystallizations

One aim of pre-production is to quantify how effectively successive re-crystallizations improve crystal purity. This aspect is tightly linked to recycling strategies designed to optimize recovery efficiency.

Both growth techniques (BG and CZ) purify the material, but impurity segregation differs. In CZ, impurities tend to concentrate in the residual melt left in the crucible. In BG, which solidifies the entire Li_2MoO_4 powder into a single ingot, impurities segregate toward the top of the ingot. If the seed is not perfectly pure, regions near the seed at the bottom can also be affected. The resulting BG ingot can thus exhibit non-uniform radiochemical purity along its length.

Recycling and purification (CZ) Because the final crystallization is always a BG growth, CZ ingots can be used as intermediate purification steps and as a source of seeds for subsequent BG growths. In this case, the main issue is recovery of the crucible leftover, where impurities are concentrated. To recycle this material, SICCAS dissolves it in distilled water, filters the solution, and then evaporates it to precipitate Li_2MoO_4 powder. This recycled powder is mixed with freshly synthesized powder for subsequent growths, increasing recovery efficiency. The CZ ingot itself is manually ground with mortar and pestle to a fine powder suitable for direct reuse in subsequent BG (or CZ) growths.

Recycling and purification (BG) For BG ingots the approach differs. If the growth is an intermediate one, cutting is not required. The top portion, where impurities accumulate, is dissolved in distilled water; the solution is filtered and evaporated to obtain new Li_2MoO_4 precipitate, which

is then combined with Li_2MoO_4 powder for subsequent growths, analogously to the CZ leftover route. If the seed remaining in the crucible is of sufficient quality, it is detached and sent to cutting for reuse as a seed in a new growth. The central section of the ingot, typically the purest, is manually ground with mortar and pestle to a fine powder for direct reuse in subsequent BG (or CZ) growths.

If the growth is final rather than intermediate, the crystal proceeds to post-growth manufacturing. The first operation is machining the ingot into the shape of a CUPID crystal. This involves two steps: cutting and sanding. These steps generate both macroscopic offcuts and sludge. The larger offcuts follow the same routes described above: material from the top region, where impurities are concentrated, is first dissolved in distilled water, filtered, and evaporated to recover Li_2MoO_4 powder, whereas more pure offcuts are directly ground to a fine powder for reuse. The sludge fraction, although small and likely the most contaminated, is under evaluation for recovery, including zone refining prior to multiple re-crystallizations. At the current early stage of pre-production, the available sludge is insufficient to grow full-size crystals; the outcome of these strategies will be addressed toward the end of this phase.

Given that suppliers are able to provide raw powders compliant with the requirements listed in Tables 3.3 and 3.4, SICCAS targets a baseline of two BG re-crystallizations for freshly synthesized Li_2MoO_4 powder. The number of re-crystallizations for recycled material will be decided by testing the resulting crystals in dedicated validation measurements performed by CUPID.

3.3.3 Post-growth manufacturing

After the final growth, the ingots are machined into $4.5 \times 4.5 \times 4.5 \text{ cm}^3$ cubic crystals with surfaces finished to the required specification. Each crystal is then packaged and delivered. These operations are not expected to alter bulk radiopurity or calorimetric performance, which are set by upstream production steps; they can, however, significantly affect surface radiopurity and may slightly influence light yield.

Crucible opening and annealing

After extraction from the BG furnaces, the crucibles must be opened to extract the ingots. The crucible is cut and peeled off the ingot (as in Fig. 3.1b), and various tools may come in contact with the crystal surface during this operation. The contamination risk of the final product, though, is limited because the outer surface of the ingot is cutted or sanded away and enters the recycling pathway. Radioactive contamination must nevertheless be minimized. To this end, CUPID is working with SICCAS to install a cleanroom immediately adjacent to the furnaces for crucible opening.

Cutting and sanding

Once removed from the crucible, the ingots are machined to the required $4.5 \times 4.5 \times 4.5 \text{ cm}^3$ cubic geometry in two steps: cutting and sanding.

- **Cutting.** Ingots are cut with a continuous-loop diamond wire saw: a 0.3 mm diamond-coated tungsten wire runs on two tensioned pulleys, producing a kerf of $\sim 0.4 \text{ mm}$ and using aviation-grade kerosene as lubricant and coolant. Materials used in this step are selected for

ultra-high radiopurity, although choices are constrained by the hardness of Li_2MoO_4 . This operation has a limited risk of leaving residual radioactive contaminants on the final crystal, because the outermost layer is removed by the subsequent surface finishing.

- **Sanding.** Cutting close to the ingot surfaces results in asymmetric stresses that may induce cracks. If the ingot geometry does not allow a direct cut to the final size without unacceptable crack risk, the ingot is first cut to a slightly larger cube and then sanded to the final dimensions. Sanding is performed with a dedicated machine using a pad, abrasive powder, and a water-free lubricant. This step is high risk for surface radiopurity, since consumables can embed contaminants deep into crystal surface. SICCAS is optimizing crucible shape and cutting machine parameters to avoid sanding whenever possible.

Both operations produce sludge containing $\text{Li}_2^{100}\text{MoO}_4$ mixed with consumables such as kerosene, abrasive powder, and lubricant. This sludge is recovered from the machines. Although sludge recovery is under evaluation (as discussed above), minimizing sludge is essential to maximize recovery efficiency.

CUPID is working with SICCAS to install a zoned cleanroom suite for these machining steps and the subsequent surface treatment and packaging operations described below. The facility will maintain low humidity to prevent moisture-induced damage arising from the hygroscopic nature of Li_2MoO_4 . Dedicated low-humidity storage cabinets will be installed to hold crystals between manufacturing steps.

Surface treatment

After cutting, the crystals are polished to the required surface finish. The final finish is obtained by adjusting pressure, and polishing time.

Polishing removes the outermost layer and potentially reduce contaminants introduced in earlier manufacturing stages. However, if the consumables are not radiopure, radioactive isotopes can be embedded into the outermost micrometers of the crystal, degrading surface radiopurity.

Polishing employs three elements: pads, an abrasive powder, and a water-free lubricant. Among the abrasives under evaluation, SiO_2 is particularly appealing. It was used to polish the CUORE TeO_2 crystals, and previous validation measurements showed excellent radiopurity. Preliminary tests on pre-production crystals indicate that irregular SiO_2 with an average grain size of $30\ \mu\text{m}$ achieves the required polishing efficiency. This powder differs slightly in grain size and shape from that used for CUORE and is currently undergoing radiopurity evaluation. A waterless silicone oil is foreseen as the lubricant to avoid moisture exposure during polishing. LAM PLAN TOUCHLAM® pads are compatible with SiO_2 and silicone oil, and are being screened for radiopurity evaluation.

After polishing, the crystal surfaces undergo an oil-based cleaning to remove residual abrasive and any polishing byproducts. Candidate solvents include solvents and gasoline. The choice is not yet finalized. An example of a good-quality crystal obtained after this processing is shown in Figure 3.1d.

Packaging and delivery

The final step is packaging. After cleaning, crystals must be enclosed in humidity-tight, ultra-radiopure containers that also prevent mechanical damage during transport. Packaging must be coordinated with the CUPID assembly line workflow, and the final solution is not yet defined. Two options are under evaluation:

- (i) **Individual packaging.** Each crystal is triple-bagged using three nested, humidity-tight bi-layer aluminum/plastic vacuum bags. The packaged crystals are then placed in vacuum-sealed boxes with polystyrene supports to prevent movement during transport. The triple layering allows sequential removal of the outer bags when entering the CUPID assembly-line cleanroom in LNGS.
- (ii) **Multi-crystal containers.** Multiple (~ 4) crystals are placed in an engineered vacuum- and humidity-tight container with fixtures that immobilize them for safe transport. The container is then triple-packaged in vacuum bags.

Option (i) involves more complex handling of each crystal, and the entire crystal surface is in direct contact with the vacuum bag. Both factors increase the probability of surface re-contamination. Its advantage is simpler immobilization during transport: external supports can contact only the outer bag to keep each crystal stable.

Option (ii) simplifies handling, since crystals are placed once into a dedicated container whose fixtures minimize contact points with the crystal surface. The trade-off is engineering complexity: custom containers must immobilize the crystals for safe intercontinental transport, provide shock protection, and ensure vacuum and humidity tightness. In addition, every material used in the internal fixtures must be radiopure to avoid introducing contaminants into the CUPID assembly-line cleanroom when the box is opened.

3.3.4 Crystal validation strategy

Because material recovery, growth, and surface treatment are still being optimized, each crystal has a slightly different history in pre-production. A rigorous, uniform protocol is therefore needed to validate and compare the crystals and to correlate growth parameters with measured performance. Chapter 4 presents the development of the CCVR program (CUPID Crystal Validation Runs) together with the measurement of the first batch of natural crystals produced by SICCAS during pre-production. Establishing and validating the CCVR protocol before the end of pre-production is essential, since the first crystals grown from recycled sludge material will be available only at the end of this phase, leaving a limited window for qualification before the start of large-scale production. Complementary ICP-MS, NAA, HPGe γ -spectroscopy, and XRD analyses will be performed on samples extracted from the ingots. Finally, because crystallization, annealing, and impurities can introduce lattice defects that may impact scintillation performance, Chapter 5 presents measurements that can be used to identify these defects and assess their impact on photophysical and scintillation properties.

3.4 Conclusion

Pre-production is a critical phase for CUPID. It lays out the key steps of the $\text{Li}_2^{100}\text{MoO}_4$ production chain on which SICCAS can act to meet the stringent requirements on calorimetric performance, scintillation performance, and radiopurity required by a large-scale next-generation $0\nu\beta\beta$ experiment. By documenting these steps, CUPID strengthens communication with SICCAS and establishes an effective feedback loop that drives the process toward the required crystal quality.

Chapter 4

Validation of Pre-Production Li_2MoO_4 Crystals

The success of a next-generation $0\nu\beta\beta$ experiment such as CUPID depends critically on the reproducibility and quality control of all its core components. Among these, the $\text{Li}_2^{100}\text{MoO}_4$ scintillating crystals play a central role. Given the scale and sensitivity of CUPID, ensuring the uniformity, radiopurity, and performance of all 1596 crystals is essential. To guarantee these requirements, CUPID is developing a comprehensive validation strategy that builds upon the experience of its predecessor, CUORE, while adapting it to the peculiarities of isotopically enriched lithium molybdate, which introduces new technical and logistical challenges. This chapter outlines the development of this validation program and presents the measurements performed on the first batch of Li_2MoO_4 crystals produced by SICCAS during the pre-production phase. The resulting protocol will serve as the basis for subsequent pre-production measurements as well as the full-scale production campaign.

4.1 The CUPID Pre-CCVR Campaign

CUPID must establish a crystal validation protocol to verify compliance with experimental requirements once large-scale production begins. This protocol, named *CUPID Crystal Validation Runs* (CCVRs), will be applied to a random subset of crystals to assess conformity with the specifications listed in Sec. 3.1. If any non-compliant crystal is identified, production will be suspended until the nonconformity is understood and corrected. This approach mirrors the procedure adopted in CUORE, where the *CUORE Crystal Validation Runs* [152, 153] served as essential diagnostic tools based on cryogenic measurements of randomly selected TeO_2 crystals. Had a crystal failed to meet the radiopurity or performance criteria, production would have been halted until the issue was resolved. The success of CUORE demonstrated that such a rigorous validation strategy effectively ensures detector quality.

In contrast to CUORE, which employed natural tellurium dioxide crystals supported by a mature and reliable industrial growth process (as TeO_2 has many non-academic purposes), the production of enriched Li_2MoO_4 crystals is intrinsically more challenging. As discussed in Chapter 3, SICCAS must establish a stable, reproducible process that minimizes costly isotope losses while maintaining radiopurity at every stage—objectives that are often in tension with one another. Consequently,

whereas CUORE crystal validation focused primarily on radiopurity, CUPID must devote equal attention to calorimetric performance and light yield during pre-production.

The pre-production phase provides the ideal framework for developing the CCVR protocol. In this stage, SICCAS is setting up a high-throughput production line capable of delivering 1596 high-quality crystals (Chapter 3), while CUPID operates these crystals as scintillating cryogenic calorimeters under experiment-like conditions to supply systematic feedback. The measurements performed on pre-production crystals thus serve a dual purpose: they provide SICCAS with detailed performance data and allow CUPID to construct and refine the CCVR protocol. This set of measurements is referred to as the *pre-CCVR* campaign.

4.1.1 Crystals in the early-stage pre-production phase

During pre-production, SICCAS grows both natural (Li_2MoO_4) and enriched ($\text{Li}_2^{100}\text{MoO}_4$) lithium molybdate crystals, all with the same dimensions as those planned for the final detector ($45 \times 45 \times 45 \text{ mm}^3$). Natural crystals differ from enriched ones only in the MoO_3 precursor used for Li_2MoO_4 synthesis, which is much less expensive than isotopically enriched $^{100}\text{MoO}_3$. Consequently, natural crystals represent the only cost-effective option for the initial optimization of growth parameters.

This chapter presents measurements on the first batch of natural crystals produced by SICCAS. At this early stage, SICCAS is demonstrating the feasibility of Li_2MoO_4 crystal growth while refining several aspects of the process, such as crucible geometry, BG growth parameters, the number of CZ/BG re-crystallizations, and the annealing temperature profile. Initial material-recycling strategies are also tested to improve recovery efficiency. For these reasons, radiopurity is not prioritized during this early phase.

The raw powders employed at this stage are not the final and do not satisfy the radiopurity specifications listed in Tables 3.3 and 3.4: the MoO_3 is not enriched and the Li_2CO_3 supplier has not been decided yet; moreover, the Li_2MoO_4 synthesis procedure is still under development (see Sec. 3.3.1). For these reasons, crystals in the early stages of pre-production are not expected to meet the requirements of Table 3.2. The main goal is therefore to characterize calorimetric performance and light yield, while defining a CCVR protocol. Although these crystals are not expected to be radiopure, radiopurity measurements remain valuable for relative comparisons between crystals, for studying the effects of re-crystallization and material recycling on contamination, and for initiating the radiopurity assessment component of the CCVR program.

4.1.2 Pre-CCVR measurement campaigns

Multiple pre-CCVR measurements have been conducted during this early-stage of the pre-production phase. These runs are performed in a low-background cryogenic facility located at the underground laboratory *Laboratori Nazionali del Gran Sasso*, using the same cryostat employed for the *CUORE Crystal Validation Runs* described in Ref. [153]. Details of the facility are provided in Sec. 4.2, and the analysis strategy used to process the data is described in Sec. 4.3. This chapter presents the following data campaigns, each with distinct objectives:

- **Pre-CCVR1** is presented in Sec. 4.4. This run involves the measurement of four $\text{Li}_2^{100}\text{MoO}_4$ crystals produced by the Nikolaev Institute of Inorganic Chemistry (NIIC), the same manufacturer that supplied the crystals for the CUPID-Mo small-scale demonstrator [150]. This run has two goals. First, characterizing the behavior of crystals from a known producer provides a standard reference for direct comparison with crystals produced by SICCAS. Second,

testing known crystals coupled to LDs enables characterization of the full system, including noise assessment. In this regard, this run also served to recommission the cryogenic setup following major hardware interventions on the cryogenic system.

- **Pre-CCVR2** is presented in Sec. 4.5. This run tests four natural Li_2MoO_4 crystals produced by SICCAS using different crystallization techniques and surface-treatment procedures. For comparison, two NIIC crystals from pre-CCVR1 were included. The measurements were conducted in three successive configurations, each featuring different thermalizations of the detectors to the heat sink.

The first batches of isotopically enriched $\text{Li}_2^{100}\text{MoO}_4$ crystals from SICCAS have recently arrived at LNGS, and the first measurement on enriched crystals (pre-CCVR3) is currently in progress; accordingly, no results are yet available and it is thus not presented in this chapter. The light detectors used in all of these runs are fabricated from 500 μm -thick high-purity Ge wafers equipped with Ge NTD thermisors coated with a 60 nm-thick SiO_2 anti-reflecting layer. No NTL assisted LDs were used in these measurements. More details about the geometry of the detectors and the assembly will be given in each section.

4.2 The Hall C Cryogenic Infrastructure for CUPID CCVR

The CCVR measurements are carried out in Hall C of the *Laboratori Nazionali del Gran Sasso* (LNGS), located in central Italy. This is the same underground laboratory that will host the CUPID experiment, which is planned to be installed in Hall A, in the same infrastructure currently used for CUORE.

4.2.1 Low-background and low-noise Hall C cryostat

Suppressing cosmic rays is a fundamental requirement for studying the bulk and surface radiopurity of $45 \times 45 \times 45 \text{ mm}^3$ Li_2MoO_4 crystals. At LNGS, the *Gran Sasso* mountain provides an overburden of approximately 1300 m of rock, which reduces the surface muon flux by about six orders of magnitude [138].

The cryogenic setup used for CCVR is housed in a dilution cryostat. More details about the working principles of dilution cryostats are given in [Appendix A](#). The Hall C cryostat is equipped with extensive external shieldings to suppress environmental backgrounds. Specifically, it includes an external 5 cm-thick lead shield and an inner 5 cm-thick layer of copper to attenuate ambient γ radiation; additionally, a layer of polyethylene pellets to moderate neutrons is also present. These external shieldings are enclosed in a plastic container that fully encapsulates the cryostat, which is continuously flushed with nitrogen gas to prevent radon accumulation, which could introduce additional γ backgrounds. A layer of lead and copper bricks are placed above the cryostat to shield γ radiation coming from above, providing close to 4π coverage. As shown in [Fig. 4.1](#), internal lead shielding positioned beneath the cryostat’s mixing chamber (MC) covers the experimental region, further suppressing background radiation from above as well as from nearby components, such as the Printed Circuit Boards (known for their lack of radiopurity) that connect the detector’s readout electronics to the cryostat wiring.

The cryostat is an Oxford Instruments TL 200 wet $^3\text{He}/^4\text{He}$ dilution refrigerator, shown in [Fig. 4.1](#).

A wet cryostat offers an important advantage over dry (cryogen-free) cryostats that rely on pulse tubes: while pulse tubes are known to introduce mechanical vibrations into the system [167], resulting in increased noise and degraded SNR, wet cryostats are free from this specific source of mechanical disturbance. However, traditional wet cryostats require periodic refilling with liquid helium, which can interrupt data-taking and reduce live-time. Additionally, occasional liquid helium shortages could delay data-taking, which is particularly undesirable for a setup intended to provide swift feedback to the producer regarding the quality of newly produced crystals, enabling prompt corrective actions when necessary.

The cryogenic facility in Hall C mitigates these limitations through the use of a dedicated liquefier, positioned near the cryostat, that continuously condenses helium gas into liquid which supplies the main bath. To minimize the transmission of mechanical vibrations, the liquefier is mechanically decoupled from the cryostat. It operates with Gifford-McMahon (GM) cryocoolers, which, like pulse tubes, exhibit a \sim Hz periodic motion. However, unlike pulse tubes, GM heads are not mechanically connected to the cryostat, significantly reducing the propagation of vibrations to the detectors.

Besides the liquefier, other components of the cryogenic infrastructure can also contribute to vibrational noise. Notably, the 1K pot is a well-known source of such disturbances [175]. These noise sources typically produce a broad harmonic spectrum that result in a $1/f$ component in the noise power spectrum of the detectors, which can strongly affect the performance of cryogenic calorimeters. To further suppress vibrational noise from the liquefier, cryogenic components, and other sources, a dedicated mechanical double stage damping system was developed for the Hall C cryostat during the CUORE R&D program and it is still in use today [175, 176]. These features make the Hall C cryogenic facility a suitable environment for conducting the CUPID CCVR validation program. All the pre-CCVR measurements presented in this chapter were acquired using this infrastructure.

4.2.2 Readout and biasing electronics

Section 2.3.1 provides a detailed description of the readout and biasing electronics used for cryogenic calorimeters equipped with NTD thermistors. The biasing circuit, shown in Fig. 2.5, generates the bias current I_{bias} through a low-noise DC power supply and a pair of identical load resistors with total impedance R_L . The power supply used to generate V_{bias} is the HP 6627A low-noise DC source [177], capable of delivering up to 20 V per output channel. Two output channels in a differential configuration are used to provide a bias voltage up to 40 V. A voltage divider is used to finely select




Figure 4.1: The Hall C cryostat used for the CUPID CCVR validation program.

V_{bias} in the range 0–39 V.

The front-end electronics, described in Refs. [131, 132], consists of three chassis, each containing six boards with two channels per board, for a total of 32 channels. Each board includes two selectable load resistors. One set of boards allows choosing between $R_L = 54 \text{ G}\Omega$ and $R_L = 4 \text{ G}\Omega$, while another set offers the choice between $R_L = 54 \text{ G}\Omega$ and $R_L = 10 \text{ G}\Omega$. Combined with the voltage divider, this configuration allows for a wide range of bias currents, ranging from $I_{\text{bias}} \sim 1 \text{ pA}$ to $I_{\text{bias}} \sim 10 \text{ nA}$, enabling flexible selection of the detectors' WP, as introduced in Sec. 2.3.1. The total gain, configurable via the preamplifier and the PGA, can be set between 448 and 10009.

The front-end and back-end electronics used in Hall C are shown in Fig. 2.7. The Bessel anti-aliasing filter consists of multiple boards, each offering a set of selectable cut-off frequencies. Three types of boards have been employed, providing the following cut-off frequency options: [16, 32, 47, 63] Hz, [16, 38, 47, 70] Hz, and [80, 120, 160, 200] Hz. The latter set, featuring higher cut-off values, is required for the LDs, which have faster signal rise times and thus broader bandwidths compared to the heat channels.

Room-temperature, voltage-sensitive differential JFET amplifiers are used to amplify the signals from the detectors. The amplifiers used in this setup exhibit a white noise level of approximately 3.2–3.3 nV/ $\sqrt{\text{Hz}}$ and a $1/f$ noise component of about 7.8 nV/ $\sqrt{\text{Hz}}$ at 1 Hz. The entire cryostat, along with the front-end boards, is enclosed in a Faraday cage to suppress electromagnetic interference and improve noise performance. The Bessel filter is housed in a separate, dedicated Faraday cage located outside the main one. The subsequent back-end electronics is unshielded.

Waveforms from the JFET amplifiers are digitized using three National Instruments PCI/PXI-6284 DAQ boards. Each board features a 18-bit 16 channels ADC capable of sampling at a total rate up to 16 kHz. When all channels are used simultaneously, the maximum achievable sampling rate is limited to 1 kHz per channel, which may be insufficient for LDs due to their faster time constants which results in a higher signal bandwidth. To accommodate this, fewer channels are read out per board, allowing LDs to be sampled at higher rates, typically 2 kHz or 4 kHz. In all the measurements presented in this chapter, heat channels were acquired with a sampling frequency of 1 kHz. The digitized waveforms are saved in a PC for offline analysis. Except for the anti-aliasing filter, all further shaping and filtering are performed digitally on the stored waveforms, as explained in Sec. 4.3.

4.3 Analysis Pipeline

To obtain excellent energy resolution spectra from raw waveforms (like the one shown in Fig. 2.10), several processing steps are required. In dual channels setups, such as scintillating cryogenic calorimeters, the LD signals also need to be processed to yield light-vs-heat scatter plots, like the ones sketched in Fig. 2.13.

This section provides an overview of the data processing steps used to analyze the measurements shown in the present chapter as well as in Chapter 6.

4.3.1 Digital signal processing

Digital signal processing is employed to trigger and filter the digitized waveforms. Except for the triggering stage, which can be performed both online and offline, all processing steps are executed offline. To ensure consistent and reliable results, pre-CCVR data are processed using APOLLO and DIANA, object-oriented C++ frameworks for digital signal processing that were originally

developed for the CUORE experiment [147, 178, 179, 103, 116], and later adapted for use with dual-channel scintillating cryogenic calorimeters by LUCIFER [180] and CUPID-0 [149]. The same tools were also used to process data acquired by CUPID-Mo [181].

DIANA is built on top of ROOTv5 [182] and integrates a PostgreSQL database [183], which enables tracking of detector configurations and electronics settings. This framework supports analyses involving coincidence data across multiple channels. The use of validated software like DIANA ensures robustness and reproducibility of the results, reducing the risk of biases related to the data processing strategy, an essential feature when assessing crystal quality during the pre-CCVR campaign. On the other hand, DIANA was primarily developed for large-scale detector arrays characterized by stable and reproducible performance. As a result, it may not natively handle unforeseen events or non-standard conditions that can arise in R&D contexts like pre-CCVR. In such cases, the development of additional custom analysis modules may be required. DIANA’s object-oriented and modular design facilitates the implementation of such extensions when needed (as discussed in Sec. 4.5.9, where customized cuts to reject cross-talk events had been introduced).

4.3.2 Treating signals from a dual-channel heat-light detector

Before describing the analysis pipeline, it is useful to outline how the heat and light channels are handled in scintillating cryogenic calorimeters with dual readout. Both heat and light channels employ NTD thermistors and operate as independent cryogenic calorimeters (as introduced in Sec. 2.4.1). The LD measures scintillation light from the facing Li_2MoO_4 crystal(s), with photons (~ 2.1 eV/photon¹) absorbed in the Ge wafer and converted to thermal pulses. Importantly, these pulses are indistinguishable from those produced by direct ionizing interactions hitting the Ge wafer. Under this working assumption, the LD is treated as an independent calorimeter with its own energy spectrum comprising both scintillation and ionization events.

Direct ionization events can be used for multiple purposes. They provide monochromatic peaks for calibrating scintillation light-signal amplitudes and increase granularity for coincidence analyses. For example, in a surface α decay within the outermost nanometers of an Li_2MoO_4 crystal, a nuclear recoil may hit the LD while the α enters the crystal, yielding a characteristic coincidence useful for identifying surface backgrounds. Consequently, the same analysis pipeline is applied to both heat and light channels, as discussed in Sec. 4.3.3.

After calibration from ionization peaks, scintillation signals associated with events in the facing Li_2MoO_4 are reconstructed. These signals are small—e.g., a 2615 keV γ typically yields ~ 750 eV in the LD—and are extracted with a dedicated highly-optimized coincidence analysis based on the optimum filter, which will be described in Sec. 4.3.4.

4.3.3 From raw waveforms to the energy spectrum

The analysis pipeline described in this section is applied to both heat and light channels. Its primary goal is the accurate reconstruction of the energy of each event while filtering away poor-quality events. An overview of the analysis pipeline is given in the flow chart shown in Fig. 4.2.

Trigger and window definitions Two types of triggers operate concurrently online on digitized waveforms during data acquisition: a periodic trigger and a derivative trigger. Each trigger generates a flag used to define windows on the raw waveforms. Each window includes fixed-duration

¹The average energy of a scintillation photon in Li_2MoO_4 is discussed in Sec. 5.4.3.

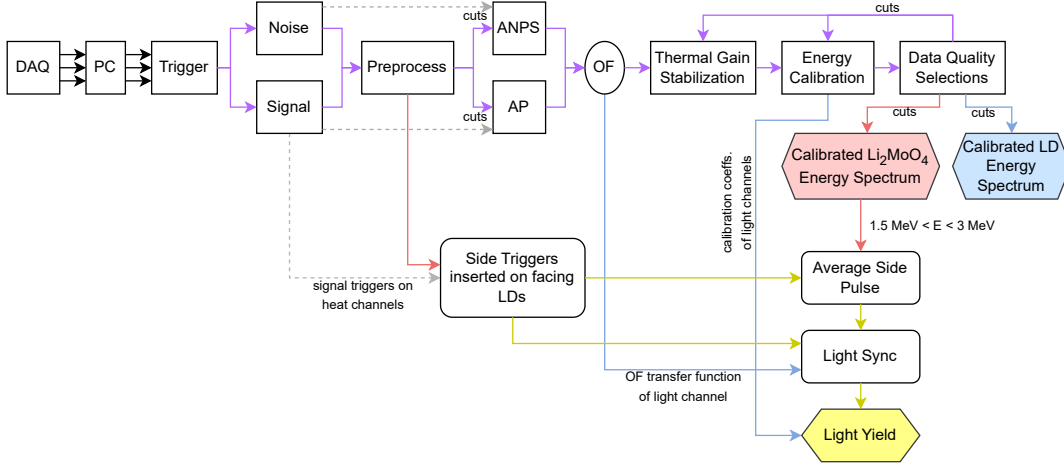


Figure 4.2: Flow chart illustrating the analysis pipeline. Violet arrows indicate steps applied to both the heat and light channels, where each is treated as an independent cryogenic calorimeter, as explained in Sec. 4.3.3. The red (blue) arrows represent processing steps on windows triggered only on the heat (light) channels. Yellow arrows indicate the analysis steps focused on heat-light coincidences, happening on the side trigger flags, as detailed in Sec. 4.3.4.

pre-trigger and post-trigger intervals. Heat channels typically have windows of approximately 6 s, with a pre-trigger of around 2 s. Due to their faster time response, light detector channels typically have 1 s windows with a 500 ms pre-trigger.

Periodic trigger flags, typically occurring at rates between 0.01 Hz and 0.1 Hz, are combined with later selections to identify pure noise windows, containing only noise without thermal pulses. These noise-only windows are used to characterize the noise affecting each channel using Fast Fourier Transform (FFT) [184]. For this reason, the periodic trigger is named *noise trigger*.

The derivative trigger algorithm exploits the fast rise time of thermal pulses to select windows containing pulses (see for example Fig. 2.4a). The threshold for the derivative trigger is set individually for each channel. One of the main limitations of the derivative trigger is that it is highly sensitive to high-frequency noise, potentially worsening the energy threshold². The derivative trigger is named *signal trigger*.

Preprocessing Once data acquisition stops, raw waveforms and trigger flags are processed using algorithms designed to compute descriptive parameters for both noise and signal-triggered windows. Examples of these parameters include rise time (RT), decay time (DT) or amplitude of signal pulses; in the pre-trigger region the waveform $V(t)$ RMS (*root mean square*) and average slope are also evaluated. These parameters are used to reject spurious events, such as derivative-triggered noise spikes and accidental noise triggers in random coincidence with genuine thermal pulses. In addition, a multi-parameter algorithm estimates the number of pulses within each triggered window

²An alternative trigger, less sensitive to this limitation, utilizes the optimum filter described later. This algorithm is called optimum trigger [144]. Needing dedicated optimization, it can practically be used only if strictly required.

to identify and reject pile-up events.³ Selecting pure signal and pure noise windows is essential for the subsequent stages of the analysis pipeline.

Another goal of preprocessing concerns heat-light coincidence events. Every event in the Li_2MoO_4 scintillating crystal coincide with an event in the associated light detector. Therefore, when a heat channel detects a signal event, an additional trigger flag is inserted into the corresponding light detector channels facing the Li_2MoO_4 crystal during the preprocessing phase. This constitutes a third trigger type with respect to signal and noise triggers introduced earlier and it is named *side trigger*. Events flagged by side triggers (side events) allow the computation of parameters such as LY and DP, as further detailed in Sec. 4.3.4.

Optimum filter To reconstruct the energy of each event accurately, triggered waveforms are filtered using the *optimum filter* (OF), a matched filter which maximizes the SNR [146]. This filter relies on the knowledge of the frequency content of both signals and noise in each channel. To construct the transfer function of the OF, two components are required: a template for the noise power spectrum (NPS) and a template for the signal. Under the key assumptions that (i) the noise remains stationary during data acquisition, and (ii) all the signals have the same shape, varying only in amplitude (i.e., all signals share the same frequency content), one can show that:

$$(\text{SNR})^2 \propto \frac{V(t_{\max})^2}{\int_{-\infty}^{+\infty} N(\omega) d\omega} = \frac{\left(\frac{1}{2\pi} \int_{-\infty}^{+\infty} V(j\omega) e^{j\omega t_{\max}} d\omega\right)^2}{\int_{-\infty}^{+\infty} N(\omega) d\omega} \leq \frac{1}{2\pi} \int_{-\infty}^{+\infty} \frac{|V(j\omega)|^2}{N(\omega)} d\omega \quad (4.1)$$

where $V(t)$ denotes the signal waveform, having amplitude $V(t_{\max})$, with t_{\max} being the time at which the signal waveform reaches its peak and $N(\omega)$ is the NPS. The last inequality results from applying the Schwarz inequality⁴, after multiplying and dividing the integrand in the numerator by $\sqrt{N(\omega)}$. Intuitively, this implies that the SNR can only be worse than the signal spectrum weighted by the inverse of the noise power at each frequency interval $d\omega$. The maximum SNR is achieved when equality holds in the final inequality. To obtain this maximum, one can imagine a filter $H(j\omega)$ that enforces equality in the expression above. This filter, known as the optimum filter, has the following transfer function:

$$H(j\omega) = K \frac{V^*(j\omega)}{N(\omega)} e^{-j\omega t_{\max}} \quad (4.2)$$

where K is a normalization constant, typically chosen so that the filter has unity gain. While a matched filter such as the OF is typically difficult to implement in analog form, the digitized windows obtained through the triggering method described earlier can be used to implement a digital version of the optimum filter transfer function given in Eq. (4.2).

The NPS template is obtained by averaging the FFTs of multiple noise-only waveforms, which are selected using the periodic noise trigger. The resulting template, called ANPS, is not only a key ingredient for constructing the OF but also a useful tool for inspecting the noise affecting each channel. To construct a representative signal template, it is essential to isolate genuine signal frequencies without contributions from the noise. Averaging multiple triggered waveforms in the time domain effectively averages away noise with random phases⁵, yielding an *average pulse* (AP)

³Only pile-up with inter-arrival times larger than the pulse rise time can be reliably identified by this algorithm; closer pulses are effectively unresolved.

⁴The Cauchy–Schwarz inequality states that $(\int f(x)g(x) dx)^2 \leq \int |f(x)|^2 dx \int |g(x)|^2 dx$.

⁵This is valid under the assumption that signal triggers are completely uncorrelated with noise phases.

unaffected by noise frequencies. The FFT of this AP provides a clean, representative template of the signal’s frequency content. This method for selecting the signal and noise templates assumes stationary noise conditions at each given detector’s WP⁶.

Since the OF operates in the frequency domain on finite waveform windows, misalignment at the window boundaries can introduce artifacts into the filter transfer function. To mitigate this issue, each waveform window is multiplied by a gating function with exponential rise and decay tails, which dampens boundary oscillations. This algorithm is named *windowed optimum filter* (wOF). All measurements presented in this thesis employ the wOF algorithm to estimate parameters such as amplitude (wOF_Amplitude), rise time (wOF_RiseTime), decay time (wOF_DecayTime), and χ^2 of each window relative to the AP template (wOF_Chi2).

Constructing the wOF with representative signal (AP) and noise (ANPS) templates is a critical step in data production. This process involves introducing careful selections based on parameters calculated during preprocessing, to identify waveform windows containing pure noise or isolated signals. Windows with unwanted features, such as pulse pile-ups, noise triggers containing a random coincidence with a signal pulse or with unstable baselines, must be excluded. For instance, waveforms with nonzero baseline slopes could artificially introduce spurious $1/f$ noise, degrading the optimum filter’s performance and worsening the SNR and thus the energy resolution. Great attention has been devoted to the definition of these selection criteria, which are crucial for accurate energy reconstruction in the analyses presented in this Chapter and in Chapter 6.

Before continuing with the analysis pipeline, it is worth emphasizing the deep physical meaning of the inequality on the left-hand side of Eq. (4.1). This inequality tells that only the noise components overlapping with the signal in the frequency domain affect the SNR. In other words, *noise frequencies that lie outside the signal bandwidth do not degrade the detector’s energy resolution* after the OF is applied. As discussed in Sec. 2.3.1, the choice of the WP changes the response of the detector. Consequently, both the signal shape and the noise power spectrum vary with the selected WP. This highlights the importance of choosing the correct WP, not simply to maximize signal amplitude as one might incorrectly infer from the plot shown in Fig. 2.6a, but rather to optimize the SNR. The key criterion should be to maximize the spectral decoupling between signal and noise while maintaining a sufficiently strong signal amplitude, leaving all the filtering work to the OF. Once this condition is met, the optimum filter can effectively suppress noise components that fall outside the signal band, thereby improving energy resolution. The choice of the WP is often non-trivial, and selecting a wrong WP can significantly degrade the detector’s energy resolution, undermining one of the key advantages of cryogenic calorimeters.

A final consideration is that the entire detector should be designed with the inequality expressed in Eq. (4.1) in mind. In fact, this equation provides a clear prescription for how the system should be built: the detector design should aim to decouple noise frequencies from those of the signal. For example, high-frequency noise is usually less relevant for heat channels, which have slower time constants, because the optimum filter can effectively suppress those frequencies. However, light detectors, which have faster time constants and therefore broader signal bandwidths, are much more sensitive to high-frequency noise, which can impact their energy resolution. At the same time, low-frequency noise, such as $1/f$ shot noise or mechanical vibrations of the cryostat (e.g., from pulse tubes, GM heads, 1K pot vibrations, or other sources), overlaps with the signal bandwidth of the heat channels as well. If not mitigated directly at the hardware level, this overlap

⁶As explained in Sec. 2.3.1, changing the WP modifies the detector’s bandwidth, influencing both the spectral shape of the noise and the characteristic time constants of the pulses.

makes the OF ineffective, degrading the detector’s energy resolution.

Thermal gain correction By combining Eq. (2.6) with Eq. (2.7), it becomes evident that even slight variations in the NTD temperature (e.g., due to temperature instabilities of the MC leading to fluctuations in the heat sink temperature) result in changes in the (intrinsic) thermal gain of the calorimeter. These gain variations manifest as drifts in the amplitude of the voltage pulse for a fixed energy deposit in the absorber, deforming the energy spectrum.

From Eq. (2.8) and Eq. (2.9), it follows that the mean value of V_{NTD} in the absence of a thermal pulse can serve as a proxy for the NTD temperature, and, assuming good thermal coupling with the absorber, for the calorimeter temperature. This mean value is estimated using the first three quarters of the pre-trigger region of each signal window and is referred to as the *Baseline*. By convention, the NTD is always reverse biased, so that an increase in the Baseline corresponds to an increase in temperature.

With this convention, combining Eqs. (2.6) and (2.7) with Eqs. (2.8) and (2.9) one can show that the amplitude of a given energy deposit in the calorimeter is anti-correlated with the Baseline value. As a result, monochromatic energy deposits trace a drift in the (Amplitude, Baseline) plane. This temperature-induced drift in pulse amplitude must be corrected to recover the true energy of the incoming particles.

This correction procedure is referred to as *thermal gain correction* or *stabilization*. It involves identifying one monochromatic line in the (Amplitude, Baseline) plane and applying the following transformation to the amplitudes of all the events:

$$\text{StabAmplitude} = k \cdot \frac{\text{Amplitude}}{p_0 + p_1 \cdot \text{Baseline}} \quad (4.3)$$

Here, StabAmplitude is the stabilized amplitude, k is an arbitrary rescaling constant, and p_0 and p_1 are parameters obtained by fitting the drift of the chosen monochromatic line.

Although some literature informally describe this as a *rotation* in the (Amplitude, Baseline) plane [135], Eq. (4.3) does not implement such a transformation. In fact, a rigid rotation would not work: as implied by Eqs. (2.6) and (2.7), monochromatic lines at different energies drift with different slopes in the (Amplitude, Baseline) plane under very small heat sink temperature variations (approximately behaving as a pencil of lines with vertex lying close to the X-axis). Thus, a rigid rotation could correct one line while over- or under-correcting others. In contrast, the transformation in Eq. (4.3) enables stabilization of the entire energy spectrum using the drift of a single monochromatic line at arbitrary energy.

If the absorbers are equipped with Si heaters, artificially injected monochromatic pulses can be used for stabilization. In their absence, monochromatic γ or α lines can also be employed. The specific stabilization strategy used in each case will be detailed in the relevant sections.

Energy calibration and thermal quenching factor Before thermal gain correction, amplitude spectra are expressed in units of millivolts. After stabilization StabAmplitude spectra in arbitrary units are obtained. The final step in energy reconstruction is the energy calibration, which consists in identifying known features in the spectra, such as γ or X-ray lines, and computing a calibration function. A linear function is generally sufficient, although higher-order polynomials may be required in the presence of non-linearities. These non-linearities can arise from the choice of the optimal WP, particularly when it lies beyond the inversion region, leading to a non-linear response over a broad energy range (see Sec. 2.3.1).

Due to the very different nature and purposes of the absorbers, light detectors (LDs) and heat channels are calibrated using different strategies, each with its own caveats:

- **Light channels:** Calibration is typically performed using ^{55}Fe sources deposited on supports facing each LD. These internal sources produce calibration X-ray lines from the decay of ^{55}Fe , namely Mn- K_α at 5.89 keV and Mn- K_β at 6.49 keV [185]. In the absence of internal ^{55}Fe sources other calibration strategies based on Cu or Mo X-rays may be employed [181]. A linear calibration is applied to all triggered events, including both signal and side trigger flags. This allows the measurement of energy released in the form of scintillation light following particle interactions in the facing Li_2MoO_4 crystal, under the assumption that the X-ray calibration can be extrapolated down to ~ 0.01 -1 keV and that thermal pulses generated by scintillation light are indistinguishable from those of ionization events. In the absence of additional higher-energy lines, this calibration can be extrapolated to estimate the energy of higher-energy events, such as nuclear recoils around ~ 100 keV. However, this extrapolation is typically unreliable. This is because LDs are designed to work in a restricted dynamic range and are optimized for low-energy interactions so they may exhibit a nonlinear response at higher energies.
- **Heat channels:** Calibration is performed during dedicated runs using external ^{232}Th sources, deployed between the external lead shielding and the OVC. These sources provide high-intensity γ lines, such as the ^{208}Tl line at 2614.5 keV. When present, internal contaminants (such as ^{40}K) can provide additional γ peaks that can be used for calibration purposes. While a linear function is often sufficient, in some cases a quadratic calibration may be required. These low-degree polynomials are enough to calibrate the detector for β/γ interaction in a wide dynamic range, typically extending from ~ 100 keV up to ~ 3 MeV. Nonetheless, neither a quadratic nor any low-degree polynomial function is adequate to calibrate the full energy spectrum up to ~ 10 MeV which includes events from β/γ , μ , α particles, nuclear recoils (NR), or a combination of them in coincident events. Importantly, the observed deviations from a single calibration curve are not attributed to residual nonlinearities in the detector's response function. Rather, they are believed to be an intrinsic feature of the detector: the interaction mechanisms contributing to the signal depend on the type of incident particle, leading to different fractions of the deposited energy being converted into phonons. Consequently, the amplitude of the signal depends on the particle type. To account for this effect, the thermal quenching factor (QF)⁷ is introduced. The QF is defined as the ratio between the amplitude of the signal induced by a given particle type and that produced by a β/γ particle of the same energy. In an ideal, non-scintillating cryogenic calorimeter, the QF is expected to be unity, as all deposited energy is converted into phonons. Under these conditions, no energy escapes the detector, and the thermometer provides an accurate measurement of the total phonon energy. However, such ideal behavior is rarely achieved in practice. Deviations from unity have been observed and, in some cases, accurately measured for a variety of crystals, including both scintillating [186, 187, 156] and non-scintillating [188, 189] compounds. In the case of Li_2MoO_4 scintillating crystals, a QF greater than unity has been consistently observed for α particles and Q-value (α +NR) events in several R&D campaigns over the past decade [190, 173, 181,

⁷In the literature, the term QF is sometimes used to describe the light quenching factor, defined as the light yield of a given particle relative to that of a β/γ of the same energy. According to Birks' law [133], this factor is typically less than one for particles with higher stopping power (dE/dx). In this thesis, the term QF will always refer to the thermal quenching factor instead.

[191]. A quenching factor $QF > 1$ indicates that energy deposited by α particles is converted into phonons more efficiently than energy deposited by β/γ interactions. In other words, it is the β/γ events that appear *quenched* relative to the α particles or Q-value lines. This observation is consistent with the fact that a larger fraction of the energy in β/γ interactions is emitted in the form of scintillation light. However, the observed values of the QF cannot be fully accounted for by light emission alone, suggesting that additional unknown mechanisms may contribute to the reduced phonon yield in β/γ interactions. To date, a conclusive study aimed at precisely describing and quantifying this effect for Li_2MoO_4 crystals is still lacking in the literature. As a consequence, reconstructing the energy of α and Q-value lines requires a separate calibration based on the identification of peaks from these event classes. Therefore, the QF must be re-evaluated for each measurement through the identification and analysis of these peaks.

Data quality cuts The wOF algorithm computes the filtered amplitude along with additional parameters that characterize the pulse shape. Pulses triggered by the derivative trigger and exhibiting shapes that deviate significantly from the AP, or with unnatural rise or decay times, are typically fake signals caused by disturbances in the electronics. A common source of such spurious signals is improper grounding of the electronics or cross-talk with diagnostic devices used to monitor the cryogenic system, such as liquid helium level meters in the main bath or the 1K pot. Cross-talk from neighboring channels can also induce spurious pulses. If not removed from the energy spectra, these events can introduce artifacts or obscure genuine peaks. The rejection of events clustering outside the physical distribution or identified as outliers in the pulse shape parameter space is an effective strategy to suppress such unwanted contributions. These advanced data quality cuts can also be used to rebuild updated AP and ANPS templates or to refine the event selections employed for performing thermal gain stabilization and energy calibration. As illustrated by the chart in Fig. 4.2, the analysis can be re-iterated multiple times with progressively tighter cuts, updating templates and selections at each pass. The final efficiencies of these cuts on signal events might be energy-dependent and has to be computed in order to estimate the level of radioactive contaminants present in the crystals.

4.3.4 Heat-light coincidences: Light Yield and Discrimination Power

When a γ interaction occurs in the Li_2MoO_4 crystal, a coincident scintillation signal is expected on the facing LD, even if the signal amplitude is below the derivative trigger threshold used to flag signal events in the light channel. In particular, before applying the OF, the scintillation signal may be buried in the detector noise, whose RMS can exceed the signal amplitude.

To capture these sub-threshold signals, side trigger flags are introduced during preprocessing. These flags are inserted into the light channel waveforms every time a signal event is triggered in the corresponding heat channel. Using these flags, one can extract the corresponding light detector waveforms and apply the OF built using the AP and the ANPS of the light detector, as introduced in Eq. (4.2). The amplitude of the filtered waveform is then evaluated at the expected time of the light pulse maximum. This procedure, used to reconstruct the light signal associated with each triggered Li_2MoO_4 event, is composed of two main steps, which are illustrated by the chart in Fig. 4.2 and described below.

Average Side Pulse Once the Li_2MoO_4 energy spectrum has been calibrated, it is possible to select a region dominated by β/γ events (typically chosen between 1.5–3 MeV), which are expected to yield the highest scintillation light output. For each selected event in the Li_2MoO_4 crystal, the corresponding side-triggered waveform from the light detector is extracted. By averaging these side waveforms, the uncorrelated noise is suppressed, and a clear average side pulse emerges. This average waveform represents the typical scintillation signal associated with the selected class of events in the heat channel and it is used to perform the light synchronization described below.

Light synchronization Each side-triggered waveform is filtered using the OF constructed from the light detector’s AP⁸ and ANPS. The same transfer function is also applied to the average side pulse, resulting in a filtered pulse with a well-defined maximum at time t_{max} . The amplitude of each filtered side window can be computed simply by evaluating its waveform $V(t)$ at the time $t = t_{\text{max}}$. This procedure is called light synchronization. The amplitudes calculated in this way allow for the assignment of a light signal to each corresponding heat event.

Finally, by applying the calibration coefficients of the light detector as described earlier, each event can be associated with a energy-calibrated light signal. The light channel energy is then used to build the scatter plots of Figures 2.13a and 2.13b.

4.4 Pre-CCVR1

This section presents the characterization measurements of four $\text{Li}_2^{100}\text{MoO}_4$ crystals produced at the Nikolaev Institute of Inorganic Chemistry, Novosibirsk, Russia (NIIC).

4.4.1 Crystals and goals

The CUPID-Mo demonstrator successfully operated $\text{Li}_2^{100}\text{MoO}_4$ crystals from NIIC, demonstrating the feasibility of producing high-quality scintillating cryogenic calorimeters based on lithium molybdate [181].

Measuring crystals grown at NIIC, provides a valuable opportunity to: (i) begin defining a measurement protocol for assessing crystal quality using material from a known and trusted producer, (ii) characterize crystals that will serve as reference samples in future cryogenic runs aimed at testing the pre-production output from SICCAS, and (iii) recommission and evaluate the performance of the cryogenic system after major cryogenic interventions. In pre-CCVR1, NIIC $\text{Li}_2^{100}\text{MoO}_4$ crystals with the following features were characterized:

- Four $\text{Li}_2^{100}\text{MoO}_4$ scintillating crystals of cubic shape, isotopically enriched in the isotope ^{100}Mo , with dimensions $4.5 \times 4.5 \times 4.5 \text{ cm}^3$ ($\pm 0.2 \text{ mm}$ tolerance) and mass $\sim 285 \text{ g}$.
- The double Czochralski (CZ) crystallization technique was employed to reduce the level of impurities (see Sec. 3.3), in particular, to reduce the ^{40}K impurity concentration in the final product.

⁸This AP is the one derived from ionization events. Alternatively, the average side pulse can be used to build a new OF transfer function; however, given the working assumption that scintillation and ionization pulses are indistinguishable in the LD, the two choices are equivalent.

Crystal label	Size (mm ³)	Net Weight (kg)	Name
LMO1	45.13 × 45.14 × 45.15	0.2857	LMO-isot-6-39-top
LMO2	45.12 × 45.09 × 45.13	0.2858	LMO-isot-6-38-bottom
LMO3	45.13 × 45.09 × 45.14	0.2853	LMO-isot-6-38-top
LMO4	44.93 × 44.82 × 44.80	0.2819	LMO-isot-6-41-top

Table 4.1: Crystals characterized in the pre-CCVR1 run.

- Extremely high chemical purity was required to the producer, with the following maximum concentrations in the final crystals: Bi, Be, Co, Cu, Sn, Sr, Ta, Ni, and Pb < 0.5 ppm; B, Al, Nb, Mn, Sb, and Zr < 1 ppm; Ba, Fe, Na, Sc, and Zn < 5 ppm; W < 10 ppm.
- The crystals were grown starting from the following raw materials:
 - High purity molybdenum trioxide MoO₃ powder, isotopically enriched to no less than 95 % of ¹⁰⁰Mo.
 - High purity lithium carbonate Li₂CO₃ powder.

Both powders were required to meet the following radiopurity requirements:

²³⁸U < 2 × 10⁻¹⁰ g/g, ²³²Th < 2 × 10⁻¹⁰ g/g, ²¹⁰Pb < 10⁻⁴ Bq/kg, ⁴⁰K < 10⁻³ Bq/kg, ⁶⁰Co < 10⁻⁵ Bq/kg
These requirements, as well as the level of isotopic enrichment, were assessed by testing random samples from the production batches with dedicated measurements.

- The radiopurity requirements imposed to the producer for the final crystals are summarized in Table 4.2. No requirement on the ⁴⁰K content of the Li₂¹⁰⁰MoO₄ crystals was imposed in this case.

²³⁸ U	²³² Th	²¹⁰ Pb	⁶⁰ Co
< 3 × 10 ⁻¹³ g/g	< 3 × 10 ⁻¹³ g/g	< 1.0 × 10 ⁻⁵ Bq/kg	< 1.0 × 10 ⁻⁶ Bq/kg

Table 4.2: Radiopurity requirements for the final Li₂¹⁰⁰MoO₄ crystals.

- Surface treatment was carried out with highly radiopure SiO₂ powder, previously validated for the treatment of CUORE TeO₂ crystals. No specific requirements on optical finish or surface roughness were defined, and consequently the crystal surfaces were not optical. It was only requested that the treatment gave surfaces without haloes or any inhomogeneities. The last surface treatment before packaging and delivery leaves a thin layer of vacuum oil on the surface of the crystals.
- The producer was given strict requirements regarding the optical quality of the final crystals, that were optically clear, colorless, and bubble-free. An absorption length higher than 50 cm in the 550-700 nm spectral range has been required to the producer.

Table 4.1 reports the precise dimensions and masses of the crystals. Photographs of the four Li₂¹⁰⁰MoO₄ crystals are presented in Figs. 4.3c, 4.4 and 4.5.

4.4.2 Detector design and assembly




Figure 4.3: (a) Bottom TSP plate with mounted copper holders for the LD frames. The square foil segments, two per light detector, are the ^{55}Fe calibration sources. (b) Bottom light detectors: two LD frames, each hosting two LDs, are mounted on the copper frame holders and vertically spaced by PTFE spacers. (c) The four $\text{Li}_2^{100}\text{MoO}_4$ crystals with the features listed in Tab. 4.1 are seated on the PTFE elements.

The detector geometry, shown in Fig. 4.4a, comprises a single floor of four crystals arranged at the vertices of a square. Each crystal has two facing LDs, one above and one below. Both crystals and LDs are secured by PTFE holders. This configuration can be assembled quickly without the need to redesign major components or fabricate new copper parts, an ideal feature for a geometry dedicated to CCVR campaigns. Moreover, each crystal has line-of-sight to at least two adjacent crystals, enabling coincidence studies for the assessment of surface contamination. The following paragraphs summarize the procedures used for cleaning, assembling and instrumenting the detector.

Detector assembly and geometry The support structure is fabricated entirely from copper and assembled with brass screws. The tower support plate (TSP)⁹ is the copper structure that holds the detector: a top and a bottom plate linked by vertical copper bars. The entire assembly is thermalized to the mixing chamber and enclosed by the cryostat’s inner thermal shield. This structure is shown in Fig. 4.5c and encloses all the detectors, which are mounted as follows. Two copper holders, that will support the LDs, are screwed to the bottom TSP plate (Fig. 4.3a). Two copper frames, each accommodating two LDs, are positioned on the copper holders; PTFE elements keep the LDs anchored to the copper frames and serve as support for the other elements, spacing the frames from the structure (Fig. 4.3b). The four $\text{Li}_2^{100}\text{MoO}_4$ crystals rest on these PTFE supports (Fig. 4.3c). Two additional frames, hosting two LDs each and identical to the lower ones, lie on top of the crystals (Fig. 4.4b), followed by two copper top-frame holders (Fig. 4.4c). Vertical copper bars on opposite sides of the assembly secure copper pins that carry the wiring for the $\text{Li}_2^{100}\text{MoO}_4$ crystals thermistors and heaters and provide thermalization (see Figs. 4.5a and 4.5c).

⁹In CUORE, “TSP” originally referred to the top copper plate, thermalized to the MC, that supports the crystal towers. In the present work, the term “TSP” is used by extension to denote the full copper support structure.

Ge NTD gluing Ge NTD thermistors are glued to both the $\text{Li}_2^{100}\text{MoO}_4$ crystals and the LDs using Araldite Rapid glue. NTD sensors for the $\text{Li}_2^{100}\text{MoO}_4$ absorber measure $3 \times 3 \times 1 \text{ mm}^3$ and its electrodes are bonded with two gold wires (diameter $25\text{--}50 \mu\text{m}$). Also Si heaters are glued on the crystals in this phase. Prior to gluing, the layer of vacuum oil is removed from the crystal surfaces with acetone in the gluing area. A reproducible procedure deposits nine drops of glue under controlled pressure, monitored by a dynamometer, to optimize thermal contact while minimizing mechanical stress from differential thermal contraction of Ge and Li_2MoO_4 . This gluing strategy has been developed after many years of R&D on cryogenic calorimeters for CUORE and CUPID. NTD sensors for the LDs measure $3 \times 1 \times 1 \text{ mm}^3$ and are glued on the LDs using three spots of glue. In this configuration, since the Ge NTDs are glued directly to a Ge absorber, no mechanical stress from differential thermal contraction occurs at low temperatures, and alternative gluing methods (for instance, a thin veil) may be adopted in future runs. Photographs of one LD NTD and $\text{Li}_2^{100}\text{MoO}_4$ NTDs are shown in Figs. 4.5b and 4.5a respectively.




Figure 4.4: Geometry of pre-CCVR1 detectors. (a) Rendering of the detector layout, with all readout channels labeled. (b) Copper frames supporting the two top LDs are placed atop the crystals and rest on PTFE spacers. (c) The top holder is installed on the PTFE elements and thermally coupled to the TSP via copper bands screwed to the lateral vertical copper bars fastened to the bottom holder.

Thermalization and wiring The top and bottom TSP plates, connected by vertical copper bars, are directly thermalized to the cryostat’s inner thermal shield and to the MC. The copper frame supports are well thermalized to the heat sink via brass screws on the bottom plate and copper thermalization bands connecting them to the lateral vertical bars. In contrast, the LD frames are not coupled by copper thermalizations and rest on PTFE spacers, rendering them thermally decoupled (floating). Similarly, the $\text{Li}_2^{100}\text{MoO}_4$ crystals rest on PTFE supports and are not in direct contact with copper. As indicated by the CUORE electrothermal model, absorber thermalization proceeds mainly via the NTDs and their readout electronics. In this setup, the $\text{Li}_2^{100}\text{MoO}_4$ crystals feature gold wires terminating on copper pins attached to the vertical bars, providing effective thermal contact with the heat sink (Fig. 4.5a). For the light detectors, however, the NTD gold wires connect to pins mounted on the floating LD frames (Fig. 4.5b), so thermalization occurs only through the PTFE elements and the subsequent wiring. Beyond the copper pins, the wiring for




Figure 4.5: (a) Close-up of the Ge NTD sensor, Si heaters, and readout wiring on the $\text{Li}_2^{100}\text{MoO}_4$ crystals. Gold wires are bonded to copper pins, which are thermally anchored to the mixing chamber via a vertical copper bar affixed to the bottom holder. (b) Close-up of the Ge NTD sensor and gold wires on a light detector; here, the copper pins are glued on the LD frame, which rests solely on PTFE supports without any direct copper thermalization. (c) Fully assembled detector tower with one half of the inner cryostat's copper thermal shield installed.

both LDs and $\text{Li}_2^{100}\text{MoO}_4$ crystals continues with twisted pairs of constantan¹⁰, insulated with a silk layer and secured along its path to the MC by copper tape (4.5c). The twisted configuration minimizes electromagnetic pickup and renders other disturbances common-mode, which are then suppressed by the voltage sensitive differential amplifier [131].

Material cleaning All surfaces facing the detectors are potential sources of background and must be rendered radiopure prior to assembly. The inner thermal shield of the cryostat, thermalized to the MC, faces the detectors directly. The same holds for the structure holding the absorbers

¹⁰Constantan is a copper–nickel alloy with stable resistance across a wide temperature range. The constantan used here has resistivities of 0.29 Ω/cm and 1.55 Ω/cm .

and the readout wiring. To remove surface radioisotopes, each component is first wiped with high-purity alcohol, then subjected to ultrasonic cleaning at 60 °C in ultra-pure water¹¹ and acid-soap bath, followed by thorough rinsing. Ultrasonic cleaning is repeated in a citric-acid bath containing sufficient H₂O₂ to produce visible bubbling (indicating mild etching), after which parts are rinsed immediately. Cleaning is completed by drying each piece with wipes, flushing with dry N₂, baking in a vacuum oven, and finally vacuum-sealing until installation. All chemicals used in this phase are of the highest available purity. The mild etching removes radioisotopes embedded in the outermost micrometers of the material and passivates the surfaces, thereby reducing oxidation.

4.4.3 Calibration sources

During dedicated calibration runs, commercially available thoriated rods are deployed outside the cryostat (within the external lead shielding) to calibrate the heat channels. These rods provide external γ lines from the ²³²Th decay chain. Data-quality cuts reject windows containing more than one pulse, so an excessively high count rate reduces selection efficiency, whereas a rate that is too low prolongs calibration runs at the expense of background measurements. To achieve an optimal rate of a few tens of mHz on the Li₂¹⁰⁰MoO₄ crystals, the rods are cut into ~ 3 cm segments.

The LDs are optimized for low-threshold operation and require a calibration curve that can be reliably extrapolated below 1 keV in order to measure the energy of scintillation light. Accordingly, ⁵⁵Fe sources are mounted on the top and bottom copper holders, directly facing each light detector. These sources emit X-ray lines from ⁵⁵Fe decay, namely Mn- K_α at 5.89 keV and Mn- K_β at 6.49 keV. Each source is prepared by depositing 10 μ L of an aqueous ⁵⁵Fe solution onto a thin foil. A single foil yields approximately 70 counts/h; multiple foils may be employed per LD to increase the rate. With a half-life of 2.737 years, these sources remain suitable for subsequent CCVR campaigns. In the present run, two ⁵⁵Fe foils were installed to face each LD. The foils are visible in Fig. 4.3a.

4.4.4 Data taking

Data-taking began at the end of September 2023 and continued for approximately three months. During the first few weeks, the readout electronics and the temperatures reached by the MC, the TSP and the detectors were characterised. These measurements were used to study the cryostat's behaviour to complete its recommissioning, and to assess the effectiveness of the thermalizations in the pre-CCVR1 assembly. Next, one week was devoted to optimizing the detectors' working point and to characterizing the response of the light detectors. Finally, 38 days were allocated to the characterization of the Li₂¹⁰⁰MoO₄ crystals, comprising 11 days of calibration (with external ²³²Th sources deployed) followed by 27 days of background measurement.

During the first few weeks, problems arose in the preamplifiers of several front-end boards, preventing data collection from four of the eight light detectors in this setup. In addition, the constantan wiring of LMO4 and LD6 disconnected during cooldown, likely due to differential thermal contraction at the soldered MC-pin connections; only the channels listed in Table 4.3 remained operational. Consequently, it was not possible to characterise LMO-isot-6-41-top (LMO4) and measure the LY of LMO-isot-6-38-bottom (LMO2).

Additionally, during the initial phase of data taking, several issues related to LD3 were observed. In particular, this channel appeared to be amplified with a gain different from the nominal value set on the preamplifier. This was revealed by an analysis of the white noise component in the ANPS,

¹¹Water with resistivity of $\sim 18.2 \text{ M}\Omega \cdot \text{cm}$ at 25 °C.

which exhibited a value nearly half of the Noise Floor predicted by Eq. 2.14, an unphysical result. As a consequence, the characterization of the WP for this detector, including the measurements of R_{base} , R_{WP} , and the intrinsic gain in $\mu\text{V}/\text{keV}$, cannot be considered reliable. Nonetheless, with an appropriate calibration, this LD can still be used for LY measurements, even in the absence of a complete WP characterization.

Top LDs	LD5	×	×	×
Heat channel	LMO1	LMO2	LMO3	×
Bottom LDs	×	×	LD3	LD4

Table 4.3: Mapping of the pre-CCVR1 operational channels; the detector geometry is depicted in Fig. 4.4a.

4.4.5 Thermalizations

The initial part of the run was dedicated to characterizing the cryogenic system and evaluating the thermalization of the detectors. The cryostat was instrumented with two calibrated NTD thermistors—one glued to the MC and the other to the TSP. Although their calibration may be subject to some degree of uncertainty, the thermometers consistently indicated a temperature of approximately 15 mK at the MC and a slightly higher value of around 17 mK at the TSP. Measurements were repeated periodically over several days during data taking, and the temperature remained stable throughout the entire run. These results suggest that the cryogenic system is performing as expected and complete the recommissioning of the Hall C cryostat.

During this stage, the base resistances of the detectors were measured. The values of R_{base} are reported in Tab. 4.4 for the LDs and in Tab. 4.5 for the Li_2MoO_4 crystals. For all channels, the measured resistances fell within the range of approximately 1–10 M Ω . In other measurements, detectors employing the same type of NTDs showed typical R_{base} values on the order of tens up to a few hundreds of M Ω for heat channels and hundreds of M Ω for light channels. Even in the absence of a calibration curve for the NTDs used in pre-CCVR1, a comparison with these measurements indicates that the observed values are incompatible with the temperatures measured on the TSP. In particular, while the heat channels exhibit only a moderate deviation from expected R_{base} values, the light detectors appear to be significantly warmer than expected. The average baseline voltage V_{NTD} remained approximately constant for all detectors throughout the data taking, with no evident long-term drifts.

To investigate the observed discrepancy between the detectors’ R_{base} and the TSP temperature, a dedicated test was performed. In this test, all NTDs were set to the lowest bias (corresponding to $P \sim 10^{-3}$ pW), except one, which was biased with high current (resulting in $P \sim 20$ pW) for a few minutes before being set back to low bias. This procedure locally increases the temperature of the selected NTD. The baseline voltages, serving as proxies for detector temperatures, were then monitored across all channels for several hours. The test was repeated using LD3 and LD5, which are mounted on different frames. The results are shown in Fig. 4.6, where the detector subjected to heating is highlighted in red. The baseline of the heated detector required more than two hours to recover its original value. Moreover, its neighboring LD on the same frame, as well as the Li_2MoO_4 crystal mounted above or below, also exhibited a temperature increase, relaxing back to the base temperature with the same time constant.

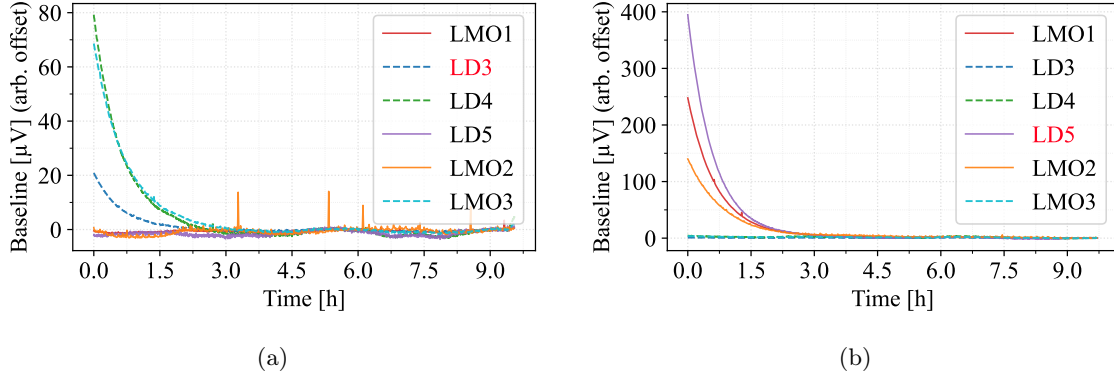


Figure 4.6: Average baseline values (V_{NTD}) of all the available channels as a function of time following the heating of (a) LD3, which shares a frame with LD4 and supports LMO3; and (b) LD5, whose frame sits on LMO1 and LMO2. The observed spikes are caused by random coincidences between baseline sampling and pulser events, and can be disregarded.

These observations indicate that the system is not well thermalized to the MC. In particular, they confirm that the detector assembly, described in Sec. 4.4.2, features LD frames that are partially decoupled from the heat sink.

The lower-than-expected R_{base} values reduce flexibility in selecting the optimal working point, thereby limiting the margin for the choice of R_{WP} and the optimization of the SNR. Moreover, increasing the bias current on one channel affects the R_{base} of other channels sharing the same frame. As a result, the working point of each channel is not fully independent, which complicates the optimization process across multiple detectors.

4.4.6 Working point and noise in the light detectors

The noise affecting the LDs is shaped by the transfer function of the detector, which depends on the choice of the WP. Although a formal description of how the transfer function evolves with the biasing current would require the development of a dedicated electro-thermal model, several general considerations hold even in its absence. In particular, the bandwidth of the detector is typically determined by the time constant $R_{\text{NTD}}C_{\text{par}}$, where C_{par} is the parasitic capacitance of the readout cables, as shown in Fig. 2.7. High values of R_{NTD} integrate low-frequency noise, making the LDs more sensitive to mechanical vibrations. Conversely, lower values of R_{NTD} push the detector response to higher frequencies, reducing sensitivity to vibrational noise.

For a given NTD type¹², lower R_{NTD} values can be achieved either by maintaining the LD at a higher temperature, thereby reducing R_{base} , or by increasing I_{bias} to select a WP corresponding to a lower R_{WP} . However, both approaches result in increased absorber temperature, reducing the intrinsic gain and consequently the signal amplitude. The best choice typically involves a compromise between the two.

The low R_{base} values observed in pre-CCVR1 (discussed in Sec. 4.4.5) indicate that the detectors

¹²NTDs with lower R_0 , which could give a lower R_{NTD} at equal absorber temperature, are not available during the pre-CCVR phase.




Figure 4.7: I - V characteristics (blue dots) and R - P curves (red squares) for (a) LD4 and (b) LD5. The noise level and spectral shape were studied at various points along the load curve; the colored circles indicate the working points at which the corresponding ANPS, shown in Fig. 4.8, was measured.




Figure 4.8: ANPS measured on (a) LD4 and (b) LD5 by scanning the biasing currents I_{bias} indicated in the legend. For some WPs, multiple measurements (Scan 2/7 and Scan 5/8) were acquired at different times during the run to monitor the stability of the noise over time. The corresponding values of R_{WP} for each configuration are shown in Fig. 4.7. The dashed lines represent the theoretical noise floor, computed using Eq. 2.14 for each configuration, while the grey shaded area indicates the contribution from the preamplifier alone.

were already quite warm. While this condition helps mitigate the impact of low-frequency noise, it also reduces the available margin of WP optimization. In particular, selecting too high a value of I_{bias} may further warm up the detector, reducing the signal amplitude without a significant improvement of the noise, which might be limited by the noise floor (Eq. 2.14) in these conditions.

For these reasons, a scan of the ANPS at different WPs was carried out on all LDs. The results

are shown in Figs. 4.8 and 4.7 (LD3 is omitted because its amplifier had unreliable gain). In all configurations, the measured noise is close to the theoretical noise floor. In strong overbias conditions (Scans 4, 5, and 8 in the plots), the noise floor is dominated by the preamplifier contribution, resulting in similar noise levels across these configurations, albeit at the cost of reduced signal amplitude and thus worse SNR. On the other hand, lower bias configurations exhibit a higher noise floor and increased sensitivity to low-frequency noise, especially for LD4. The ANPS of LD5 exhibits peaks between 30-100 Hz indicating a higher electromagnetic pick-up noise.

Scan 6 represents a compromise between these two regimes and additionally features an R_{WP} close to the inversion region, as it can be deduced from the load curve shown in Fig. 4.7. Moreover, in the pre-CCVR1 assembly, the detectors are not thermally independent (see Fig. 4.6), and maintaining too high a value of I_{bias} on one LD may further reduce the already low R_{base} of other detectors sharing the same frame. For these reasons, this working point—named WP4—was selected for the characterization of the LY of the $\text{Li}_2^{100}\text{MoO}_4$ crystals. The value of R_{WP} is shown in Tab. 4.4. All LDs were acquired with a sampling frequency of 5 kHz and triggered using 3 s windows with a 1 s pretrigger. An anti-aliasing Bessel filter with a cutoff frequency of 80 Hz was applied to LD3 and LD4, while a 120 Hz cutoff was used for LD5.

4.4.7 Light detectors calibration and performances

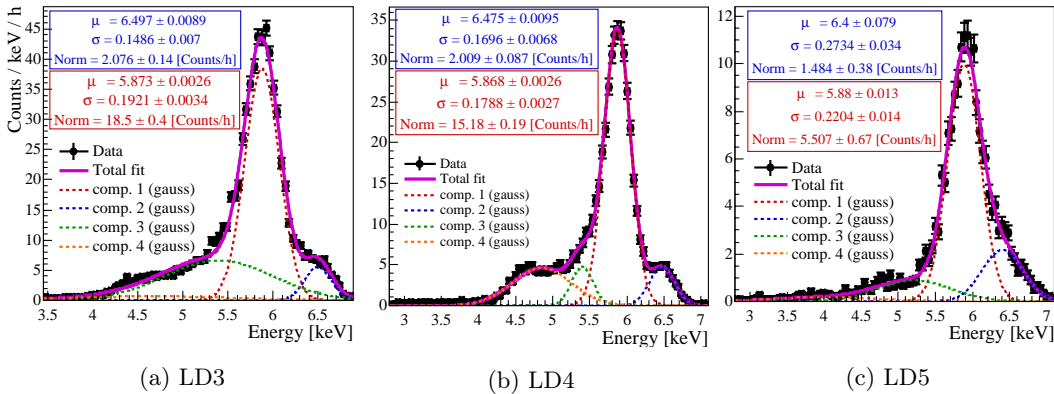


Figure 4.9: Calibrated ^{55}Fe spectra measured on the available light detectors (rates not corrected for cut efficiencies). The data were fitted with a multi-Gaussian model to extract the energy resolution at the $\text{Mn-}K_{\alpha}$ line (red). The $\text{Mn-}K_{\beta}$ line (blue) was also identified.

The analysis strategy described in Sec. 4.3 was used to obtain calibrated ionization spectra for the LDs. Thermal gain stabilization was not required for LDs. Calibration was performed using only the $\text{Mn-}K_{\alpha}$ line at 5.89 keV on the complete set of background data. The calibrated spectra, uncorrected for cut efficiencies, are shown in Fig. 4.9. The performance of the LDs was assessed through the energy resolution at the $\text{Mn-}K_{\alpha}$ line and at the baseline; estimates of the LDs’ intrinsic gain (sensitivity) and pulse rise and decay times were also obtained. These results are summarized in Tab. 4.4 and were determined by fitting the ^{55}Fe peak. Unrecognized backgrounds affecting its low-energy tail were modeled with multiple Gaussian components. The results of the fit are

Channel	Bsl. res. σ [eV]	^{55}Fe Mn- K_α σ [eV]	Sensitivity [$\mu\text{V}/\text{keV}$]	R_{base} [M Ω]	R_{WP} [M Ω]	RT [ms]	DT [ms]
LD3	157.7 ± 0.4	192 ± 3	×	×	×	~ 4.8	~ 10
LD4	148.0 ± 0.5	179 ± 3	~ 2.76	~ 3	~ 2.62	~ 4.8	~ 14
LD5	198 ± 1	220 ± 14	~ 3.49	~ 4.22	~ 4.10	~ 4	~ 11

Table 4.4: Performance of the light detectors at the selected WP. Due to the unreliable amplifier gain of LD3, a precise evaluation of its resistance and intrinsic gain was not possible. The RT in LD3 and LD4 are limited by the Bessel cutoff frequency.

presented in Fig. 4.9. The Mn- K_β line at 6.49 keV was reconstructed at its nominal energy in all LDs, confirming the correctness of the calibration, although exhibiting a slightly worse resolution than the K_α line. The ratio between the observed intensities of the two lines roughly matches the expected value ($R_{K_\alpha/K_\beta} \sim 7.2$), with the small discrepancies to be attributed to the naïve modeling of backgrounds. The measured rate, uncorrected for cut efficiencies, was of the order of a few tens of counts per hour, in rough agreement with expectations for the ^{55}Fe sources described in Sec. 4.4.3.

The detectors showed a baseline resolution $\sigma \sim 150$ eV. The slightly worse resolution observed on LD5 is attributed to increased pick-up noise in its ANPS, which includes higher-frequency components that may overlap with the signal bandwidth, which extends to higher frequencies for the LDs. This overlap degrades the SNR even after wOF is applied, as described by Eq. 4.1.

The low values of R_{base} , and the correspondingly low R_{WP} achieved in this run, push the signal bandwidth to higher frequencies. For LD3 and LD4, the measured RT is compatible with ~ 4.2 ms, which corresponds to the minimum measurable RT given the 80 Hz Bessel filter cutoff¹³. Their signal bandwidth is therefore limited by the filter, with pulse shapes determined by its low cutoff frequency, leaving limited room for improvement by the wOF. LD5, on the other hand, is operated with a higher Bessel cutoff of 120 Hz (minimum measurable RT ~ 2.8 ms). The RT measured on LD5 exceeds this limit and thus reflects the intrinsic response of the system.

4.4.8 Working point for the heat channels

Noise studies and WP optimization were also carried out on the heat channels. Unlike for the LDs, in this case the Si heater glued onto the crystal surface allows to inject artificial monochromatic pulses which can be used to study the evolution of signal amplitude as a function of the WP. The results of this study, along with the corresponding I - V and R - P curves, are presented in Fig. 4.10.

As a general rule of thumb, it is often possible to exploit the favourable Debye temperature of Li_2MoO_4 crystals, which leads to higher intrinsic gain compared to TeO_2 absorbers (see table 2.1), and operate the detectors in slight overbias—i.e., at lower R_{WP} and lower pulser amplitude—where they become less sensitive to low-frequency noise such as mechanical vibrations. However, during pre-CCVR1, the low values of R_{base} observed in the heat channels indicate that the absorbers are already quite warm, which limits the margin for WP optimization. In particular, higher values of I_{bias} risk further increasing the absorber temperature, thereby reducing too much the signal amplitude without a corresponding improvement in the noise levels. To better understand the impact of the noise, dedicated studies were performed. These showed that all detectors exhibited

¹³For a Bessel filter with cutoff frequency f_c , the minimum measurable 90%–10% RT is $\sim 0.339/f_c$ [192].




Figure 4.10: Solid lines: I - V characteristics (blue dots) and R - P curves (red dots). Dashed lines: unfiltered amplitudes of artificial monochromatic pulses, generated using a Si heater, as a function of V_{NTD} (blue crosses) and P (red crosses). The selected WP is marked by a green circle on all four curves. The value of R_{WP} measured close to the end of the data-taking (day 27) lies close to the curves acquired at the beginning of the run, indicating stable operating conditions. The closest measurement to the selected WP taken at the beginning of data taking (day 0) is highlighted by a purple square.

a $1/f$ noise excess below 2 Hz in the ANPS, exceeding what is expected from the preamplifier alone. Furthermore, the white component of the ANPS was often limited by the noise floor defined in Eq. 2.14, with the preamplifier contribution typically dominating. This behavior arises from the low R_{WP} values achievable when starting from such low R_{base} : the Johnson-Nyquist noise is strongly suppressed at low R_{NTD} . As a result, under these warm-detector conditions, increasing the biasing current has negligible impact on the overall noise level. This implies that the optimal SNR can, to first approximation, be found by simply maximizing the signal amplitude. It should be noted, however, that this is not a general result: if the detectors were colder and exhibited higher R_{NTD} values, the low-frequency noise would be amplified and the signal bandwidth would be shifted, making the noise behavior (and the SNR) more sensitive to WP choice.

Moreover, in the pre-CCVR1 assembly, the detectors were not thermally independent (see Fig. 4.6), and maintaining too high a value of I_{bias} on one detector may further reduce the already low R_{base} of other detectors sharing the same frame.

For these reasons, the working point selected for the characterization of the $\text{Li}_2^{100}\text{MoO}_4$ crystals, named WP4, was chosen near the inversion region, close to the maximum pulser signal amplitude. The value of R_{WP} is reported in Tab. 4.5, and the corresponding values of I_{bias} , pulser amplitude, and V_{NTD} are highlighted in Fig. 4.10. All LMOs were acquired using a Bessel filter with a 47 Hz cutoff frequency and a 2 kHz sampling rate, triggered using 10 s windows with a 2.5 s pretrigger.

4.4.9 Characterization of $\text{Li}_2^{100}\text{MoO}_4$ NIIC crystals

The first four days of data taking were dedicated to the calibration of the $\text{Li}_2^{100}\text{MoO}_4$ crystals. The analysis strategy described in Sec. 4.3 was used to reconstruct the energy of each event. The wOF was built using AP and ANPS templates extracted from the first ten days of background data, which benefit from higher statistics and are less affected by pile-up. The calibration data

were then filtered with the wOF built with this templates and stabilized against thermal gain drifts using monochromatic pulses injected every 8 minutes via the Si heaters. Basic data-quality cuts were applied to reject pile-up events and pulses growing on unstable baselines.

The following γ peaks were identified in the filtered, stabilized calibration spectrum:

511 keV, 2614.5 keV (^{208}Tl), 911.2 keV (^{228}Ac), 969 keV (^{228}Ac)

and used to perform a linear energy calibration. The resulting calibration coefficients were then applied to the stabilized background data. The calibrated background spectrum for all $\text{Li}_2^{100}\text{MoO}_4$ detectors is shown in Fig. 4.11. It exhibits γ lines from the ^{238}U chain (in particular from isotope ^{214}Bi) and evident ^{210}Po α and Q-value ($\alpha+\text{NR}$) peaks. No additional α or Q-value lines are observed. These observations are discussed below in the paragraph on radiopurity.

Quenchig factor

As discussed in Sec. 4.3.3, β/γ events are quenched relative to α events, and a $\text{QF} > 1$ must be applied to restore the correct energy scale for α interactions after performing a γ -based calibration. Since no clear α peaks are visible in the calibration data, the QF is estimated by fitting the ^{210}Po α line observed in the γ -calibrated background spectrum with a Gaussian model. This estimate relies on two key assumptions: (a) that the γ calibration can be reliably extrapolated up to ~ 5.3 MeV, and (b) that the few counts observed at the α line (presumably due to surface contamination on the crystal or nearby materials) can be adequately described by a Gaussian peak. Both assumptions are very strong and are valid only within a margin of approximation: the calibration extrapolation introduces non-negligible uncertainty at such high energies, and the surface nature of α events make the Gaussian approximation questionable. However, no alternative method to estimate the QF is available. The values obtained for the QF are reported in Tab. 4.5.

The energy spectrum in the α region, corrected for the estimated QF, is shown in Fig. 4.13. This correction also relies on the additional assumption that the QF is constant with energy¹⁴.

Data quality cuts and selections efficiency

The derivative trigger efficiency is assumed to be ~ 1 for all channels above ~ 400 keV. The efficiency of the data quality selections is estimated as the ratio of the number of pulser events before and after cuts, under the assumption that the cut efficiency is flat in energy. Since the only cuts applied are designed to reject pile-up events or pulses growing on unstable baselines, this assumption holds to a reasonable approximation. The resulting cut efficiencies are 72.5(7)% for LMO1, 72.0(7)% for LMO2, and 79.1(8)% for LMO3; uncertainties are statistical only (binomial).

Pulser events are not flagged by the readout system. However, due to their significantly shorter RT compared to particle-induced pulses, a cut on RT is effective in removing pulser events while retaining a signal efficiency close to 1. The equivalent energies of pulser events are above 6 MeV in all channels.

The spectra shown in Figs. 4.11 and 4.13 are corrected for cut efficiency computed as described above. These selections efficiencies are also used in the radiopurity analysis.




Figure 4.11: Background spectrum corresponding to a total live time of ~ 24 days. The spectrum is corrected by cuts efficiencies. The α events are not corrected by QF. For visual clarity, the *total* spectrum (sum of the three crystal spectra) is not normalized to the total mass ($3\times$ the mass of a single crystal); consequently its rate (in counts/keV/kg/yr) is three times higher than the actual value. This convention is purely for visualization and is not used for any quantitative estimates.

Ch.	R_{base} [M Ω]	R_{WP} [M Ω]	Bsl. res. FWHM [keV]	^{208}Tl FWHM [keV]	Sensitivity [nV/keV]	RT [ms]	DT [ms]	QF α ^{210}Po
LMO1	~ 1.1	~ 1.08	7.21 ± 0.47	6.51 ± 1.29	~ 4.58	~ 24	~ 270	1.0807(8)
LMO2	~ 7	~ 5.6	6.44 ± 0.09	8.20 ± 1.57	~ 12.76	~ 34	~ 430	1.0773(9)
LMO3	~ 1	~ 0.94	19.82 ± 0.19	22.90 ± 4.00	~ 1.93	~ 24	~ 430	1.0813(3)

Table 4.5: Performance of the NIIC $\text{Li}_2^{100}\text{MoO}_4$ crystals tested in pre-CCVR1. The QF values for α particles were estimated by fitting the ^{210}Po α line in the background spectrum with a Gaussian model. Only statistical uncertainties are reported.

Calorimetric performance

The calibration data were used to assess the calorimetric performance of the $\text{Li}_2^{100}\text{MoO}_4$ crystals. The energy resolution and intrinsic gain (sensitivity) were estimated by fitting the ^{208}Tl line at 2614.5 keV. Estimates of the baseline resolution, as well as the pulse rise and decay times, were also obtained. The results are summarized in Tab. 4.5.

Both LMO1 and LMO2 achieved good FWHM energy resolutions at the ^{208}Tl line, compatible within 2σ with the CUPID target of 5 keV. However, due to the short time dedicated to calibration (~ 4 days), the available statistics at the ^{208}Tl line is limited, and the resulting resolutions at 2614.5 keV are affected by large statistical uncertainties. In particular, LMO1 shows a baseline resolution worse than its resolution at 2614.5 keV –but still statistically compatible– reflecting this limited precision. Baseline resolutions, which are known with much smaller statistical uncertainties, indicate that under the operating conditions of this run, the detectors do not meet the 5 keV CUPID

¹⁴An energy dependence of the QF has been observed in other measurements performed on Li_2MoO_4 crystals during the CUPID R&D phase, although no publication is currently available.

requirement.

Although LMO1 and LMO3 have similar R_{base} values and comparable noise levels (both limited by the noise floor which is dominated by the preamplifier term) the FWHM resolution of LMO3 is significantly worse, around ~ 20 keV. This is attributed to a sub-optimal WP: the lower I_{bias} resulted in a reduced signal amplitude, as shown by the dashed lines in Fig. 4.10. This also explains why LMO1 exhibits a lower intrinsic gain than LMO3, despite having similar R_{base} and R_{WP} .

In contrast, the higher intrinsic gain observed in LMO2 is attributed to both a higher R_{base} , suggesting a colder absorber with more favourable heat capacity, and to a WP chosen close to the maximum signal amplitude, again visible from the dashed lines in Fig. 4.10.

As also observed for the LDs, the faster RTs measured on LMO1 and LMO2 with respect to LMO3 can be interpreted as a consequence of the increased detector bandwidth resulting from their lower R_{WP} .

These considerations show that the performance differences observed among the three $\text{Li}_2^{100}\text{MoO}_4$ crystals can be fully explained by the combined effects of WP selection and R_{base} . The collected data do not provide any experimental evidence of intrinsic differences in crystal quality. Deviations from the CUPID 5 keV FWHM energy resolution target are attributed to sub-optimal operating conditions achieved in this run, as indicated by the low R_{base} values, which reflect a relatively warm absorber temperature.

Light Yield and Discrimination Power

The LY was estimated using analysis strategy described in Sec. 4.3.4. Although calibration data offer higher statistics at the ^{208}Tl line, they contain very few α events, limiting their usefulness for evaluating the discrimination power. For this reason, the results presented in the following are based on the background dataset.

The LY for each event was computed as the light-to-heat ratio, following the assumptions introduced in Sec. 2.4.1. The resulting light-vs-heat and LY-vs-heat distributions for LMO1 and LMO3 are shown in Fig. 4.12a. While the width of the LY distribution depends on energy (see Sec. 5.4.3), in order to increase the statistics available for estimating the LY and the DP, a wide energy window between 2 MeV and 6 MeV was selected. This range includes both higher-statistics β/γ events as well as degraded α events from surface interactions and the Q-value and α lines of ^{210}Po . A fit to the LY projection within this range was used to simultaneously extract both the average LY for β/γ interactions and the corresponding $\text{DP}(E > 2 \text{ MeV})$, calculated using Eq. 2.18.

The results of this analysis are presented in Fig. 4.12b and summarized in Tab. 4.6. The quoted uncertainties are statistical only and do not include possible systematic effects arising from

Crystal (LD)	LY (β/γ) [keV/MeV]	LY width (β/γ) [keV/MeV]	DP ($E > 2 \text{ MeV}$)
LMO1 (LD5 - Top)	0.285(8)	0.100(8)	2.1(8)
LMO3 (LD3 - Bottom)	0.307(5)	0.073(4)	3.0(2)

Table 4.6: LY distribution mean and width, and DP measured on the facing LD for the NIIC $\text{Li}_2^{100}\text{MoO}_4$ crystals as measured in pre-CCVR1. LD3 (LD5) faces the bottom (top) face. Only statistical uncertainties from the fit are reported.

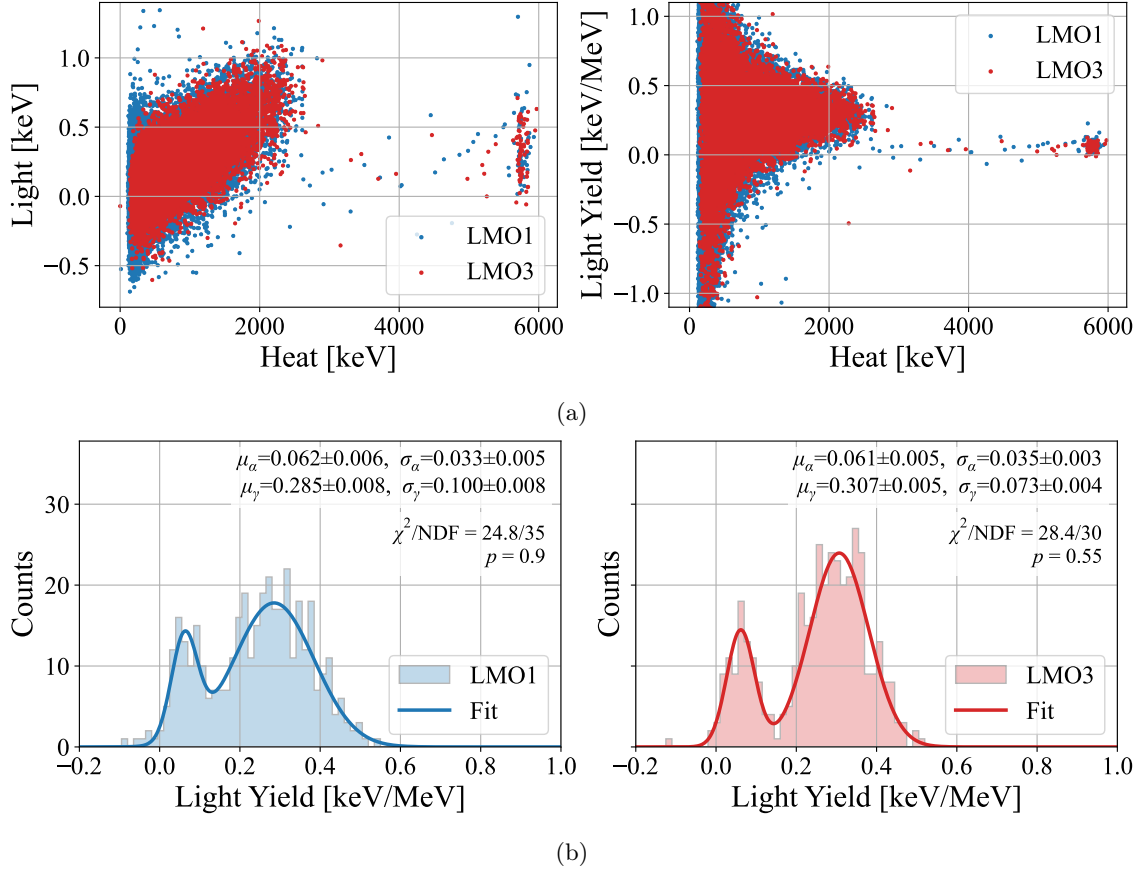


Figure 4.12: (a) Light-vs-heat (left) and LY-vs-heat (right) distributions for LMO1 and LMO3. (b) Fit of the LY distributions obtained events having $E > 2$ MeV in LMO1 (left) and LMO3 (right).

the energy reconstruction strategy used for light signals¹⁵. The LY values obtained for the two NIIC $\text{Li}_2^{100}\text{MoO}_4$ crystals (uncorrected for geometric light collection efficiency) are consistent with the ~ 0.3 keV/MeV reported in the literature for Li_2MoO_4 crystals operated with similar detector geometries and light detector configurations, as shown in Fig. 11 of Ref. [155].

The measured DP does not reach the CUPID goal of $\text{DP}(\sim 3 \text{ MeV}) > 5$, but in this run, the result is limited by the performance of the LDs discussed above and by the availability of a single light detector. In any case, CUPID will employ NTL assisted LDs, and a direct comparison with CUPID expectations is therefore not meaningful. Moreover, the DP depends on the crystals only through their LY, which, as noted, is in agreement with values reported in the literature.

¹⁵As is the case for the thermal quenching factor in heat channels, different types of incident radiation, such as X-rays or scintillation photons, may not follow the same calibration curve.

Radiopurity

The background data were used to assess the radiopurity of the crystals. The total background spectrum, corrected for cut efficiencies, is shown in Fig. 4.11, and a zoom on the QF-corrected α region is presented in Fig. 4.13. The corresponding livetime is approximately 24 days (576.943 hours) for LMO1 and LMO3, and 23 days 22 hours (574.166 hours) for LMO2. In the γ region, clear peaks from ^{214}Bi (^{238}U chain) are observed with relatively high intensity. If originating from surface or bulk crystal contamination, α lines from the fast chain segment (^{226}Ra to ^{214}Po) should also appear between ~ 4.5 MeV and ~ 7.5 MeV with comparable intensity. Since these are not observed, the ^{214}Bi lines are attributed to external sources, likely due to the Rn box of the cryostat not being flushed with N_2 during most of the background run.

It should also be noted that the primary goal of the pre-CCVR1 run was not to perform a radiopurity evaluation of NIIC-grown crystals. Since CUPID will adopt crystals produced by SIC-CAS, these crystals were used as a proof of concept and to begin defining a radiopurity validation protocol. The NIIC crystals, grown by a trusted supplier, also serve as reference standards for the upcoming pre-CCVR and CCVR campaigns on SICCAS samples.

For these reasons, the β/γ region is excluded from the radiopurity analysis, and no study of ^{40}K or ^{60}Co contamination was performed. Instead, the α region was used to derive limits on bulk contamination from ^{232}Th , ^{238}U and ^{235}U progenitors, and to assess the presence of surface and bulk contamination from ^{210}Po .




Figure 4.13: Zoom on the α region of the background energy spectrum, corrected for QF. The vertical shaded bands mark the Q -values of the main progenitors of interest: ^{210}Po , ^{232}Th , ^{238}U , ^{235}U . Shaded areas represent the $\pm 3.5\sigma_\alpha$ integration windows used for estimating upper limits on bulk contamination.

^{210}Po contamination The spectrum in Fig. 4.13 clearly shows both the full-energy α peak and the Q -value peak from ^{210}Po . This isotope is a common contaminant, as ^{222}Rn gas tends to accumulate on surfaces and may also become embedded in the crystal bulk during growth. With a short half-life of approximately 3.8 days, ^{222}Rn quickly decays into long-lived progeny, most notably ^{210}Pb (half-life ~ 22.2 years), which produces ^{210}Po through its decay chain.

To derive an upper limit on bulk ^{210}Po , a conservative hypothesis is adopted in which *all* events observed under the Q -value peak are attributed to internal decays. Under this assumption,

the spectrum is integrated over $Q \pm 3.5\sigma_\alpha$, where σ_α is determined by fitting the ^{210}Po α peak with a Gaussian model in the worst-performing heat channel to ensure maximal conservativeness. The corresponding integration window is shown in Fig. 4.13 as a shaded area, and the resulting measurements are reported in Tab. 4.7. These values already approach the radiopurity requirements even under the deliberately conservative assumption that the entire Q -value line arises from bulk contamination.

The presence of a distinct α line indicates that at least part of the activity likely originates from surface-related events, either at the crystal surface or on nearby passive elements. Quantifying surface activity requires additional assumptions and dedicated simulations or coincidence studies. Given the goals of pre-CCVR1, such studies were not pursued here; only bulk limits are thus quoted.

^{232}Th , ^{238}U , and ^{235}U progenitors Conservative upper limits on bulk contamination from ^{238}U , ^{235}U , and ^{232}Th progenitors are obtained by integrating in $Q \pm 3.5\sigma_\alpha$ around the respective Q -values, using the same σ_α definition as above. These very conservative integration ranges are shown by the shaded bands in Fig. 4.13.

When no counts are observed in the interval, a zero-background hypothesis is adopted, and the Feldman–Cousins method [193] is used to compute 90% C.L. upper limits (applied whenever fewer than 10 counts fall in the window). The results are summarized in Tab. 4.7. The specific activities of ^{232}Th and ^{238}U half-lives can be used to convert the requirements imposed to the producer (Table 4.2) into equivalent activity limits. These correspond to $1.2 \mu\text{Bq/kg}$ for ^{232}Th and $3.7 \mu\text{Bq/kg}$ for ^{238}U . The conservative limits in Table 4.7 indicate that, when the three crystals are considered individually, the ~ 28 -day exposure yields upper limits above the targets, implying that compliance cannot be tested with the current statistical sensitivity.

In order to assess whether the accumulated exposure suffices to test the requirements in Table 4.2, the spectra of the three operational crystals are summed, assuming uniform bulk contamination across all crystals. The total spectrum is then analyzed with the Feldman–Cousins method is performed adopting the same conservative assumptions as above (zero background in the $Q \pm 3.5\sigma_\alpha$ integration window). For ^{238}U , the combined 90% C.L. upper limit is $< 2.94 \mu\text{Bq/kg}$, satisfying the $< 3.7 \mu\text{Bq/kg}$ requirement derived from Table 4.2. For ^{232}Th , the 90% C.L. upper limit is $< 6.04 \mu\text{Bq/kg}$, meaning that under these assumptions there is not the sensitivity required to test the $< 1.2 \mu\text{Bq/kg}$ specification derived from Table 4.2. However, this upper limit is already close to the requirement even under the very conservative hypothesis that all counts falling in the integration window are attributed to this contamination.

Crystal	^{210}Po (bulk) [$\mu\text{Bq/kg}$]	^{235}U (bulk) [$\mu\text{Bq/kg}$]	^{232}Th (bulk) [$\mu\text{Bq/kg}$]	^{238}U (bulk) [$\mu\text{Bq/kg}$]
LMO1	41.9 ± 9.9	≤ 12.4	≤ 12.4	≤ 9.06
LMO2	28.3 ± 8.2	≤ 18.8	≤ 5.4	≤ 5.4
LMO3	44.8 ± 9.8	≤ 4.9	≤ 11.4	≤ 4.9

Table 4.7: Estimated levels of radioactive contamination in the NIIC $\text{Li}_2^{100}\text{MoO}_4$ crystals tested in pre-CCVR1. Since no peaks are observed in $Q \pm 3.5\sigma_\alpha$, limits are quoted at 90% C.L. using the Feldman–Cousins method under a zero-background hypothesis.

4.5 Pre-CCVR2

This section presents the characterization measurements of four natural Li_2MoO_4 crystals produced at SICCAS, named pre-CCVR2.

4.5.1 Goals of the pre-CCVR2 campaign

As outlined in Sec. 4.1.1, during the early pre-production stage radiopurity is not the primary concern, and the stringent requirements in Table 3.2 are not expected to be met. Within this framework, pre-CCVR2 was structured to (i) quantify light yield and calorimetric performance of the first batch of natural Li_2MoO_4 crystals and (ii) supply SICCAS with targeted feedback to optimize growth and processing.

Background data were nevertheless recorded and analyzed to enable relative comparisons of radiopurity among crystals with different crystallization and re-crystallization histories.

Secondary objectives were strictly instrumental to future CCVR operations: testing alternative thermalization schemes and wiring, evaluating the impact of vibrational noise on heat and light channels, and defining a reproducible strategy to assess and compare calorimetric performance of multiple crystals.

4.5.2 First batch of SICCAS natural Li_2MoO_4 crystals

Four natural crystals were tested in pre-CCVR2. In all cases, the final growth was performed by the BG method, but with different numbers and types of re-crystallizations. The raw powders were provided by two producers with different chemical purity levels. Their complete growth history is summarized in Fig. 4.14, where powders from producer "O" and producer "A" are shown in blue and violet, respectively. In all cases, the first synthesis of Li_2MoO_4 powder was performed by wet chemistry.

- 816 (LMO4): This crystal was obtained after two crystallizations, an initial CZ growth followed by a BG growth. After growth, the crystal was cut to $45 \times 45 \times 45 \text{ mm}^3$ and polished with Al_2O_3 powder and a kerosene lubricant using a jeans tissue on the polishing machine.
- 817 (LMO3): This crystal underwent three successive crystallizations. The first was CZ, followed by two BG re-crystallizations. The final crystal was cut to $45 \times 45 \times 45 \text{ mm}^3$ and polished with Al_2O_3 powder and a kerosene lubricant using a jeans tissue on the polishing machine.
- 818 (LMO5): This crystal was produced from two different re-crystallization paths. The main component derived from a double BG re-crystallization, while an additional fraction came from material recycled from a previous CZ growth. This strategy was adopted by SICCAS to reduce isotope losses: leftover ingots from one growth were reused in subsequent ones. This allowed SICCAS to collect data necessary to evaluate the recovery efficiency when recycling strategies are used. The two components were combined in a final BG growth. Afterward, the crystal was cut to $45 \times 45 \times 45 \text{ mm}^3$ and polished with SiO_2 powder and a silicon oil lubricant using a jeans tissue on the polishing machine.
- 819 (LMO6): This crystal was produced from three different precursor sources. About 80% of the material originated from powders of producer "O", split between two BG crystallizations:

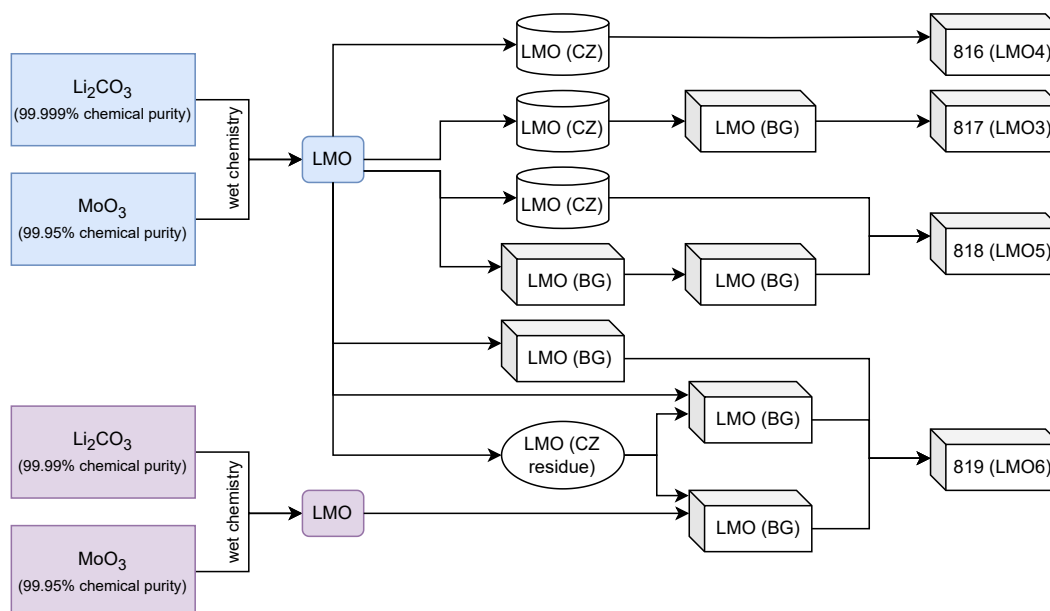


Figure 4.14: Visual scheme representing the raw materials and the re-crystallization history of each of the natural crystals produced by SICCAS and tested in pre-CCVR2. The powders from the two producers are highlighted in blue for producer "O" and violet for producer "A". The cylindrical (parallelepipedal) shapes indicate CZ (BG) growths.

one using purely Li_2MoO_4 obtained by wet chemistry, and one including material recycled from residues left in the crucible after previous CZ growths. The remaining 20% of the material came from a single BG crystallization performed with powders from producer "A", mixed with a small fraction of CZ crucible residues (from producer "O"). SICCAS reported that these CZ residues exhibited a yellowish coloration, consistent with a higher concentration of impurities, as expected from the leftovers of the CZ process. Before being added to the mixture, the residues were dissolved in distilled water, filtered, and dried to obtain Li_2MoO_4 powder (this component is indicated as *CZ residue* in Fig. 4.14). Also in this case, these procedures are used by SICCAS to estimate the recovery efficiency and evaluate the impacts of recycling strategies on crystal quality. After the final growth, the crystal was cut to $45 \times 45 \times 45 \text{ mm}^3$ and polished with SiO_2 powder and a silicon oil lubricant using a jeans tissue on the polishing machine.

Although this growth was carried out with natural MoO_3 powder, the recycling strategies adopted for crystals 818 and 819 are essential for SICCAS to optimize recovery efficiency, a key step to keep costs under control once (costly) enriched powders will be used. At the same time, recycling may compromise crystal quality or radiopurity, introducing non-reproducible effects or instabilities in $\text{Li}_2^{100}\text{MoO}_4$ production batches. This highlights one of the most critical aspects of the pre-CCVR campaign. Testing crystals grown with heterogeneous combinations of precursors provides direct feedback on the reproducibility of crystal quality and radiopurity when recycling conditions are

Powder	Producer-batch	^{232}Th [ng/kg]	^{238}U [ng/kg]	K [$\mu\text{g}/\text{kg}$]
Li_2CO_3	O-20211021	219	213	< 5000
	A-20200907	33722	54.0	28163
MoO_3	O-20211122	1036	3975	94332
	A-20200907	4764	22850	24948

Table 4.8: Radiopurity (ICP-MS) of different powder batches. These measurements were commissioned by SICCAS to the Shanghai Institute of Applied Physics (SINAP).

used. The results of these measurements provide SICCAS with valuable feedback to validate or improve recycling strategies and to recompute production yield estimates in case some recycling paths do not meet the required calorimetric or radiopurity standards.

SICCAS also provided the results of ICP-MS measurements performed on the powders of the two producers, labeled "O" and "A", to assess the content of ^{232}Th , ^{238}U , and K in the raw powders. The results are shown in table 4.8. After cutting and polishing, all crystals were finally cleaned with kerosene and packed for delivery.

The crystals LMO-isot-6-39-top (LMO1) and LMO-isot-6-38-bottom (LMO2) from pre-CCVR1 were also installed in the assembly and served as references. This allowed a direct comparison of performance under identical conditions, reducing uncertainties that may arise when comparing results across different runs, where variations in base temperature or in the noise level of the setup can occur.

4.5.3 Detector design and assembly

The assembly design adopted in pre-CCVR2 differed from that used in pre-CCVR1 described in Sec. 4.4.2. While pre-CCVR1 focused exclusively on testing crystal quality, the pre-CCVR2 campaign also pursued additional CUPID objectives related to thermalization and noise characterization. In particular, one of the goals was to better understand some of the noise features observed in the *Gravity Design Prototype Tower* (GDPT), described in detail in Ref. [155]. The modular design developed for pre-CCVR1 allowed the detector modules to be stacked vertically instead of horizontally, thereby reproducing the configuration of the GDPT. The resulting structure, shown in Fig. 4.15a, consisted of a three-floor tower comprising the six Li_2MoO_4 crystals described earlier and eight interleaved LDs, with the LDs on the intermediate floors shared between two crystals. The spacing of each Li_2MoO_4 and its bottom (top) LD is 0.5 mm (4 mm).

Each Li_2MoO_4 was equipped with Ge NTD thermistors of the same type used in the GDPT demonstrator [155], having dimensions $3 \times 3 \times 1 \text{ mm}^3$, as well as Si heaters. The LDs employed thermistors of the same type, cut to dimensions $3 \times 0.5 \times 1 \text{ mm}^3$, and were also equipped with Si heaters. Two copper plates, placed at the top and bottom, completed the mechanical structure. On the sides, two vertical copper spines were anchored to these plates and supported the total weight of the tower. Fig. 4.15c depicts the tower during construction, after the three modules had been stacked but before mounting the frame with the two top LDs and installing the copper spines, while Fig. 4.15d provides a close-up view highlighting the PTFE elements and copper frames and Fig. 4.15e shows the tower installed inside the Hall C cryostat.

The spines were designed to avoid any mechanical contact with the copper frames, which were intended to remain in position solely by gravity. However, In Ref. [155], a design oversight had caused slight friction between the spines and the frames, which was suspected to induce excess

noise in the detector, particularly in the LDs. This issue was corrected in the pre-CCVR2 assembly to enable a direct comparison of the noise performance of the two configurations. Since this study is not directly related to crystal quality and doesn't affect the evaluation of the light yield, it will not be discussed in great detail in this work.

The vertical spines also supported the wiring of the NTDs and Si heaters. Two different wiring layouts were implemented for the two columns, as described in the next section. All materials used in the assembly were cleaned following the same protocol described in Sec. 4.4.2.

The channel mapping and labeling scheme used in pre-CCVR2 is shown in Fig. 4.15b, and the detectors will be referred to according to these labels in the following. Due to connection issues on LMO3 and LD3, it was not possible to evaluate the calorimetric and scintillation performance or the radiopurity of crystal 817. For the same reason, LD3 could not be used to measure the LY of LMO-isot-6-38-bottom (LMO2).

4.5.4 Wiring and different thermalizations

In the pre-CCVR2 setup, the two columns were equipped with different wiring schemes. The column on the right in Fig. 4.15a used CuPEN (copper on a polyethylene-2,6-naphthalate substrate) wiring, similar to the configuration described in Ref. [155]. A close-up of the CuPEN wiring is shown in Fig. 4.15d. The other column employed constantan wiring, as in the pre-CCVR1 setup (Sec. 4.4.2).

The choice of two different wiring schemes served two purposes. First, while the performance of CuPEN wiring for NTDs coupled to crystals is well established from CUORE, its effect on noise in the LDs with this assembly still needed to be compared against constantan wiring. Second, it allowed a direct comparison of the thermal coupling of the detectors to the MC when using CuPEN versus constantan.

In the detector assembly described in Sec. 4.5.3, the LD frames were not directly connected to any copper elements thermalized to the MC. The only nearby copper structures were the vertical spines, which were mechanically separated from the frames. As a result, the copper frames were thermally floating, with their only effective thermalization paths being the contact between the Ge of the LD and the frame followed by the readout wiring, coupled to the vertical copper spines. The contribution of the PTFE lockers is considered negligible due to its low thermal conductance. In this configuration, the readout wiring is the dominant thermalization path, which differed between the two columns. Unlike in the GDPT demonstrator of Ref. [155], the gold wires connecting the NTDs were bonded directly to the vertical copper spines, without thermal contact with the LD frames.

Three different thermalization schemes were tested in three successive cooldowns:

- **Just readout wiring (v1):** The most basic configuration, in which the only heat-extraction path is through the readout wiring (copper spines), the Ge NTD, the Ge LD, and finally the copper frame.
- **Copper bands (v2):** As in v1, but with additional copper bands connecting the bottom copper plate to the frames, thereby providing a direct thermal link between the frames and the MC. A picture of this configuration is shown in Fig. 4.15e.
- **Gold wires (v3):** Same as v1, but with gold wires linking the vertical copper spines to the frames. This configuration is similar to v2 but uses gold wires instead of copper bands for the thermal link.




Figure 4.15: Geometry of the detector assembly used in the pre-CCVR2 campaign. (a) Full tower. (b) Rendering and detector-channel mapping of the crystals and LDs. (c) Tower during assembly, before positioning the last LD frame. (d) Detail of the ^{55}Fe sources (small copper pieces positioned vertically on the side LD frames). (e) Pre-CCVR2 detectors installed in the Hall C cryostat with the v2 thermalization scheme.

Results

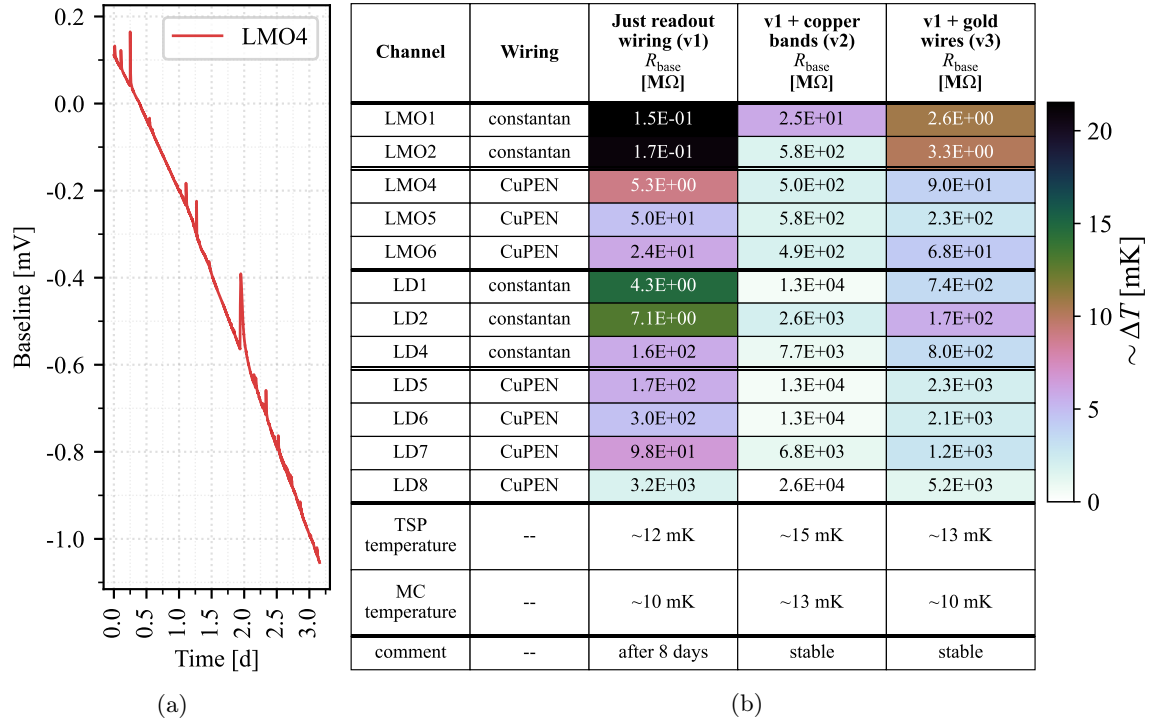


Figure 4.16: (a) V_{NTD} baseline of the NTD on LMO4 (already divided by the preamplifier gain) observed over ~ 3 days (with the thermalization scheme v1). (b) Base resistances (R_{base}) measured for each detector under the three thermalization schemes. Cells are color-coded by the corresponding ΔT , defined with respect to the detector with the highest R_{base} . The ΔT values are not direct temperature measurements but qualitative indicators of thermal coupling. They were computed reversing Eq. (2.8) and using $T_0 = 3.6$ K and $R_0 = 2.4 \Omega$ for NTDs on the LMOs (values from Ref. [155]), and $R_0 = 14.4 \Omega$ for the NTDs on the LDs (accounting for their smaller cross section, which is 1/6 of the LMO NTDs). The table also lists the TSP and MC temperatures for each configuration. For v1, where R_{base} values continued to rise, the values after ~ 8 days are reported.

The base resistances measured in the three configurations are reported in Fig. 4.16b. Although slight differences in MC temperature prevent a strict one-to-one comparison, the observed R_{base} values differ enough to draw clear conclusions.

With thermalization v1, the detectors did not reach high R_{base} . A clear asymmetry was observed between the two columns. On the constantan side, detectors never reached equilibrium with the MC even after several days, showing baselines with a linear descending trend (short spikes due to cryogenic instabilities were disregarded) and continuously rising R_{base} . An example is shown in Fig. 4.16a, where the V_{NTD} (at low bias $I_{\text{bias}} \sim 4$ pA) is plotted over ~ 3 days. Both LDs and LMOs on the constantan side showed significantly lower R_{base} than those on the CuPEN side. Conversely, the CuPEN detectors reached equilibrium with no long-term drift, although their R_{base}

values remained lower than those obtained in the GDPT demonstrator [155] and in configurations v2 and v3.

With thermalization v2, the highest R_{base} values were obtained, demonstrating that thermalizing the copper frames to the MC reduces the thermal load on the detectors and allows them to reach lower temperatures. No long-term drift was observed in this case. A small asymmetry between the two columns persisted: the LMOs had similar R_{base} on both sides, while the LDs on the CuPEN side were slightly colder than those on the constantan side. The highest R_{base} among all configurations was observed in LD8, with $R_{\text{base}} \sim 26 \text{ G}\Omega$.

Thermalization v3 performed slightly worse than v2. All detectors reached somewhat lower R_{base} values, and the asymmetry between columns became more pronounced, especially in the LMOs but also visible in the LDs. As in v2, no long-term drift was observed.

In summary, the assembly could not reach thermal equilibrium with the heat sink unless the copper frames were thermally connected to the MC. Among the tested options, copper bands (v2) provided better thermalization than gold wires (v3).

Calibration data were acquired in all three configurations, and the analysis followed the strategy used for pre-CCVR1 (Sec. 4.4). Background data, however, were collected only in configuration v2. Since the focus here is on the performance of the Li_2MoO_4 crystals, the results discussed in the following refer exclusively to data acquired in v2. The higher R_{base} values in this configuration allowed greater margin for WP optimization and enabled a consistent comparison of crystal performance in terms of sensitivity and energy resolution at different R_{WP} values.

4.5.5 Pre-CCVR2 (v2) data taking

Data-taking for thermalization v2 began in early July 2024 and continued for approximately two months. The first days were dedicated to updating the front-end boards and the electronics required to operate the Si heaters. Due to connection issues, only the heaters on LDs 1–5 and on LMOs 4, 5, and 6 were operational. During this phase, the base resistances of all detectors were repeatedly measured to verify temperature stability. The values of R_{base} for all detectors are reported in the table in Fig. 4.16b.

Two weeks were then devoted to investigating the noise observed in the LDs with the GDPT assembly. Since these studies are not directly related to the crystal quality, only their main results are briefly mentioned in the following. In this period, the working point of the LDs was also optimized to achieve the best possible energy resolution, a crucial step for accurately measuring the LY of the crystals. A comparison of LD performance at different working points is presented in Sec. 4.5.6, while the LY values obtained under optimal conditions are reported in Sec. 4.5.8.

In parallel, the calorimetric performance of the Li_2MoO_4 crystals was characterized by measuring the detector response at different working points. These measurements were performed using the external ^{232}Th calibration sources described in Sec. 4.4.3. The results are presented in Sec. 4.5.7.

Finally, about 12 days of background data were collected to evaluate the radiopurity of the crystals. The results are discussed in Sec. 4.5.9.

Internal ^{55}Fe sources were deployed inside the assembly to calibrate the LDs, an essential step for measuring the LY of the crystals. These sources are similar to those described in Sec. 4.4.3, but due to geometrical constraints the calibration sources for the central and bottom LDs were positioned vertically on the copper frames (see Fig. 4.15d). This placement reduced the geometric efficiency, leading to lower calibration rates for these channels. In contrast, the sources for the top LDs were positioned directly above them, resulting in a more favorable geometry and thus higher

calibration rates.

All measurements were acquired with a sampling frequency of 2 kHz for the LDs and 1 kHz for the LMOs. The Bessel cut-off frequencies were adjusted several times during data taking; when relevant, the chosen values will be reported in the following discussion.

4.5.6 Working point selection and noise in the LDs

Owing to the high R_{base} values (0.2–20 G Ω) of the LDs in the thermalization v2 configuration, the detectors display increased sensitivity to (typically low-frequency) mechanical vibrations together with a higher intrinsic gain. Maximizing the SNR therefore requires finding a trade-off between signal amplitude and detector’s bandwidth. According to the prescriptions derived from the inequality in Eq. (4.1), the objective is to decouple noise from signal by shifting the detector bandwidth to higher frequencies, that is, by reducing the time constant $R_{\text{NTD}}C_{\text{par}}$. This can be obtained either by lowering R_{base} (physically warming the LDs with the heaters) or by increasing I_{bias} to select a WP with lower R_{WP} . These two options¹⁶ were explored, and two WPs were compared:

- **WP5:** the detectors were operated in strong overbias, with a large biasing current ($I_{\text{bias}} = 7.2$ nA (28 V/4 G Ω)). In this configuration, the absorber is warmed by the NTD itself and operates at a higher temperature, resulting in lower intrinsic gain (sensitivity). The strong overbias reduces the working resistance R_{WP} , pushing the detector response to higher frequencies and leading to faster time constants. The pulser cable in this case is unplugged.
- **WP4:** the absorber temperature was artificially increased by injecting constant power through the Si heaters in LDs 1–5. The lower R_{base} in this case allowed for a lower R_{WP} at relatively low I_{bias} . In contrast to WP5, this made it possible to operate the NTDs close to inversion ($R_{\text{dyn}} \rightarrow 0$) without having excessively high working resistances. The chosen biasing current was $I_{\text{bias}} = 0.14$ nA (6 V/54 G Ω). For LDs 6–8, which lacked operational Si heaters, WP4 is identical to WP5, with the only difference that in WP5 the pulser cable is physically disconnected.

In both configurations, the Bessel cut-off frequencies were 160 Hz for LDs 1–5 and 120 Hz for LDs 6–8. The values of I_{bias} , R_{base} , and R_{WP} , together with the corresponding LD performance (intrinsic gain, energy resolution, and time constants) for WP4 and WP5, are reported in Table 4.10.

A few days of data taking were acquired at each WP. The data processing strategy described in Sec. 4.3 was applied to reconstruct the event energies. The baseline and ^{55}Fe energy resolutions were extracted by fitting the calibrated energy spectra of noise and ionization events. An example of such a fit is shown in Fig. 4.17a, corresponding to the energy spectrum of LD8 at WP5. Both the Mn- K_{α} line at 5.89 keV and the Mn- K_{β} line at 6.49 keV were identified in LD8. Unknown background contributions were observed in the left tail of the 5.89 keV peak¹⁷. When correctly modeled with an additional Gaussian component, these contributions do not affect the estimated energy resolution or the centroid position.

¹⁶NTDs with lower R_0 would yield a smaller R_{NTD} at the same absorber temperature, reducing noise impact at fixed signal amplitude and thus improving SNR. Such NTDs were not available during the CUPID pre-CCVR phase.

¹⁷This background is likely due to Auger electrons from the ^{55}Fe source (Mn K- L_1L_2 , Mn K- L_1L_3 , Mn K- L_2L_3 , Mn K- L_3L_3). However, the multi-Gaussian model plus flat background provides a good description of the data in the 3.5–7 keV region for LD8. For this reason, no further studies were performed.

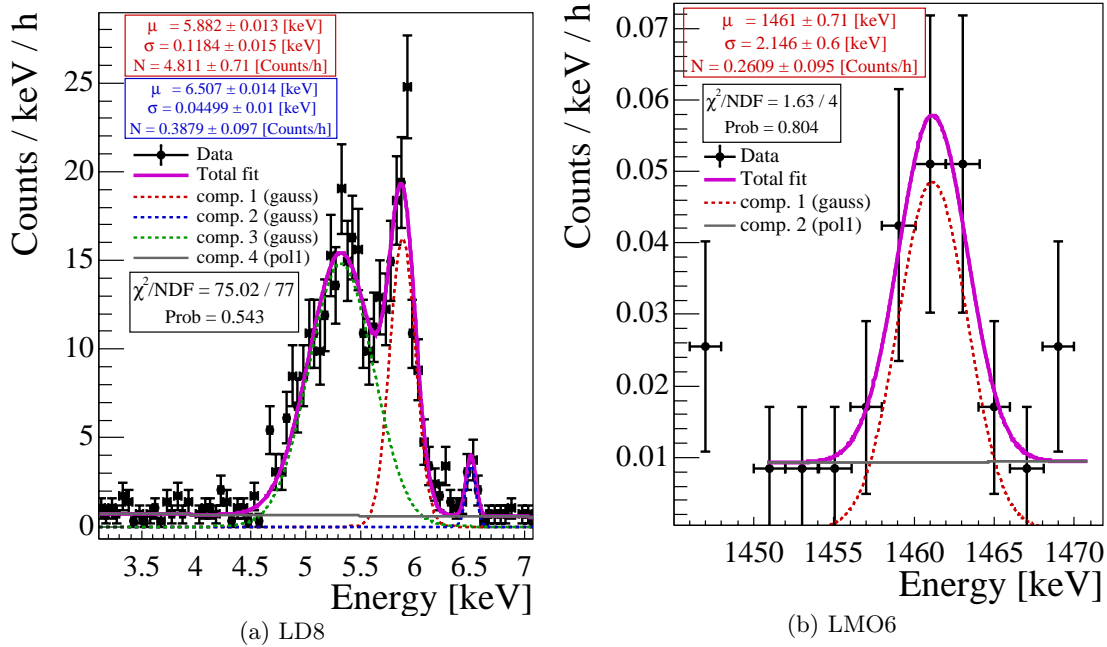


Figure 4.17: Calibration and resolution at the best working point (WP5) of (a) LD8 using ^{55}Fe line and (b) LMO6 using ^{40}K line. Both correspond to a total of ~ 2 days of calibration runs and are not normalized by cuts efficiencies.

Some caveats must be considered before comparing the performance achieved at the two WPs. During the LD noise characterization it was noted that the pulser lines introduced a 1 Hz noise component and higher-order harmonics in the detectors, as shown in Fig. 4.18a. When the pulser cable was disconnected, the 1 Hz harmonics disappeared. A further test was performed by keeping the cable connected but switching off the liquefier GM heads. In this case, the harmonics were reduced, suggesting that the noise was picked up by the pulser cable and related to the operation of the GM heads. This effect was visible in both the constantan and CuPEN columns, but not in the heat channels. WP4 data were acquired with the pulser cable connected, whereas in WP5 the cable was disconnected. This difference must be taken into account in the comparison between the two WPs.

Table 4.10 shows that the performance of LDs 6–8 (without the Si heaters) is consistent between WP4 and WP5, as expected since these detectors were operated in the same conditions. The small differences in energy resolution between the two WPs indicate that, even though the ANPS shows a clear effect from the 1 Hz harmonics, the wOF effectively suppresses these spikes, making them negligible in the filtered signals thereby minimizing their impact on the energy resolution. For LDs 1–5 (having the Si heaters), the baseline energy resolutions obtained at WP5 are significantly better than those at WP4. This indicates that in strong overbias the contribution of noise within the signal bandwidth is reduced, and the wOF is effective in minimizing its impact. At the ^{55}Fe line, the resolutions obtained at WP5 are similar or better than those at WP4, depending on the

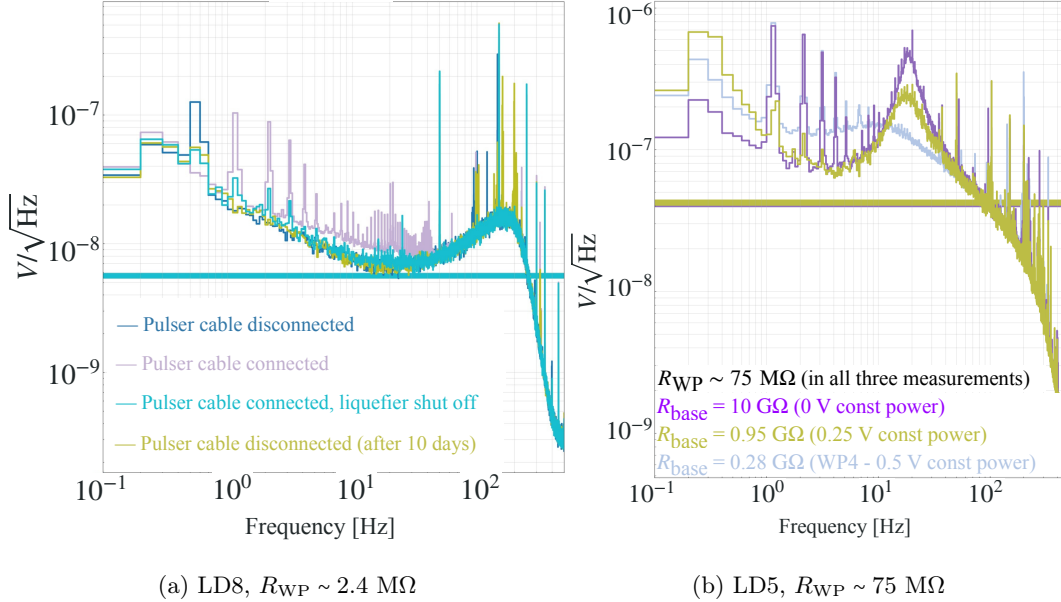


Figure 4.18: (a) ANPS in LD8 at WP5 ($I_{\text{bias}} = 7.2 \text{ nA}$ ($28 \text{ V}/4 \text{ G}\Omega$)) under different configurations of the pulser cable and liquefier GM head operation. The case with both disconnected was repeated after 10 days to check reproducibility. (b) ANPS in LD5 at different R_{base} and similar R_{WP} , obtained by tuning the Si heater power and adjusting bias currents. The pulser cable remained connected. In both panels, the solid horizontal line indicates the noise floor from Eq. (2.14).

detector.

The intrinsic gain observed in WP4 is on average higher than in WP5, despite the higher R_{base} in the latter. This is consistent with the fact that in WP5 the strong overbias warms the absorber, reducing signal amplitude. However, the electrothermal feedback increases the detector bandwidth, as shown by the shorter RT and DT observed in WP5 compared to WP4. This is also consistent with the fact that WP5 has lower R_{WP} . It should be noted that in WP5, the RT values are limited by the Bessel cut-off frequency, corresponding to $\sim 2.1 \text{ ms}$ for LDs 1–5 (160 Hz) and $\sim 2.9 \text{ ms}$ for LDs 6–8 (120 Hz), so the intrinsic detector RT may be even shorter. These considerations indicate that the LDs perform better at WP5 than at WP4.

A direct comparison between the two WPs, however, is complicated by the fact that WP4 was designed to operate the detectors close to inversion, whereas WP5 was chosen as the strongest overbias achievable with the available setup. As a result, the working resistances R_{WP} are different in the two WPs. One might argue that the better performance observed in WP5 is solely due to the lower R_{WP} . A fairer comparison would be between two WPs with the same R_{WP} but different R_{base} . Nevertheless, due to electrothermal feedback, the detector bandwidth is determined not only by R_{WP} but also by R_{dyn} . Two NTDs operating at the same R_{WP} but with different R_{base} are effectively in different working conditions, corresponding to different positions on the $I - V$ curve. To qualitatively illustrate this effect, the ANPS for LD5 was compared among three WPs where $R_{\text{WP}} \sim 75 \text{ M}\Omega$ but at different R_{base} values of 0.28, 0.95, and 10 G Ω (set by varying the

constant power on the Si heaters). The pulser cable was kept connected in all three cases, so that the effect of the 1 Hz harmonics was identical, allowing for direct comparison. The results, shown in Fig. 4.18b, indicate that the detector frequency response strongly depends on R_{base} . In particular, when R_{WP} is obtained by operating in strong overbias, the response is pushed to higher frequencies, reducing sensitivity to low-frequency noise such as vibrations. In contrast, when R_{WP} is obtained by lowering R_{base} (warming the absorber with the Si heater and operating at lower I_{bias}), the bandwidth is reduced, integrating more low-frequency noise and worsening the impact of vibrations. This can be appreciated by looking at the flat component of the noise at lower frequencies in the ANPS in the three cases.

The evidence collected in pre-CCVR2 suggests a consistent picture: with the available NTDs, the high R_{base} values achieved under good thermalization to the MC make the LDs sensitive to low-frequency noise, but this can be mitigated by lowering R_{NTD} . Although the 1 Hz harmonics in the WP4 data complicate the comparison, there is experimental evidence that operating NTDs in strong overbias at higher R_{base} is preferable over biasing near the inversion region at lower R_{base} . The LY data of the Li_2MoO_4 crystals were therefore collected with the LDs operated at WP5.

This study thus conclusively defined the operating conditions for future pre-CCVR runs and CCVR measurements with these LDs and NTDs, thereby finalizing the LD-related component of the CCVR protocol and enabling reliable LY measurements of the Li_2MoO_4 crystals produced by SICCAS.

4.5.7 Li_2MoO_4 calorimetric performance and WP optimization

Although a systematic procedure for optimizing NTD-equipped cryogenic calorimeters was already developed for CUORE TeO_2 crystals [124], Li_2MoO_4 absorbers differ from TeO_2 (as Li_2MoO_4 has a more favourable thermal capacity) and therefore require an ad hoc protocol. Crucially, the detector-response characterization needed for WP optimization coincides with that required to assess calorimetric performance, which is the central goal of the CCVR (and pre-CCVR) campaign. The measurements presented below thus serve a dual role: they provide the basis for a systematic WP optimization of Li_2MoO_4 crystals and establish a calorimetric-performance protocol to be adopted as the standard for crystal validation in the CUPID CCVR program.

Detector response and WP optimization

In the pre-CCVR1 campaign (Sec. 4.4.8), the low R_{base} values deriving from poor thermalization left little room for WP optimization, given the already warm operating conditions of the detectors. By contrast, in pre-CCVR2 with the thermalizations v2, the R_{base} values are of the order of 0.5 G Ω , which allows to perform systematic scan of the detector performance as a function of the WP and a more complete assessment of the calorimetric quality of the crystals. This also enables a clearer disentanglement of NTD response from intrinsic crystal properties, improving the reliability of crystal-to-crystal comparisons. All the heat channels were therefore characterized through noise scans, I - V and R - P curves, and monochromatic pulser injections with the Si heaters. In the following, results are shown for LMO5 as a representative example; the same measurements were carried out on all crystals, which exhibited comparable behavior.

Figure 4.19 shows the evolution of the AP and the ANPS as a function of the working point. The corresponding WP specifications (I_{bias} , V_{bias} , R_L , V_{NTD} , and working resistance) are reported

LMO	WP1 – Scan 6: 0.54 nA (28.9 V / 54 GΩ)						WP5 – Scan 4: 0.24 nA (12.9 V / 54 GΩ)						WP4 – Scan 2: 0.03 nA (1.82 V / 54 GΩ)						QF _α
	R _{base} [MΩ]	R _{WP} [MΩ]	FWHM bsl. [keV]	Peak FWHM res. [keV]	Sens. [nV/keV]	RT [ms]	DT [ms]	R _{WP} [MΩ]	FWHM bsl. [keV]	Peak FWHM res. [keV]	RT [ms]	DT [ms]	R _{WP} [MΩ]	FWHM bsl. [keV]	Peak FWHM res. [keV]	RT [ms]	DT [ms]		
1	25.1	11.1	6.58 ± 0.10	9.6 ± 2.6 @ 2615	13	31	220	11.7	3.65 ± 0.03	x	40	232	23.7	35.8 ± 0.18	x	50	230	1.054	
2	579.6	22.2	2.98 ± 0.05	5.8 ± 0.9 @ 2615	23	40	300	50.8	x	x	x	x	273.7	9.1 ± 0.2	11.6 ± 2.8 @ 618	112	378	1.046	
4	498.3	20.1	6.15 ± 0.10	6.9 ± 0.9 @ 2615	12	25	200	44.7	5.09 ± 0.06	5.9 ± 2.2 @ 2615	36.5	222	242.5	21.3 ± 0.2	34.1 ± 8.9 @ 1460	93	340	1.076	
5	579.6	20.6	5.20 ± 0.09	6.3 ± 0.7 @ 2615	17	30	350	47.3	3.89 ± 0.04	7.8 ± 3.9 @ 2615	43.2	390	259.7	10.2 ± 0.2	15.8 ± 2.3 @ 1460	120	560	1.057	
6	487.4	19.5	4.31 ± 0.37	6.3 ± 2.4 @ 2615	25	25	150	42.9	3.17 ± 0.07	5.0 ± 1.2 @ 1460	35.6	174	229	8.2 ± 0.2	10.9 ± 1.9 @ 1460	76	260	1.078	

Table 4.9: Comparison of the performance of LMOs (LMO1 and LMO2 from NIIC and LMO4, LMO5, LMO6 from SICCAS) operated at WP1, WP4 and WP5. The empty cells indicate that limited statistics prevented to perform reliable fits of the corresponding lines. Where no uncertainties are reported, the values should be considered as rough indications of the performance. The uncertainties for these values typically range between 1% and 20%, depending on the shape of the distributions and the available statistics.

LD	WP4								WP5							
	R _{base} [GΩ]	R _{WP} [MΩ]	WP: $I_{\text{bias}}(V_{\text{bias}}/R_L) + \text{heater}$	Sens. [μV/keV]	σ bsl [eV]	σ ⁵⁵ Fe [eV]	RT [ms]	DT [ms]	R _{base} [GΩ]	R _{WP} [MΩ]	WP: $I_{\text{bias}}(V_{\text{bias}}/R_L)$	Sens. [μV/keV]	σ bsl [eV]	σ ⁵⁵ Fe [eV]	RT [ms]	DT [ms]
1	0.31	107	0.14 nA (6 V / 54 GΩ) + const. power	2.45	295	320	10.6	26.5	13.1	2.06	7.23 nA (28 V / 4 GΩ)	1.23	85.6	338	2.5	5.1
2	0.03	27	0.14 nA (6 V / 54 GΩ) + const. power	1.53	280	537	7.0	19.5	2.6	1.30	7.23 nA (28 V / 4 GΩ)	1.18	78.5	200	2.5	7.5
4	0.12	77	0.14 nA (6 V / 54 GΩ) + const. power	2.45	276	265	9.1	25.0	7.7	1.80	7.23 nA (28 V / 4 GΩ)	1.19	83.1	120	2.45	5.1
5	0.28	76	0.14 nA (6 V / 54 GΩ) + const. power	2.43	290	271	10.5	26.8	12.9	1.50	7.23 nA (28 V / 4 GΩ)	x	x	x	2.15	5.45
6	13.25	1.5	7.23 nA (28 V / 4 GΩ)	0.92	86.6	100	3.3	6.0	13.3	1.50	7.23 nA (28 V / 4 GΩ)	0.86	84.5	264	2.8	5.9
7	6.80	1.2	7.23 nA (28 V / 4 GΩ)	0.50	86.7	362	3.3	7.0	6.8	1.22	7.23 nA (28 V / 4 GΩ)	0.46	86.3	362	2.8	7.0
8	26.00	2.4	7.23 nA (28 V / 4 GΩ)	1.01	100	143	3.3	6.0	26.0	2.39	7.23 nA (28 V / 4 GΩ)	1.01	87.1	119	2.75	5.9

Table 4.10: Comparison of the performance of LDs operated at WP4 and WP5. For LD5, the limited statistics at the ⁵⁵Fe peak prevented the determination of the intrinsic gain (sensitivity) and the energy resolution. Where no uncertainties are reported, the values should be considered as rough indications of the performance. The uncertainties typically range between 1% and 20%, depending on the shape of the distributions and the available statistics. These results are discussed in Sec. 4.5.6. Note that the RT observed at WP5 is limited by the Bessel cutoff. LD6, LD7 and LD8 did not have operational Si heaters.

in Fig. 4.20, whose legend uses the same colour code for direct comparison. The left panel of that figure shows a base resistance of $R_{\text{base}} \sim 579 \text{ M}\Omega$ for LMO5.

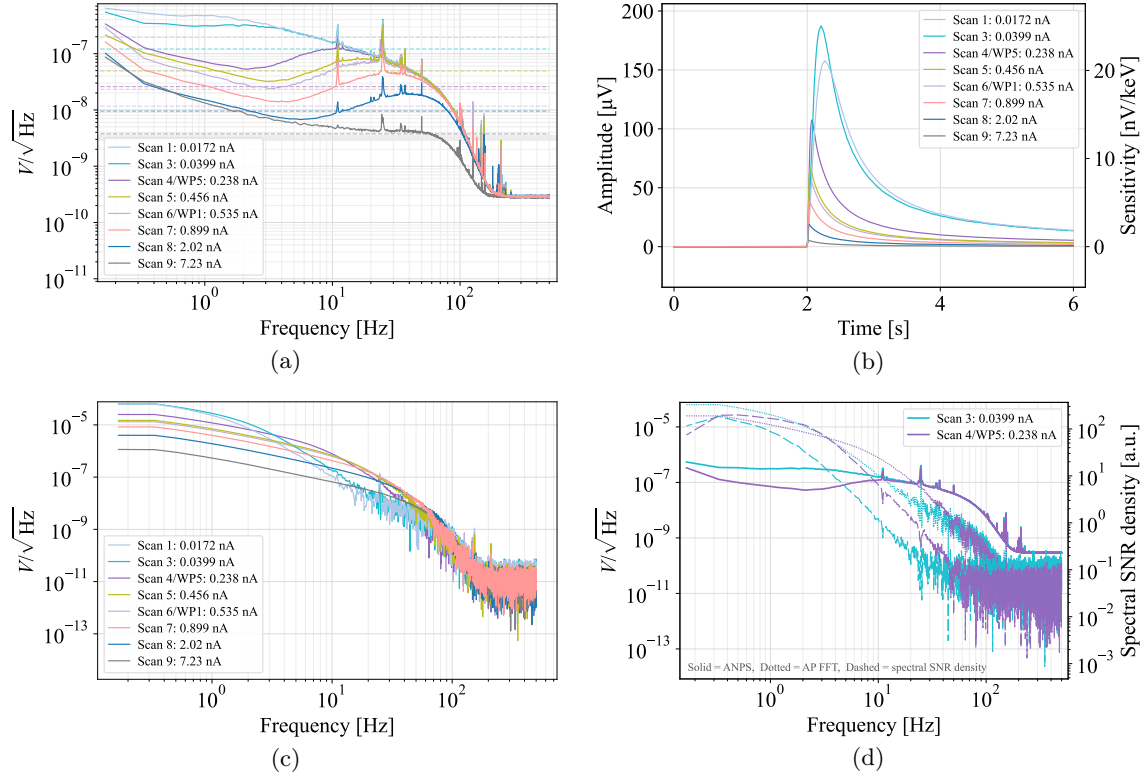


Figure 4.19: Evolution with WP of (a) the ANPS, (b) the pulser AP in the time domain, and (c) in the frequency domain for LMO5. The I_{bias} for each WP is shown in the legend. In (a), the grey band represents the white component of the pre-amplifier noise, and dashed lines show the noise floor from Eq. (2.14). The AP in (b) and (c) was obtained by injecting monochromatic pulser pulses corresponding to an equivalent particle energy of $\sim 7480 \text{ keV}$. (d) Comparison of ANPS (solid) and AP (dotted) for Scan 4 (WP5) and Scan 3 (close to WP4). The spectral SNR density $S(\omega)/\sqrt{N(\omega)}$ (dashed) shows that decoupling noise and signal frequencies enhances the SNR, which in Scan 4 is higher than in Scan 3 across almost the entire bandwidth. The values of working resistance and V_{NTD} , as well as V_{bias} and R_L , are reported in Fig. 4.20, whose legend uses matching colours.

Figure 4.19b illustrates the evolution of the AP for a monochromatic pulse injected by the pulser, corresponding to an equivalent particle energy of approximately 7480 keV ¹⁸. Both the intrinsic gain and the pulse shape vary with the WP. For the lowest I_{bias} (Scan 1, $R_{\text{WP}} \sim 350 \text{ M}\Omega$), the average pulser amplitude is about $155 \mu\text{V}$ (corresponding to $\sim 20.7 \text{ nV/keV}$), and the pulse exhibits slower time constants. Slightly increasing I_{bias} (Scan 3, $R_{\text{WP}} \sim 225 \text{ M}\Omega$) leads to a higher amplitude, which peaks at about $180 \mu\text{V}$ (corresponding to $\sim 24 \text{ nV/keV}$). Further increasing I_{bias} makes the pulses

¹⁸The energy was reconstructed from the pulser peak observed in the spectrum at WP4 calibrated with γ lines (Fig. 4.26).




Figure 4.20: *Left panel*: R - P curve, pulser amplitude, and relative FWHM energy resolution for monochromatic pulses injected by the Si heater. The coloured vertical bands indicate the I_{bias} (V_{NTD} / R_L) configurations reported in the legend on the right. The outlier close to the minimum of the FWHM resolution curve is attributed to the switch from the large to the small load resistance R_L . *Right panel*: I - V curve and corresponding pulser amplitude evolution. The points where a full characterization of AP, ANPS, and FWHM pulser resolution was performed are marked with coloured circles matching the WP legend. Both panels refer to LMO5. The pulser amplitude and FWHM energy resolution correspond to monochromatic pulses of ~ 7480 keV, so the amplitude in both panels is proportional to the intrinsic gain (sensitivity) of the crystal at each WP in the scan.

faster and reduces the signal amplitude as the electrothermal feedback starts kicking in, reducing R_{WP} . This behavior is clearly visible in Fig. 4.20, where the evolution of the pulser amplitude, measured after the OF, is shown as a function of the power P (left panel) and of V_{NTD} (right panel). Figure 4.19c shows the evolution of the AP in the frequency domain, where it can be seen that the spectral content of the signal shifts toward higher frequencies as I_{bias} increases, indicating that pulses become faster and their time constants shorter.

These observations show that the pulse is shaped by the electrothermal transfer function of the detector. The same transfer function, and its evolution with WP, also governs the detector response to noise, thereby defining the shape of the ANPS. Since, to first approximation, the energy resolution is determined by the SNR, the I_{bias} that maximizes the signal amplitude (Scan 3) does not necessarily correspond to the optimal resolution. For this reason, the noise and its evolution with WP must also be examined. Figure 4.19a shows the evolution of the ANPS at different WPs. A clear trend is observed: at lower bias currents (e.g., Scans 1, 2, and 3), the noise level at low frequencies lies above the noise floor (dashed lines). As I_{bias} increases, R_{WP} reduces, the detector bandwidth widens, the time constants shorten, and the sensitivity to low-frequency noise decreases, leaving a residual $1/f$ -like contribution. This trend continues up to the strongest overbias (Scan 9), where the noise spectrum is ultimately shaped by the Bessel cut-off, which was set to 47 Hz for all these measurements.

Figure 4.19d compares the ANPS and the AP in the frequency domain for Scan 3 ($R_{WP} \sim$

225 M Ω), which corresponds to the maximum signal amplitude, and Scan 4 ($R_{\text{WP}} \sim 47.3$ M Ω), where the NTD operates at much lower resistance, at its inversion region. The same plot also shows a figure of merit, $S(\omega)/\sqrt{N(\omega)}$, which can be interpreted as an approximate measure of the SNR spectral density. Since this quantity is higher for Scan 4 than for Scan 3, it can be qualitatively inferred that shifting the detector response toward higher frequencies improves the SNR, even at the cost of a strongly suppressed signal amplitude. This behavior contrasts with that observed in TeO₂ crystals, where the selected optimal WP was found at a bias point with only a moderate reduction in signal amplitude [124]. However, differences in absorber mass, NTD type, and base temperature in Ref. [124] preclude a straightforward comparison with TeO₂, highlighting the need for custom configuration-specific studies like the one presented here.

These qualitative results indicate that WP optimization is nontrivial, as it requires a trade-off between signal amplitude and frequency response. To quantitatively determine the optimal WP, a dedicated measurement campaign was carried out by operating the detectors in Scans 1 to 9 for several hours each. The analysis steps described in Sec. 4.3 were applied to all acquired data, and the ANPS and AP shown in Fig. 4.19 were used to build the OF. After stabilization, the pulser peaks in the spectra were fitted with a Gaussian model to extract the relative FWHM pulser energy resolution. This procedure allows the identification of the optimal WP and provides a quantitative comparison of the effects of signal amplitude and detector transfer function in terms of SNR.

The orange curve in the left panel of Fig. 4.20 shows the relative pulser resolution as a function of the power P , with the curve exhibiting a clear minimum that indicates the existence of an optimal WP (in terms of energy resolution). It is evident that the relative energy resolution near the maximum intrinsic gain (Scan 2 and 3) is far from optimal. By contrast, increasing I_{bias} reduces the signal amplitude but, by minimizing the overlap of noise with the signal bandwidth, improves the pulser energy resolution after application of the OF¹⁹.

The experimental evidence presented here shows that WP optimization requires a trade-off: higher bias currents reduce the intrinsic gain but at the same time reshape the detector response, mitigating the impact of low-frequency noise. The optimal WP is therefore not the one with the largest signal amplitude, but the one that yields the best energy resolution by balancing intrinsic gain and noise-shaping effects. Accordingly, the methodology presented here provides a well-defined reproducible procedure for determining the optimal working-point configuration in following pre-CCVR campaigns and in future CCVR measurements.

Calorimetric performance of natural pre-production SICCAS Li₂MoO₄

The FWHM obtained from pulser events is useful for studying the evolution of the detector resolution with WP and can provide an indirect estimate of the intrinsic gain through the equivalent energy. However, the true detector energy resolution must be extracted from real particle interactions. For this reason, three WPs were selected to operate the detectors for multiple days to acquire calibration spectra with sufficient statistics. These are referred to as WP1 (Scan 6), WP5 (Scan 4), and WP4 (Scan 2, very close to Scan 3). The corresponding R_{WP} for all the detectors are listed in Table 4.9. WP4 corresponds to the maximum signal amplitude, WP5 lies close to the inversion point, and WP1 to the strongest overbias achievable with the largest available load resistor²⁰, lying

¹⁹The same plot was also built before applying the optimum filter. In that case, the trend of the pulser resolution is less clear, highlighting the role of the OF in these studies.

²⁰Using the lowest load resistor would allow a further increase of I_{bias} , but at the cost of a higher noise floor.

beyond inversion. These WPs are highlighted in both the I - V characteristic and the R - P curve in Fig. 4.20. In all cases, data were processed with the analysis flow of Sec. 4.3 up to the heat-channel calibration.




Figure 4.21: Calibration spectrum in SICCAS crystals (γ region) operated at WP1 corresponding to a total live time of ~ 4 days, acquired with external ^{232}Th sources. The rate is not corrected by cuts efficiencies.

The calibration spectrum obtained after ~ 4 days at WP1 is shown in Fig. 4.21. Similar spectra were obtained at WP4 and WP5.

Energy resolution The calorimetric performance of the SICCAS crystals was evaluated against the NIIC reference by fitting the gamma lines measured at each WP. Priority was given to the ^{208}Tl 2614.5 keV line; when statistics were insufficient, alternative γ peaks were fitted. Figure 4.17b shows a fit of the ^{40}K line for LMO6 at WP5. The results of all fits are summarized in Table 4.9. In some cases limited statistics resulted in large uncertainties. Moreover, comparing energy resolutions obtained at different gamma energies is not formally correct. Even with these limitations, several observations can be made:

- Both the average gamma-line and the baseline FWHM at WP4 are significantly worse than at WP1 and WP5. Within the available uncertainties, the resolution at gamma-lines at WP1 is similar to that at WP5, consistent with the trend of the relative pulser resolution shown in the left panel of Fig. 4.20. The FWHM baseline resolution at WP5 is systematically better than the one at WP1 for all channels.
- WP1, which is close to the optimal WP and has sufficient statistics at ^{208}Tl 2614.5 keV, allows a direct comparison among SICCAS crystals and with NIIC. The resolutions are comparable across all crystals, with the only exception of LMO1 (NIIC). The poorer performance of LMO1 is attributed to its lower R_{base} rather than to crystal quality. The three SICCAS crystals are mutually consistent.
- FWHM baseline resolution differs across crystals and varies with WP. For instance, at WP1 LMO1 shows a worse baseline than LMO4 but remains statistically compatible, whereas at WP5 LMO1 is better and statistically incompatible with LMO4. Since the baseline reflects

the residual noise after the OF, these variations are attributed to the WP dependence of the noise; ANPS studies analogous to Fig. 4.19a support this. Thus, the observed differences in baseline resolution not linked to crystal quality and have limited impact on gamma-line resolutions, which in macrocalorimeters are often dominated by additional non-idealities beyond the baseline noise.

Overall, the SICCAS Li_2MoO_4 crystals exhibit mutually consistent energy resolutions despite their very different growth histories, indicating that differences in material recycling and re-crystallizations (Figure 4.14) do not significantly affect calorimetric performance. Their resolution is comparable to the NIIC reference crystals, at the level of ~ 6.5 keV FWHM at WP1, close to the CUPID goal. Further improvement is expected under reduced noise conditions in the CUPID cryostat. Although limited statistics increase the uncertainties, detectors operated at WP5 perform slightly better than at WP1; notably, LMO6 achieves a FWHM of ~ 5 keV at the ^{40}K γ line.

Intrinsic gain (sensitivity) The sensitivity was first evaluated at WP1, chosen because it lies close to the optimal WP and provides sufficient statistics at the ^{208}Tl 2614.5 keV line in all channels. Since, as discussed above, sensitivity depends strongly on the WP, these values are meaningful only when considered together with the R_{base} and R_{WP} of the configuration. The results are reported in Table 4.9. Among the NIIC crystals, the lower sensitivity of LMO1 compared to LMO2 is attributed to its lower R_{base} and R_{WP} . LMO2, having values closer to those of the SICCAS crystals, is therefore taken as the appropriate NIIC reference. LMO2 shows sensitivity comparable to LMO6, while LMO4 and LMO5 have slightly lower values. However, the R_{base} values are not identical and the sensitivity uncertainty is difficult to quantify. Within these limitations, LMO6 appears marginally better than the other two SICCAS crystals.

To verify that this result is not just a fluctuation related to the WP choice, the pulser equivalent energy was also used to study the evolution of the sensitivity as a function of the working point. This cross-check could be performed only on the SICCAS crystals (LMO4, LMO5, and LMO6), since the NIIC reference crystals had non-operational Si heaters. In this case, the sensitivity was obtained by fitting the pulser peak in the γ -calibrated spectra at WP4 and converting it to the equivalent energy. Although heater pulses are not fully identical to particle interactions, they provide sufficient statistics in a short measurement time to study the evolution of sensitivity with WP. The results of this scan (Scans 1–9) are shown in Fig. 4.22. Despite some fluctuations, particularly for LMO4 and LMO6, the trend confirms that LMO6 consistently maintains slightly higher sensitivity across all WPs.

Time constants As one can read from Table 4.9, WP1 exhibits the shortest RT. WP5 shows slightly longer RT on average, and WP4 the longest. This is consistent with the faster detector response expected from the wider bandwidth at lower working resistance, which to first approximation scales with the time constant $R_{\text{WP}}C_{\text{par}}$. The wider bandwidth at WP1 and WP5 also results in slightly shorter DTs. Across WPs, the three SICCAS crystals show small but systematic differences: LMO6 has the shortest DT, LMO4 is similar but slightly longer, and LMO5 has the longest DT. LMO5 also shows consistently longer RT, while RT values for LMO4 and LMO6 are similar. These differences persist systematically across all WPs, indicating that they are not merely an artifact of working point selection. However, with the available data it is not possible to identify a unique cause. The three crystals share the same CuPEN wiring, so differences in the thermal conductance to the heat sink could arise from small variations in the NTD gluing to the crystal




Figure 4.22: Evolution of the intrinsic gain (sensitivity) with WP, shown as a function of the power P (left) and the working resistance R_{WP} (right) for LMO4 ($R_{base} \sim 498 \text{ M}\Omega$), LMO5 ($R_{base} \sim 579 \text{ M}\Omega$) and LMO6 ($R_{base} \sim 487 \text{ M}\Omega$). Sensitivities were derived from Si-heater pulses by determining their equivalent energy from the fitted pulser lines in the γ -calibrated spectrum at WP4.

surface. Different exposure to humidity during surface treatment and packaging could also play a role, since no strict humidity control was enforced. Differences due to intrinsic crystal properties cannot be excluded, but present data are not conclusive. LMO2, the only NIIC crystal with similar R_{base} , shows time constants in line with SICCAS crystals; however, a direct comparison is limited by its constantan wiring, while SICCAS crystals use CuPEN.

These considerations indicate that the first batch of pre-production SICCAS crystals achieves calorimetric performance comparable to NIIC references, with no significant differences among the three SICCAS crystals despite their different re-crystallization paths.

Constructing the plots in Fig. 4.22 that track sensitivity versus WP, expressed both in terms of power (left) and in terms of R_{WP} (right), proved effective for comparing the intrinsic calorimetric performance of the crystals, provided that the detectors share the same NTD type and operate at the same absorber temperature (that is, the same R_{base}). For robust assessment in future pre-CCVR campaigns and during CCVR measurements, at least one reference crystal of known quality should be included and instrumented with the same NTDs and Si heaters, so that identical WP scans can be performed at similar R_{base} . This procedure provides a consistent and reproducible basis for assessing the intrinsic calorimetric performance of Li_2MoO_4 crystals and for isolating WP effects from detector-to-detector variations. Comparisons based solely on γ -line energy resolutions are not reliable indicators of crystal quality, since the resolution depends on both signal amplitude and noise, the latter arising from extrinsic sources (see Section 2.3.1) that may vary from channel to channel. Nevertheless, γ -line resolutions offer complementary information that, when combined with the sensitivity study described above, enables a complete characterization of detector performance with

respect to the CUPID requirements. This finalizes the definition of the component of the CCVR protocol dedicated to assessing the intrinsic calorimetric performance of the crystals.

4.5.8 Light Yield

The crystal LY was measured at WP5, where the LDs achieved their best performance. The LDs were calibrated on the ^{55}Fe line; an example fit for LD8 is shown in Fig. 4.17a. The calibration accuracy was strongly affected by the placement of the internal sources: when mounted vertically on the copper frames (Figure 4.15d), the geometric efficiency dropped, reducing the rate and making several calibrations difficult or uncertain. This is critical, since the LD calibration enters directly in the LY determination. In particular, LD2 could not be calibrated, and with LD3 not operational, the LY of LMO2 could not be measured. The low ^{55}Fe rate also impacted LDs 1, 5, 6, and 7, where limited statistics hindered background modeling near the left tail of the $\text{Mn-}K_{\alpha}$ line. Residual events on that tail can bias the centroid, introducing a systematic uncertainty on the LD calibration and thus on LY. Quantifying this effect is challenging. As a rough estimate, one can fit the calibration peak of LD8 and LD4 with a single-Gaussian model, instead of the multi-component model used in Figure 4.17a, and compare the results. The resulting energy-scale shift is of order $\sim 10\%$. Although qualitative, this shift is much larger than the statistical uncertainties and limits the sensitivity of quantitative comparisons between crystals.

A few days of data were acquired with external ^{232}Th sources to increase β/γ statistics. After applying the full processing sequence shown in Fig. 4.2, the LY vs heat scatter plots were produced, following the same procedure as in pre-CCVR1 (Sec. 4.4.9). Due to the lower statistics compared to pre-CCVR1, the energy cut used to build the LY distributions was set to $E > 1.5$ MeV. The results are shown in Fig. 4.23, where the top row displays the LY vs heat scatter plots and the other rows the LY distributions for each channel. The LY for each crystal is reported in Table 4.11, with statistical uncertainties only. The table also lists the width of the β/γ distribution; this quantity directly influences the DP (Eq. 2.18) and depends on both the LD energy resolution and the LY itself, as it will be discussed in Section 5.4.3. Table 4.11 also reports the DP for each heat–light channel pair, as well as the DP obtained by summing the light collected by the top and bottom LDs, which is the strategy that will be adopted by CUPID.

The measured DP (top + bottom) is close to the CUPID target of $\text{DP} > 5$ around 3 MeV. Since CUPID will employ NTL-assisted LDs, this value is expected to further increase.

Although the achieved discrimination power is very good, the DP is not an effective metric for validating crystal quality as it depends on the LD performance in addition to the light yield of the Li_2MoO_4 . Since the only crystal-dependent contribution to DP is the LY, this is the only meaningful quantity to consider for validating Li_2MoO_4 in the CCVR program.

The mean LY (uncorrected for light-collection efficiency) measured by the bottom (top) LD for the first batch of natural pre-production SICCAS crystals is ~ 0.30 (0.27) keV/MeV. A direct comparison with the ~ 0.35 keV/MeV requirement in Section 3.1 provides limited information, since this target value was derived from the GDPT demonstrator [155] without accounting for position-dependent effects in the setup, such as reflections inside the inner copper thermal shield of the cryostat and other passive elements. For this reason a NIIC reference crystal is always included in each run and serves as an internal anchor for relative LY comparisons. The GDPT demonstrator, which featured modules of comparable geometry, documented crystal-to-crystal variations at the level of ~ 0.04 keV/MeV [155]. Considering also the LD calibration uncertainty induced by the reduced ^{55}Fe rate, the LY measured for the SICCAS crystals (LMO4, LMO5, LMO6) is consistent

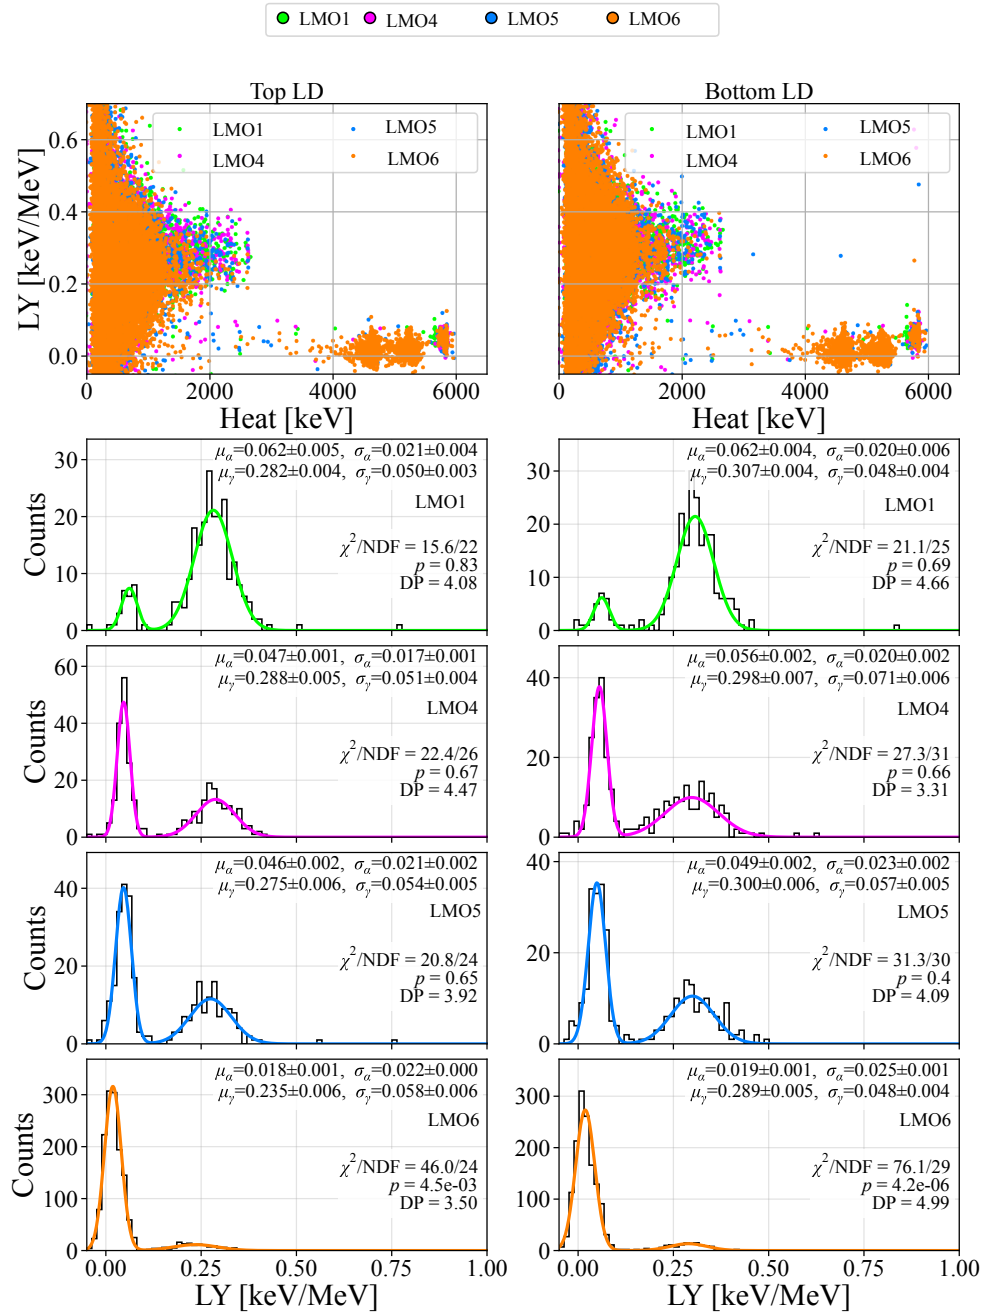


Figure 4.23: Left column: top LD; Right column: bottom LD. LY vs heat (top row) and LY distributions for $E > 1.5$ MeV (bottom four rows).

Crystal	LD	LY (β/γ) [keV/MeV]	LY width (β/γ) [keV/MeV]	DP	DP (Top+Bottom)
LMO1	Top – LD2	0.282(4)	0.050(3)	4.08	5.60
	Bottom – LD1	0.307(4)	0.048(4)	4.66	
LMO4	Top – LD6	0.288(5)	0.051(4)	4.47	4.97
	Bottom – LD5	0.298(7)	0.071(6)	3.31	
LMO5	Top – LD7	0.275(6)	0.054(5)	3.92	4.89
	Bottom – LD6	0.300(6)	0.057(5)	4.09	
LMO6	Top – LD8	0.235(6)	0.058(6)	3.50	5.56
	Bottom – LD7	0.289(5)	0.048(4)	4.99	

Table 4.11: LY distribution mean and width, and DP for the SICCAS Li_2MoO_4 crystals (LMO4, LMO5 and LMO6) and for one reference NIIC crystal (LMO1) as measured in CCVR2. The values are computed selecting events with $E > 1.5$ MeV. The reported uncertainties are statistical only. The last column lists the DP obtained considering the total light collected simultaneously by the top and bottom light detectors is considered.

with the NIIC reference in this run (LMO1) and with the pre-CCVR1 values reported in Table 4.6, within the aforementioned crystal-to-crystal spread. These values are also compatible, within about two times this spread, with the top and bottom LY reported for the GDPT geometry [155].

4.5.9 Background measurement and radiopurity

As discussed in Sec. 4.5.1, no strict radiopurity protocol was applied by SICCAS for the growth, surface treatment and packaging of this first batch of pre-production natural crystals. The raw powders used for this Li_2MoO_4 synthesis were not meeting the CUPID radiopurity requirements. For these reasons, the main goal of this run was assessing calorimetric performance rather than radiopurity. However, a background measurement can still provide an initial feedback about the radiopurity observed for crystals produced by different recycling strategies and re-crystallization histories.

The external calibration sources had been removed and a total livetime of approximately 12 days of background data had been acquired with the detectors set at WP4²¹. More precisely, the acquired live times were: LMO1 (295.996 h), LMO2 (268.568 h), LMO4 (282.770 h), LMO5 (283.881 h), and LMO6 (284.436 h). The data processing strategy described in Sec. 4.3 was then used to produce calibrated energy spectra for all the heat channels. The LY was not used to discriminate α interactions in this case.

Data quality cuts and selection efficiency

Above threshold, the efficiency of the derivative trigger is assumed to be approximately unity for all channels. For LMO1 the achieved threshold was particularly poor, and an efficiency close to unity is considered only for energies above 1300 keV. LMO6 exhibited a better threshold, so unit

²¹Due to time constraints dictated by the cryogenic system, the background measurement was acquired before completing the WP scan described above, resulting in a sub-optimal WP choice.




Figure 4.24: Selection efficiency of the pile-up cuts. The dashed line shows the weighted average, with the shaded band indicating its statistical uncertainty. The uncertainties are dominated by the limited statistics of the line fits. Error bands extending into non-physical regions ($\epsilon > 1$) are a consequence of low statistics and carry negligible weight in the weighted average; they are therefore still shown in the plot. The centroids of the α peaks (above 5 MeV) are not corrected for QF in this figure.

efficiency is assumed above ~ 150 keV. For the remaining LMOs the trigger efficiency is assumed to be unity above ~ 300 keV.

Two sets of data quality selections are applied to produce the background spectra.

Pile-up The first selections are aimed at rejecting pile-up events. Only windows with a single signal trigger are kept, and a pulse-recognition algorithm is used to discard cases where two pulses occur too close in time for the trigger to fire twice. This algorithm still leaves some residual pile-up when two pulses fall within a time window of the same order as the RT.

These cuts are assumed to have a flat efficiency in energy. This assumption was tested by evaluating the selection efficiency on γ lines with sufficient statistics for a meaningful fit. Each line was fitted with a Gaussian plus a linear background, and the efficiency was derived from the ratio of the peak areas before and after the cuts. The ^{210}Po Q -value and α lines were also fitted with a Gaussian model which, given the limited statistics, showed good agreement with the data and were therefore included in the efficiency evaluation. The resulting efficiencies for each channel are shown in Fig. 4.24. In all cases, uncertainties are dominated by the statistical errors on the fit parameters. The results are consistent with a flat energy dependence, so the efficiency is taken as the weighted average across all lines.

To maximize statistics, no further data quality selections (such as those rejecting events on drifting baselines) are applied. This choice may slightly degrade the energy resolution, but the compromise is acceptable given the goals of this study and the relatively short time dedicated to acquire background data.

Cross-talk The second set of selections target cross-talk events. In pre-CCVR2 (v2), a wiring issue introduced significant cross-talk. The differential amplifiers used in the CUPID FE boards pro-

vide a high common-mode rejection ratio (CMRR), which suppresses cross-talk when it is induced symmetrically on both polarities of the NTD readout circuit. The wiring fault, however, coupled the cross-talk only to one polarity of neighboring channels, preventing the CMRR from cancelling it. Both heat and light channels were affected, with cross-talk induced in both directions.

Cross-talk pulses can be distinguished from real interactions by their shapes. An example is shown in the top plot in Fig. 4.25, where a real event in LMO6 generates a spurious pulse in LD2 characterized by an anomalous shape. The wOF algorithm provides several pulse-shape parameters that can be exploited for rejection, but many of them (e.g. wOF_Chi2) are energy-dependent, and cross-talk shapes differ across channels. No single set of cuts is therefore valid for all detectors and the following customized channel-by-channel strategy is implemented in the DIANA analysis pipeline:

1. A given pulse-shape parameter is selected and plotted against energy. Cross-talk events are seen as distinct outlier populations (Fig. 4.25, left).
2. The energy dependence of the parameter is modeled with a function $\mu(E)$, which can be piecewise defined and include both linear and nonlinear components. Its energy-dependent width, $\sigma(E)$, is also modeled with an arbitrarily complex piecewise defined function.
3. The parameter is normalized as:

$$\text{Normalized parameter} = \frac{\text{Original parameter} - \mu(E)}{\sigma(E)} \quad (4.4)$$

so that the resulting distribution for real signal events approximates a standard Gaussian $N(0,1)$.

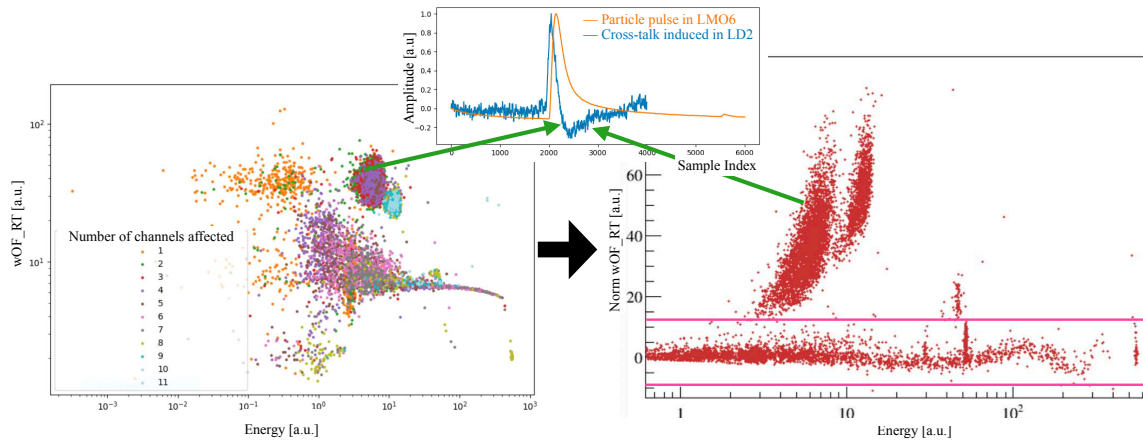


Figure 4.25: *Left*: wOF_RiseTime vs. energy in LD2. Cross-talk events form distinct outlier populations. Colors indicate the number of coincident cross-talk events generated by the same aggressor, in this case a real interaction in LMO6. *Top center*: an example of a cross-talk pulse induced by in LD2 by LMO6. *Right*: the same parameter normalized according to Eq. (4.4), where a flat cut (pink lines) cleanly separates signal from cross-talk, with efficiency close to 1.

4. The normalized parameter is plotted against energy (Fig. 4.25, right). In this representation, flat cuts can be applied with efficiency assumed to be energy-independent.

The separation power achieved with this strategy depends on the channel and on the parameter selected. In some cases, multiple parameters must be combined in order to reject all cross-talk events while maintaining unitary signal efficiency.

Each channel was treated individually, and it was always possible to define suitable cuts after normalization, all with signal efficiencies close to unity. This was verified with the same strategy used for pile-up cuts, confirming flat efficiencies consistent with unity within the available uncertainties.

Quenching factor

The thermal quenching factor between β/γ and α interactions in each channel was computed following the same procedure used in pre-CCVR1 (Sec. 4.4.9). The resulting values are listed in Table 4.9 under the WP4 column.

The only difference with respect to pre-CCVR1 is that, in this case, the QF was determined using the ^{210}Po Q -value line instead of a pure α line. Since the Q -value peak originates from a coincident α and nuclear recoil (NR), the extracted QF can differ because NRs interact with the crystal differently due to their much higher stopping power. This, combined with the different WP selection in pre-CCVR2, explains why the QF measured in pre-CCVR1 for the two reference NIIC crystals (Table 4.5) is not identical to that observed here.

Due to the extrapolations required in the calculation, a reliable estimate of the statistical uncertainty on QF is difficult to obtain. For this reason, even though the SICCAS crystals exhibit on average a higher QF than the reference NIIC crystals, the difference may still be compatible with statistical fluctuations, and no firm conclusion can be drawn. In any case, the values reported in Table 4.9 are consistently used for correcting α and Q -value energies in the background spectra presented in the following.

Gamma region

The spectra in the γ region for all LMOs, normalized by selection efficiency, are shown in Fig. 4.26. The upper panel compares the three SICCAS crystals, while the two lower rows compare each SICCAS crystal with the facing NIIC crystal on the same floor of the tower.

LMO6 shows a strong ^{40}K contamination. The high intensity of the β^- continuum indicates that this contamination is predominantly internal. LMO5 also exhibits a pronounced 1460.8 keV γ line from the ^{40}K EC branch, but the smaller ratio of the line intensity to the internal β^- continuum suggests that an external component may be present (likely due to the highly contaminated LMO6 which is separated by LMO5 only by the thin Ge LD).

The (enriched) NIIC crystals display an enhanced continuum due to the ^{100}Mo $2\nu\beta\beta$ spectrum, which provides a broad background in the β/γ region. This contribution acts as a background under the 1460.8 keV γ line from ^{40}K . No such line is visible in the NIIC spectra collected during pre-CCVR1 (Fig. 4.11), while here it can be seen in the same crystals. The 1460.8 keV peak observed in LMO2 in pre-CCVR2 is therefore likely due to an external contribution, plausibly from the facing SICCAS crystal (LMO5). A direct comparison with pre-CCVR1, however, is complicated by the elevated background levels present in that run because the Rn-box flushing was absent.

Additional γ lines from the ^{238}U chain are identifiable in some channels (mainly 609.3 keV from ^{214}Bi and 351.9 keV from ^{214}Pb). Given the short half-lives of the isotopes in the mid-part of the

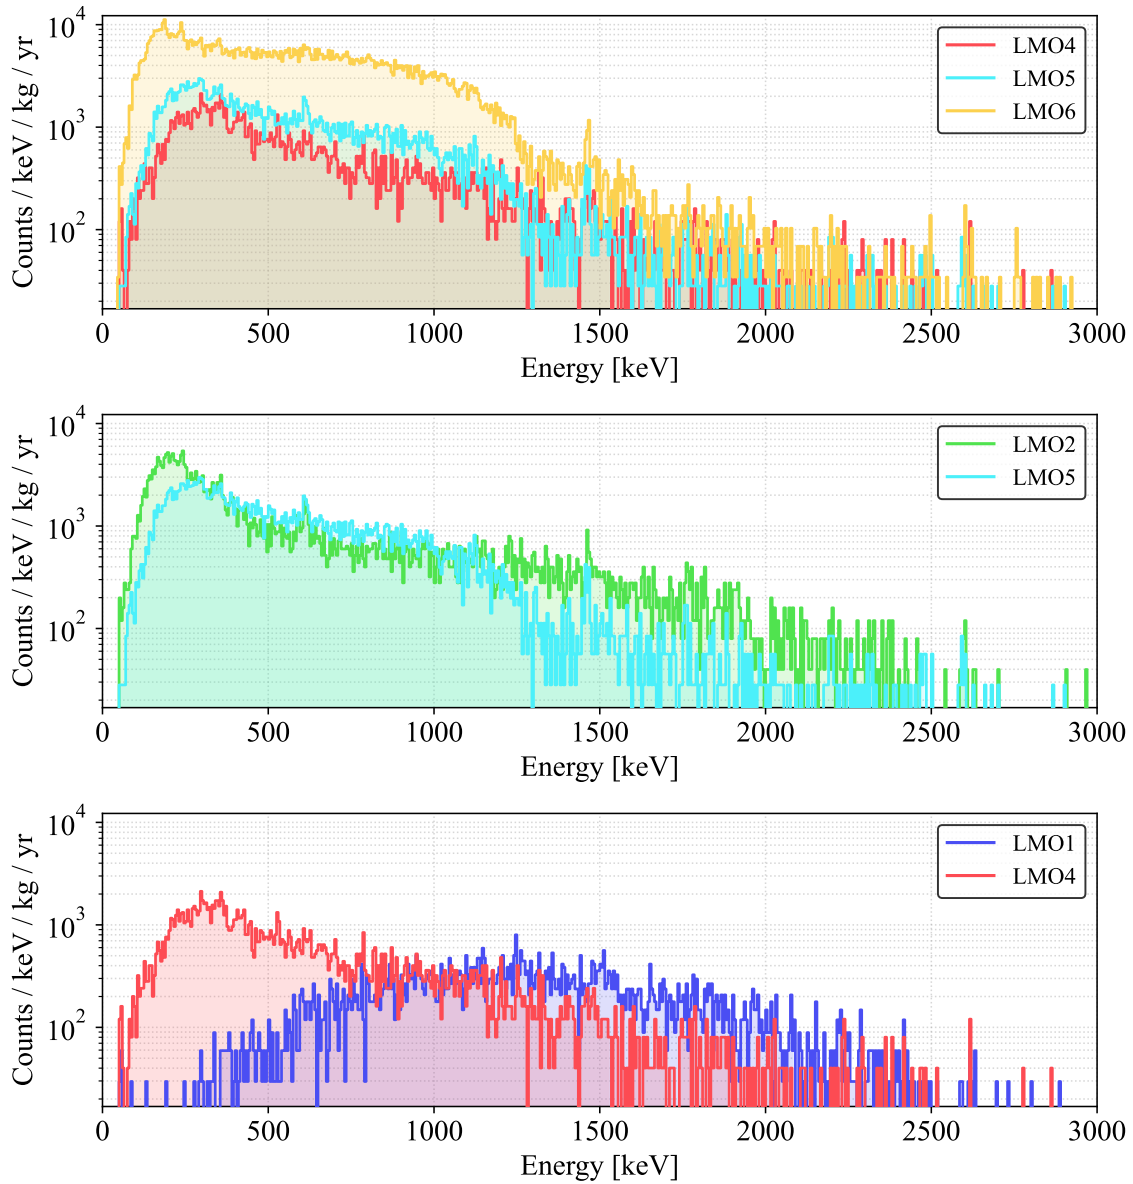


Figure 4.26: Background spectra in the γ region normalized by selection efficiency. *Top*: comparison of the three SICCRAS crystals. *Middle and bottom*: comparison of each SICCRAS crystal with the facing NIIC crystal on the same tower floor. A description of the main features of the spectra (^{40}K peak and continuum and lower energy γ lines from ^{238}U chain) is given in the text.

chain, a crystal origin (bulk or surface) would imply accompanying α lines above 6 MeV. No such α peaks are observed. It is therefore concluded that these γ lines arise from external backgrounds rather than from crystal impurities. Similar considerations apply for the few counts observed at the ^{208}Tl 2615 keV line. Since these contaminants are not crystal-related, they are not discussed further. However, in future validation runs, when radiopurity will be a primary objective, external contributions will need to be modeled more accurately to better constrain the crystals' radiopurity.

^{40}K contamination To quantify the ^{40}K contamination in the SICCAS crystals and to distinguish internal from external contributions, a dedicated MC simulation was performed using the official CUPID software based on Geant4. It includes the detector response and accounts for the finite energy resolution and the threshold. The model was adapted from that developed for the GDPT demonstrator, which features a comparable geometry, and modified to reproduce the smaller tower used in pre-CCVR2. Only the detectors and the tower assembly were included, while the cryostat shielding and other external components were excluded. As a consequence, backgrounds originating outside the cryostat or from cryostat materials were not simulated.

The simulations support the hypothesis that LMO6 contains an internal ^{40}K contamination. As shown in the left panel of Fig. 4.27, the simulated internal β^- continuum up to $E_0 \simeq 1311$ keV is not exceeding the data provided the entire 1460.8 keV γ line is attributed to internal activity from the EC branch. The ^{40}K activity is evaluated using the most conservative approach, that is, assuming that *all* counts under this peak originate from internal contamination, thereby providing an upper-limit estimate; the true activity may be lower if part of the signal arises from neighboring crystals (e.g., LMO3, whose spectrum could not be measured).




Figure 4.27: Comparison of data with ^{40}K simulations. *Left:* LMO6, where the simulated internal β^- continuum is compatible with the data if the entire 1460.8 keV γ line is assumed internal. *Right:* LMO5, where attributing the full 1460.8 keV γ line to internal activity would overestimate the internal β^- continuum.

In contrast, for LMO5 and LMO4, attributing the full 1460.8 keV γ line to a bulk contamination would yield a simulated internal β^- continuum that exceeds the measured spectrum, as illustrated for LMO5 in the right panel of Fig. 4.27. For these crystals, the activity is evaluated using the most conservative approach, namely assuming that all counts in the 1000–1200 keV interval originate

from the internal ^{40}K β^- continuum. This neglects possible contributions from other backgrounds in this range (for example the Compton continuum of external 1460.8 keV γ rays) and therefore overestimates the true internal activity. The resulting activities are reported in Table 4.12.

Crystal	Conservative hypothesis for ^{40}K	Activity of internal ^{40}K contamination [mBq/kg]
LMO4	Limited by β^- continuum	1.2 ± 0.4
LMO5	Limited by β^- continuum	3.9 ± 1.0
LMO6	All 1460.8 keV γ attributed to internal	10.7 ± 2.5

Table 4.12: Activity of internal ^{40}K contamination in the SICCAS crystals from CCVR2 under the most conservative assumptions. For LMO6, the full 1460.8 keV γ line is assumed internal. For LMO4 and LMO5, the values are obtained under the conservative hypothesis that all counts in [1000–1200] keV originate from the internal β^- continuum, neglecting other β/γ backgrounds. Uncertainties are statistical only.

These results indicate that LMO6 is more contaminated by ^{40}K than LMO4 and LMO5. Although such a trend could be compatible with an effect of recycling steps, an unambiguous attribution is not possible, since LMO6 was produced starting from two different precursor powders. It is therefore concluded that both the recycling protocol and the choice of precursor powders must be tightly controlled to mitigate the risk of ^{40}K accumulation. This observation has been communicated to SICCAS as input to refine the growth process.

LMO4 and LMO5, although not grown from recycled material, exhibit mild contamination. These levels must be interpreted in light of the intrinsic radiopurity of the precursor powders. The powders used for this test growth were not those intended for enriched production. For example, the ICP–MS results in Table 4.8, combined with the specific activity of natural K, imply that the natural MoO_3 powder from producer “O” (“A”) corresponds to a ^{40}K contamination of ~ 3000 mBq/kg (~ 800 mBq/kg), which exceeds the CUPID requirement on MoO_3 in Table 3.4. Despite this, the re-crystallization history of LMO4 and LMO5 achieves substantial purification, by at least a factor > 100 . Therefore, the crystals grown from enriched $^{100}\text{MoO}_3$ powder meeting the requirements specified in Table 3.4 are expected to meet the CUPID specifications for the final $\text{Li}_2^{100}\text{MoO}_4$ (Table 3.2), provided no additional impurities are introduced during growth.

Alpha region

The spectra in the α region for all LMOs, normalized by selection efficiency, are shown in Fig. 4.29 over the interval 3–6 MeV. The upper left (right) panel compares the three SICCAS (NIIC) crystals, while the bottom panels compare each SICCAS crystal with the facing NIIC crystal on the same floor of the tower. Above 6 MeV, no α lines are observed in any crystal.

^{210}Po contamination Figure 4.29 shows that all SICCAS crystals exhibit a ^{210}Po contamination, visible as a Q -value line at $Q_\alpha \simeq 5407$ keV which is more intense than the corresponding α peak at $E_\alpha \simeq 5304$ keV. In contrast, the NIIC crystals show the opposite behavior, consistent with pre-CCVR1 observations (Fig. 4.13).

Due to the relatively high trigger thresholds in the heat channels, coincidence studies provide limited information: nuclear recoils from low-energy surface events are not efficiently triggered. It

is therefore not possible to quantitatively distinguish bulk from surface decays, nor to estimate the penetration depth of surface contaminations using coincidences alone. For this reason, the activities of the ^{210}Po contaminant are quoted under the following worst-case scenario assumptions:

- **Bulk contamination:** all counts observed at the Q -value are interpreted as internal.
- **Surface contamination:** MC spectra are simulated assuming contaminants uniformly distributed on the crystal surfaces, with exponential depth profiles of mean penetration $1\ \mu\text{m}$, $0.1\ \mu\text{m}$, and $0.001\ \mu\text{m}$. Each simulation is normalized to the observed data under the conservative hypothesis that the entire Q -value line is explained by surface events.

All the above hypotheses, including the different surface penetration depths, are mutually exclusive and are treated independently. This approach provides the most conservative evaluation of the possible ^{210}Po activities under distinct assumptions about the origin of the contamination.

An example of data-simulation comparison for LMO5 is shown in Fig. 4.28. The activities computed under these conservative assumptions are reported in Table 4.13.




Figure 4.28: Comparison of data with ^{210}Po simulations for LMO5. *Left:* simulation normalized to the integral of the Q -value line, under the conservative assumption that all events are due to bulk contamination. *Right:* simulations for different surface depth profiles, each normalized to the Q -value line, corresponding to the conservative hypothesis that all counts at the peak arise from surface contamination.

Other α lines Figure 4.29 shows that LMO6 exhibits two intense α structures near $\sim 4.8\ \text{MeV}$ and $\sim 4.3\ \text{MeV}$. The pronounced low-energy tails and the absence of a corresponding Q -value peak strongly suggest a surface origin rather than bulk crystal contamination. Although coincidences among LMO heat channels are limited by thresholds, the LDs operate at much lower thresholds and always trigger on NR and on (saturated) α pulses. No events from these two structures in LMO6 are in coincidence with LD8 or LD7, indicating cleanliness of top and bottom crystal faces, thereby

Crystal	Bulk [mBq/kg]	Surface (1 μm) [$\mu\text{Bq}/\text{cm}^2$]	Surface (0.1 μm) [$\mu\text{Bq}/\text{cm}^2$]	Surface (0.01 μm) [$\mu\text{Bq}/\text{cm}^2$]
LMO4	1.8 ± 0.42	8.1 ± 1.2	8.7 ± 1.9	9.8 ± 2.1
LMO5	1.4 ± 0.3	6.5 ± 1.3	6.9 ± 1.4	7.8 ± 1.6
LMO6	2.4 ± 0.4	10.9 ± 1.7	11.6 ± 1.8	13.2 ± 2.1

Table 4.13: Activity on ^{210}Po contamination in SICCAS crystals, assuming either bulk contamination or surface contaminations with different depth profiles. Each value is obtained under a separate conservative assumption about the origin of the contaminant, with the different cases treated independently.




Figure 4.29: Spectra in the α region (3–6 MeV), normalized by selection efficiency. *Top left*: comparison of the three SICCAS crystals. *Top right*: comparison of the NIIC crystals. *Bottom panels*: each SICCAS crystal compared with its facing NIIC counterpart on the same floor. No α lines are observed above 6 MeV, where only a flat continuum is present.

suggesting an external origin. A localized contamination on copper facing LMO6 could plausibly account for this behavior, including the smaller peaks observed at the same energies in LMO5 and at even lower intensity in LMO4, which would be consistent with the reduced line of sight expected

from a local spot.

Since this contamination is external, it is not related to crystal quality and will not be further discussed here²². This, however, prevents setting robust limits on bulk ^{238}U and ^{232}Th progenitors, since their α energies overlap with backgrounds from degraded α from this contaminant. For the same reason, it is not possible to draw firm conclusions on potential differences related to the two surface treatment methodologies used for these crystals; however, the absence of α lines above 6 MeV in all three samples can be interpreted as an indication that the two polishing methods result in comparable radiopurity. Given the qualitative nature of these results, only feedback on ^{210}Po is provided to SICCAS in this case. The same crystals are being currently operated in a new CUPID demonstrator (Vertical Slice Test Tower), and updated feedback on radiopurity in this region will be provided once new data are available.

4.6 Conclusion

During pre-production CUPID must simultaneously provide feedback to the producer and establish a rigorous protocol for crystal validation. While immediate measurements guide the optimization of growth procedures, the same data also define the standardized CUPID Crystal Validation Runs (CCVRs) that will be applied during full-scale production to monitor crystal quality. To this end, pre-CCVR campaigns in the underground low-background and low-noise LNGS Hall C CUPID cryogenic facility were conducted. The measurements presented in this chapter demonstrate that the calorimetric performance of the first batch of natural pre-production SICCAS crystals already fulfills CUPID expectations and establish the methodological basis for assessing calorimetric response and light yield. The observed light yields are consistent with those previously obtained in crystals from the same producer used in successful CUPID demonstrators.

The SICCAS pre-production natural crystals, cut to CUPID dimensions and produced via BG growth along different re-crystallization paths, were intentionally realized at this early stage without a finalized radiopurity protocol and using powders not yet meeting the final CUPID requirements. Two crystals were grown solely from natural powders, while others incorporated recycled fractions—a necessary step for evaluating and maximizing recovery efficiency before the transition to enriched isotopes. Calorimetric performance and light yield proved stable across all growth strategies. Radiopurity studies, although performed under conservative assumptions, highlighted that the re-crystallizations employed provide a good purification factor. An enhanced ^{40}K content was found in crystals containing recycled material. This provided essential feedback to SICCAS, who can now refine growth parameters to mitigate such effects.

Taken together, these results provide both direct input for improving crystal growth and the basis for a standardized CCVR protocol, ensuring reproducible validation of performance and radiopurity and establishing the framework required for reliable full-scale production of CUPID crystals.

²²The two structures near 4.8 and 4.3 MeV, together with the absence of features above 6 MeV, could be ascribed to the ^{238}U chain with a break at ^{226}Ra . In this scenario, however, the QF-corrected energies do not match the nominal α lines of ^{238}U (4.19 MeV) and ^{234}U (4.77 MeV) without fine tuning of the contaminant depth and position or invoking a strong detector nonlinearity. In any case, the pronounced left-tail shape makes it difficult to assign precise nominal energies to the degraded α particles producing these structures. Since the contamination does not appear to originate from the crystal surfaces, no further investigation is presented here.

Chapter 5

Systematic Scintillation and Thermoluminescence Characterization of Li_2MoO_4 Crystals

The scintillation of Li_2MoO_4 crystals is central to background mitigation in next-generation $0\nu\beta\beta$ searches since it enables efficient α and pile-up rejection. Given the moderate intrinsic scintillation and low light yield of Li_2MoO_4 , systematic studies that quantify the light production and its correlation with growth-induced defects or impurities are needed. The following measurements examine the role of defects in scintillation and provide guidelines to SICCAS for selecting growth parameters that optimize the scintillation performance.

5.1 Radioluminescence and Thermoluminescence Characterization of Scintillating Li_2MoO_4 Crystals

The performance of scintillating crystals is strongly influenced by crystal quality. Uncontrolled impurities and intrinsic defects create color centers and trap states within the bandgap that modify optical and scintillation properties, in particular the scintillation yield [158]. Although growth technologies continue to improve, a finite density of point defects at the atomic scale is unavoidable. During pre-production, SICCAS is still defining a standardized growth protocol that must optimize multiple aspects at the same time: radiopurity, crystal quality, and calorimetric and scintillation performance. If correlations between crystal defects and scintillation properties were established, they would provide supplementary input that can guide SICCAS in finalizing the growth protocol. The study of the role of defects is therefore a complementary strategy to calorimetric measurements at ~ 15 mK for assessing crystal quality and provides practical guidance for the producer [194, 195]. This is especially relevant in pre-production, where the growth process varies from crystal to crystal (see Chapter 3).

Refs. [158, 160, 196, 197, 198] show that the scintillation properties and the effects of defects in

Li_2MoO_4 can be effectively investigated by testing small samples over a wide temperature range, from 320 K to 10 K, combining radioluminescence RL, both as a function of temperature and of the cumulated dose (*bright burn*), and wavelength-resolved thermally stimulated luminescence (TSL), also known as thermoluminescence. In the early stage of pre-production, seven crystal samples are cut from BG ingots produced at SICCAS to perform these studies. The results, together with a summary of the experimental methodology, are presented below.

5.2 Crystal Samples

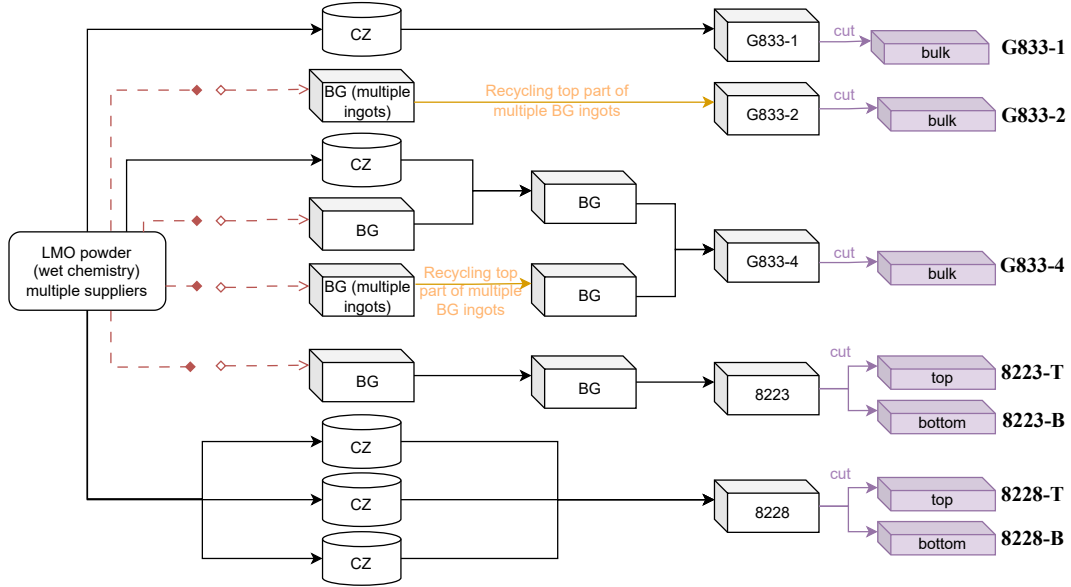


Figure 5.1: Schematic of the tested Li_2MoO_4 samples and their growth histories. Cylinders (parallelepipeds) denote CZ (BG) growths. Red dashed arrows indicate material derived from a number of multiple re-crystallizations; intermediate steps are not detailed, since only the last 2–3 iterations typically determine the final quality. Orange arrows mark recycling of the top part of the BG ingot, where impurities segregate, into subsequent growths (see Sec. 3.3.2 for details about recycling strategies). Violet arrows show the cuts used to obtain the studied samples. For some ingots a single bulk sample was prepared; for others, two samples were cut from the top and bottom to probe possible longitudinal non-uniformities. The label given to each sample is shown on the right.

A total of seven samples were available for testing, all extracted from BG ingots. Each sample originates from a distinct crystallization sequence, shown in Fig. 5.1. In addition to number and types of re-crystallizations, SICCAS varied the growth parameters, including furnace temperature and time profiles, crucible geometry, and annealing program (see Sec. 3.3 for details on crystal growth). The specific settings used for each ingot are not relevant for CUPID, since the optimization of growth parameters is entirely delegated to SICCAS. It is nevertheless important to note that each

crystal has a slightly different crystallization history, which can account for differences observed in the results of the scintillation characterization presented below. The detailed experimental results have been provided to SICCAS for interpretation in light of the parameters used in each growth.

Most precursor material comes Li_2MoO_4 synthesized from Li_2CO_3 and natural MoO_3 powders (wet chemistry) supplied by different producers. Most crystals used powders from producer "O", whose chemical purity is shown in the blue boxes of Fig. 4.14; only a tiny fraction of raw powders comes from other producers, all having similar chemical purity.

Samples G833-1/2/4 were cut from the bulk of BG ingots to a nominal size of $7 \times 7 \times 1 \text{ mm}^3$, then polished to optical quality with parallel faces. Ingots 8223 and 8228 were of sufficient quality to yield a standard CUPID crystal of $45 \times 45 \times 45 \text{ mm}^3$. The top and bottom samples (8223-T/B and 8228-T/B) were taken as (irregular) flat slices immediately above and below the CUPID-sized crystal, i.e. separated by about 45 mm along the growth axis. These slices had irregular shape and exceeded the dimensions suitable for the experimental setup described below; they were therefore cut to approximately $13 \times 5 \times 2 \text{ mm}^3$. Their faces are unpolished and not perfectly parallel, a condition that may slightly influence the results presented below. The two CUPID-sized crystals (8223 and 8228) were tested in an above-ground cryostat at IJCLab¹ to measure the light yield on a facing LD in a CUPID-like geometry (see, e.g., Fig. 2.12b). The measurement strategy followed the methodology presented in Sec. 4.5.8, although a different analysis pipeline was used. The resulting LY values are reported in the following list, together with details about each sample.

- **G833-1:** Bulk cut from a BG ingot obtained after two re-crystallizations, CZ followed by BG. The resulting ingot appears transparent.
- **G833-2:** Bulk cut from a BG ingot grown entirely from recycled material taken from the top parts of multiple previous BG ingots. The head of each ingot was dissolved in distilled water, the solutions were combined, paper-filtered, and evaporated to saturation to recover the powder used for the BG growth. The BG ingot shows a faint yellow tint.
- **G833-4:** Bulk cut from a BG ingot grown from a mixture of two components, both coming from the bulk of BG ingots. The first ingot comes from combining the bulk of a BG ingot and material from a previous CZ growth. The second ingot is obtained by recycling the heads of multiple BG ingots (as for G833-2). The resulting ingot is fully transparent, suggesting that combining recycled material with bulk material from higher-purity ingots can improve optical quality. A comparison of the RL and TSL results for G833-2 and G833-4 is therefore very informative for SICCAS.
- **8223-T and 8223-B:** Top and bottom slices from a BG ingot that also yielded a standard CUPID crystal. The BG was obtained from the bulk material of at least three consecutive BG re-crystallizations. The two slices bracket the CUPID crystal along the growth axis and are used to probe longitudinal non-uniformities. The LY of the CUPID crystal measured at IJCLab is $\sim 0.23 \text{ keV/MeV}^2$.
- **8228-T and 8228-B:** Top and bottom slices from a BG ingot that also yielded a standard CUPID crystal. The components used for the BG growths originate from three CZ ingots. Similarly to 8223-T/B, these samples are obtained from top and bottom of the CUPID crystal 8228. The LY of the CUPID crystal measured at IJCLab is $\sim 0.15 \text{ keV/MeV}^2$, so this crystal emits less scintillation light than all the other crystals produced at SICCAS.

¹Laboratoire de physique des deux infinis Irène Joliot-Curie

The results presented below can be interpreted with reference to the crystallization history summarized here, but the complete set of growth parameters is unknown to CUPID, leaving a margin of uncertainty. The aim of these measurements is to deliver a complete and accurate characterization of all samples under reproducible conditions. The detailed interpretation and correlation with growth parameters is entirely delegated to SICCAS, which holds the complete details of the production process. The same measurement protocol will be used in the future on other samples extracted from the pre-production line.

5.3 Experimental Setup

RL and TSL measurements are performed in the setup shown in Fig. 5.2. The samples are excited by unfiltered X rays from a Philips PW2274 tube with tungsten target and beryllium window, operated at 20 kV and 20 mA and placed about 13 cm from the crystal. At this voltage the spectrum is dominated by Bremsstrahlung with superimposed tungsten L and M characteristic lines, produced by electrons emitted by thermionic effect and accelerated onto the target. The crystal is mounted in an optical cryostat with a base temperature of 10 K. The cryostat is equipped with a heater driven by a PID (proportional-integral-derivative) controller, which allows setting temperatures between 10 K and 320 K. The scintillation light emerging from the cryostat window is focused by two parabolic mirrors onto the entrance slit of a monochromator (Jobin-Yvon Triax 180) equipped with a 100 lines/mm grating and coupled to a UV-enhanced, back illuminated, Horiba Symphony II CCD [199] operated with liquid-nitrogen cooling, enabling high-resolution spectroscopy. The spectral response and wavelength scale of the spectrometer are custom calibrated using two Oriel spectral calibration lamps; a spectral correction curve is available and applied to the collected data.

Direct measurement of the crystal temperature is not feasible with this apparatus since it is

²These values are determined selecting events releasing [2.5, 3.5] MeV in the heat channel, a region dominated by β/γ interactions.




Figure 5.2: Setup used to acquire RL and TSL measurements. In the left panel, a rectangular sample is shown on the cryostat cold finger and fixed with silver paint.

not equipped with a thermometer. Only the cryostat's cold finger temperature is monitored with a calibrated thermometer. Thermal contact between the crystal and the cold finger is ensured with silver paint. One face of each sample is coated with liquid silver paint and pressed firmly onto the cold finger; after drying the paint acts as an adhesive, providing high thermal conductance for the Li_2MoO_4 samples. Under these conditions, any temperature offset between the crystal and the cold finger is expected to be negligible. In the following, the crystal temperature is assumed to coincide with the cold-finger temperature measured by the thermometer. During data taking, temperature instability was observed below 20 K. This instability was attributed to component aging in the cryostat. The affected data were rejected from the analysis and are not included in the following discussion.

5.4 Radioluminescence and Low Temperature Scintillation Properties

5.4.1 Methodology

Radioluminescence (RL) spectra are acquired between ~ 15 K and 320 K. The crystal is first warmed to 320 K. At each selected temperature $T < 320$ K, the sample is irradiated with X-rays and the scintillation spectrum is recorded; irradiation is then stopped, a lower temperature is set, and the measurement is repeated. The temperature is decreased monotonically from high to low values to avoid overlap with thermally stimulated luminescence from previously trapped charge carriers, ensuring that the recorded emission is pure RL at each T . An example of the RL emission as a function of the temperature is shown in Fig. 5.3 for sample 8223-B.

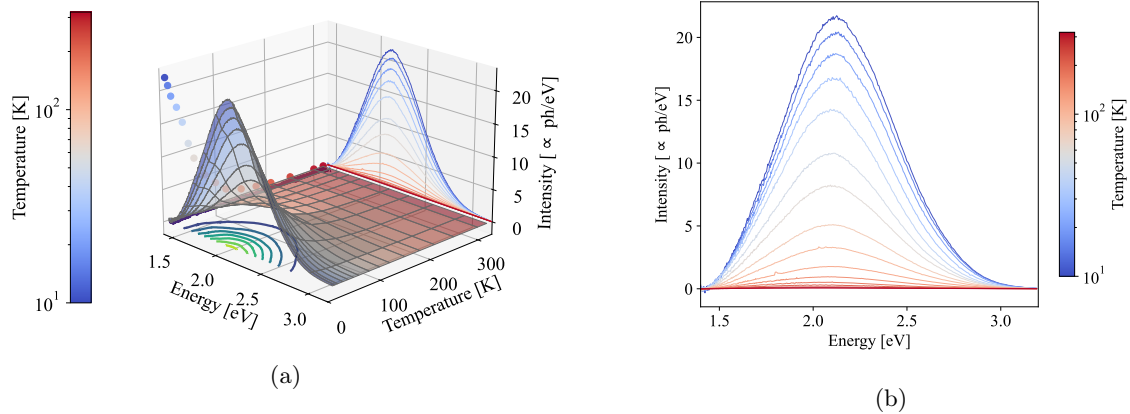


Figure 5.3: Radioluminescence as a function of temperature for sample 8223-B. In (a) the surface is colour-mapped by temperature (log scale). The vertical walls plot the emission spectrum at each temperature (lines) and the corresponding maximum intensity (dots); floor contours are colour-mapped by intensity. (b) shows the the RL emission spectra at different temperatures.

5.4.2 Results

RL measurements were acquired for all seven samples. The emission spectra at each temperature display bands with features similar to those reported for crystals grown with other techniques and by different producers [158, 160]. In particular, a single broad band is observed (see Figure 5.3b), with peak energy and width evolving with temperature. No additional band in the blue region around $T \sim 85$ K was detected in any of the samples, in agreement with Ref. [160] but in contrast with Ref. [195].

For each sample, all spectra acquired at different temperatures were processed to extract the peak emission energy, the FWHM, and the integrated intensity. The results are shown in Fig. 5.4. The right panels demonstrate that both peak energy and FWHM evolve with temperature. The FWHM narrows on cooling, from ~ 0.90 eV at 320 K to ~ 0.70 eV at 15 K. Upon cooling, the peak emission shifts from ~ 2.05 eV at $T = 320$ K to ~ 2.13 eV at $T = 10$ K, with only minor differences among samples despite their distinct growth histories (only sample G833-1 shows a nearly constant peak at ~ 2.1 eV). The shift of the peak emission energy suggests that the emission is composed of multiple overlapping bands, consistent with previous observations in Li_2MoO_4 crystals grown with different techniques [158, 196]. In particular, Ref. [160] attributes the high-energy band at 2.15 eV to radiative annihilation of excitons self-trapped on MoO_4 groups (intrinsic scintillation), and the lower-energy band at 1.9 eV to defect-related luminescence. This interpretation of the higher energy band is consistent with the scintillation mechanism of molybdates, where intrinsic emission is generally ascribed to excitons self-trapped on isolated MoO_4 complexes [200, 201]. The composite nature of the emission is also supported by the non-monotonic trend of the FWHM. When the temperature increases from ~ 150 K to ~ 200 K, the FWHM decreases. This behaviour is consistent with temperature quenching of one overlapping band, which causes an apparent narrowing of the composite broad-band emission.




Figure 5.4: RL intensity (left), peak emission photon energy (top right), and FWHM of the RL distributions (bottom right) as a function of temperature for all samples.

Isolating the defect-related component, however, requires Gaussian decomposition of the overlapping bands as it is done in Ref. [158], which is non-trivial given how much those overlap, requiring the development of a simultaneous fit. The exact origin of defect-related emission in these samples therefore remains an open point.

The integrated RL intensity (left panel of Fig. 5.4) increases with decreasing temperature until a plateau is reached at $T \sim 45$ K. Interestingly, the FWHM exhibits a similar plateau under the same temperature. This behaviour is consistent with the literature for Li_2MoO_4 and is favourable for its use as a scintillating cryogenic calorimeter.

Direct comparison of intensities between samples is non-trivial, as realistic uncertainties on the absolute values are difficult to assign for this experimental setup. Nevertheless, the measurements are systematically repeated with a reproducible methodology for all samples. The geometry of the system (focal point of the parabolic mirror and irradiation from the X-ray tube) is designed to minimize systematic differences due to sample size, allowing to get self-consistent results. Within these conditions, all samples show comparable intensity except crystal 8223-T, which consistently exhibits about one-third of the emission observed in the other samples at $T = 10$ K. Although this difference is too large to be attributed to systematic uncertainties in the measurement procedure, dedicated measurements to quantify potential systematic effects³ are planned and will be performed in the near future. Assuming that this difference is not due to systematic uncertainties in the measurement procedure, no clear correlation with the crystallization history of 8223-T is found (see Fig. 5.1).

5.4.3 Light Yield and RL properties

Finding direct correlations between RL emission and the LY in the CUPID setup is non-trivial, because the current apparatus allows measurements only down to $T = 10$ K, whereas CUPID detectors operate at $T \sim 0.015$ K. The behaviour of RL intensity in the range 10 K–15 mK is unknown for these samples, introducing a relevant source of uncertainty. Photoluminescence measurements (under VUV and UV excitation), performed down to $T \simeq 2$ K on crystals grown with other techniques, show that Li_2MoO_4 luminescence increases abruptly below $T = 10$ K [160]. This suggests that extrapolations to lower temperatures are unreliable (especially for the intensity), and it cannot be assumed that the RL intensity plateau observed below 40 K in Fig. 5.4 persists down to ~ 15 mK. Nevertheless, in absence of data at lower T , the spectral shape of the RL emission at 10 K (especially the peak emission energy and width) is typically used as a first-order indicator of the scintillation properties at ~ 15 mK temperatures. The RL peak emission energy and FWHM measurements at 10 K for all samples are listed in Table 5.1.

Several other factors further complicate RL intensity and LY comparisons. The LY measured in the CUPID geometry represents the fraction of scintillation photons collected by the light detector per unit of energy measured in the heat channel. It depends on the light collection efficiency, surface quality, and internal reflections of the setup. Moreover, β/γ interactions used to evaluate LY involve higher energies and deeper penetration than the soft X-rays used for RL measurements. A direct quantitative comparison between RL intensity and LY is therefore not experimentally accessible. Nevertheless, testing a large set of pre-production samples may reveal correlations or trends with growth parameters, motivating these studies even when all these caveats are considered.

³The 8223-T/B and 8228-T/B samples are thin, rectangular pieces measuring 13 mm \times 5 mm. Unlike the 7 mm \times 7 mm squares, their orientation on the cryostat’s cold finger (Figure 5.2) may influence the light collection efficiency.

Sample	RL Peak Emission Energy $E_{\text{peak}}[\text{eV}]$	RL width FWHM [eV]
8223-B	2.12	0.70
8223-T	2.14	0.70
8228-B	2.13	0.70
8228-T	2.15	0.71
G833-1	2.13	0.69
G833-2	2.12	0.69
G833-4	2.14	0.71
Average	2.13	0.70

Table 5.1: Radioluminescence peak emission energy and FWHM for each sample at $T = 10$ K.

In this set, the CUPID crystal from ingot 8223 shows a higher LY than the one from ingot 8228 (see Sec. 5.2), while in RL the 8223-T sample exhibits the lowest emission, with 8223-B and 8228-T/B showing similar values. No evident correlation between LY and RL intensity at 10 K is therefore observed for these ingots.

RL measurements determine the typical energy of a scintillation photon, which provides the necessary scale to estimate the width of the LY distribution and its energy dependence. The emission at $T = 10$ K is therefore taken as a reference, under the working assumption that the RL spectral shape does not change appreciably down to $T = 15$ mK (the luminescence intensity is irrelevant for the following discussion). The peak energies and FWHM values are summarized in Tab. 5.1. The average peak emission energy of all the samples is $E_{\text{peak}} \sim 2.13$ eV (582 nm, yellow). With this value, one can estimate the impact of statistical fluctuations in the number of collected photons on the width of the β/γ band in the LY distribution. This width, denoted $\sigma_{\gamma(\beta)}$, was introduced in Eq. (2.18), which defines the discrimination power (DP). From Eq. (2.18) it follows that $\sigma_{\gamma(\beta)}$ directly controls the DP. The $\sigma_{\gamma(\beta)}$ values measured for SICCAS crystals in pre-CCVR2 are reported in Tab. 4.11. $\sigma_{\gamma(\beta)}$ is energy dependent. Recalling that $\text{LY} \equiv L/H$ (light-to-heat ratio), a first-order error propagation gives

$$\text{LY} = \frac{L}{H} \Rightarrow \sigma_{\gamma(\beta)}^2(E) = (\text{LY})^2 \left[\left(\frac{\sigma_L}{L} \right)^2 + \left(\frac{\sigma_H}{H} \right)^2 \right] = (\text{LY})^2 \left[\left(\frac{\sqrt{\sigma_{\text{Pois}}^2(E) + \sigma_{\text{LD}}^2}}{E \text{LY}} \right)^2 + (\Delta E)^2 \right] \quad (5.1)$$

where E is the energy deposited in the crystal by the β/γ interaction, $L = E \text{LY}$ is the light energy reaching the LD by definition, $\Delta E \equiv \sigma_E/E$ is the relative energy resolution of the heat channel⁴, σ_{LD} is the LD baseline resolution, and $\sigma_{\text{Pois}}(E)$ accounts for Poisson fluctuations in the number of collected photons. The latter can be written as

$$\sigma_{\text{Pois}}(E) = \sqrt{N_{\text{photons}}} E_{\text{ph}} = \sqrt{\frac{E \text{LY}}{E_{\text{ph}}}} E_{\text{ph}} = \sqrt{E \text{LY}} E_{\text{ph}} \quad (5.2)$$

where E_{ph} is the average scintillation photon energy, which, in first approximation can be taken as the average value of E_{peak} from Tab. 5.1.

⁴In cryogenic calorimeters ΔE is approximately constant with energy and is $\mathcal{O}(0.1\%)$ for CUPID heat channels.




Figure 5.5: Light Yield distribution for the LMO5–LD6 pair from pre-CCVR2v2 (details in Sec. 4.5.8), overlaid with the $\pm 2\sigma$ total expected width calculated using Eq. (5.3) and assuming an average energy required to produce a scintillation photon of $E_{\text{ph}} = 2.13$ eV. The data points include all recorded interactions, comprising both α and β/γ events. The individual contributions from the LD baseline resolution, Poisson fluctuations, and heat-channel resolution are shown separately. Their quadratic sum well describes the distribution of β/γ interactions. The α population is clearly discriminated (see Tab. 4.11), and its observed width is dominated by the LD baseline resolution.

Substituting Eq. (5.2) into Eq. (5.1) and simplifying yields

$$\sigma_{\gamma(\beta)}(E) = \sqrt{\left(\frac{\sigma_{\text{LD}}}{E}\right)^2 + \left(\frac{\sigma_{\text{Pois}}(E)}{E}\right)^2 + (LY)^2 (\Delta E)^2} = \sqrt{\frac{\sigma_{\text{LD}}^2}{E^2} + \frac{LY E_{\text{ph}}}{E} + (LY)^2 (\Delta E)^2} \quad (5.3)$$

To gauge the relative sizes of the terms, one can consider LD6 coupled to LMO5 in pre-CCVR2 at WP5 (chapter 4), with $\sigma_{\text{LD}} \simeq 84.5$ eV from Tab. 4.10 and $\Delta E \simeq 0.12\%$ inferred from the peak FWHM in Tab. 4.9. Using $E_{\text{ph}} = E_{\text{peak}} = 2.13$ eV (Tab. 5.1) and $LY \simeq 0.3$ keV/MeV (Tab. 4.11), a ^{208}Tl 2615 keV γ event produces on average

$$N_{\text{photons}} \simeq \frac{E LY}{E_{\text{peak}}} \simeq \frac{2615 \text{ keV} \times 0.3 \text{ keV/MeV}}{2.13 \text{ eV}} \simeq 368 \text{ photons}$$

on one LD, which corresponds to $\sigma_{\text{Pois}}(2615 \text{ keV}) \simeq 40$ eV. This is smaller than σ_{LD} but not negligible, especially at energies near $Q_{\beta\beta}$, where the term σ_{LD} is suppressed by E^2 at the denominator. By contrast, the term $(LY)^2 (\Delta E)^2$ remains negligible over the entire dynamic range of the detector. Fig. 5.5 shows Eq. 5.3 and its individual components compared with data, demonstrating that this simple model has a good agreement.

The same reasoning extends to α events. In this case, since α interactions produce much lower scintillation $\sigma_{\alpha}(E)$ is dominated by the LD baseline term, σ_{LD} .

In summary, while σ_{Pois} contributes to the DP, the dominant limitation arises from the LD baseline term in Eq. (5.3). Reducing LD noise is therefore critical for CUPID.

5.5 Thermoluminescence

5.5.1 Methodology

Defects in the crystalline lattice influence luminescence and may limit scintillation performance. Thermally stimulated luminescence (TSL), also called thermoluminescence, probes the interplay between recombination centers and traps [202]. TSL is the light emitted upon heating a scintillator that was previously exposed to ionizing radiation at a lower temperature. After irradiation ceases, the temperature is increased linearly while monitoring the emission arising from carriers released from traps and radiatively recombining at luminescence centers.

Each sample is first warmed to 320 K to empty traps, then cooled to the cryostat base temperature (~ 10 K). In this state, traps are assumed empty and can be refilled by charge carriers. The sample is irradiated with soft X-rays for 15 min without spectral acquisition. Immediately after irradiation, the temperature is raised at a constant rate of 6 K/min, precisely controlled by the PID. During the warm up, the light emission is monitored by acquiring consecutive spectra with 5 s integration time, providing fine temperature granularity of the TSL signal.

5.5.2 Results

The data acquired with the procedure described above are displayed as TSL contour plots shown in Fig. 5.6 for all samples. All crystals, except 8228-T, exhibit a non-null TSL signal only below 60 K, where the efficiency of the luminescence centers is relatively high. The emission appears as a broad band centered at ~ 2.1 eV, consistent with the behavior reported for Li_2MoO_4 samples grown by NIIC using different techniques [158].

In contrast, sample 8228-T shows an additional, weaker, and narrower TSL band between ~ 100 K and ~ 150 K, centered at about 1.8 eV. The appearance of these higher-temperature peaks suggests the presence of more stable traps that release carriers only when the crystal gets more thermal energy. The fact that this TSL emission is narrower than the RL spectrum and the lower temperature trap implies that the released carriers mainly recombine through a single-component radiative emission. This may indicate the presence of a defect in the lattice, linked with the presence of a trap and a luminescence center, the two being geometrically correlated. As a consequence, when the trapped carrier is freed, it recombines almost immediately with the same defect that emits light, thus acting as a luminescence center. A similar feature was observed for the NIIC sample in Ref. [158], which was interpreted as the evidence of spatial correlation between that specific trap and a luminescence center.

The TSL contour plots were used to construct the *glow curves*, obtained by integrating each spectrum to derive the total TSL intensity as a function of temperature. Because the TSL signal depends both on the number of electron being released by the traps and on the efficiency of the luminescence centers, in order to directly compare the TSL emission of different samples their glow curves have to be normalized so as to factor away these effects. A conventional way to perform this normalization is by dividing the integrated TSL signal by the intensity of the RL emission at the same temperature. As a consequence, the RL intensity of each sample was interpolated numerically and the normalized TSL/RL glow curves were computed. This procedure is meaningful only where a non-zero TSL signal is observed (below 60 K in this case). Both the TSL and normalized TSL/RL glow curves for all samples are shown in Fig. 5.7.

From the glow curves, all samples display a broad low-temperature TSL peak. Sample 8228-T additionally exhibits weaker peaks at higher temperatures, consistent with the discussion above.

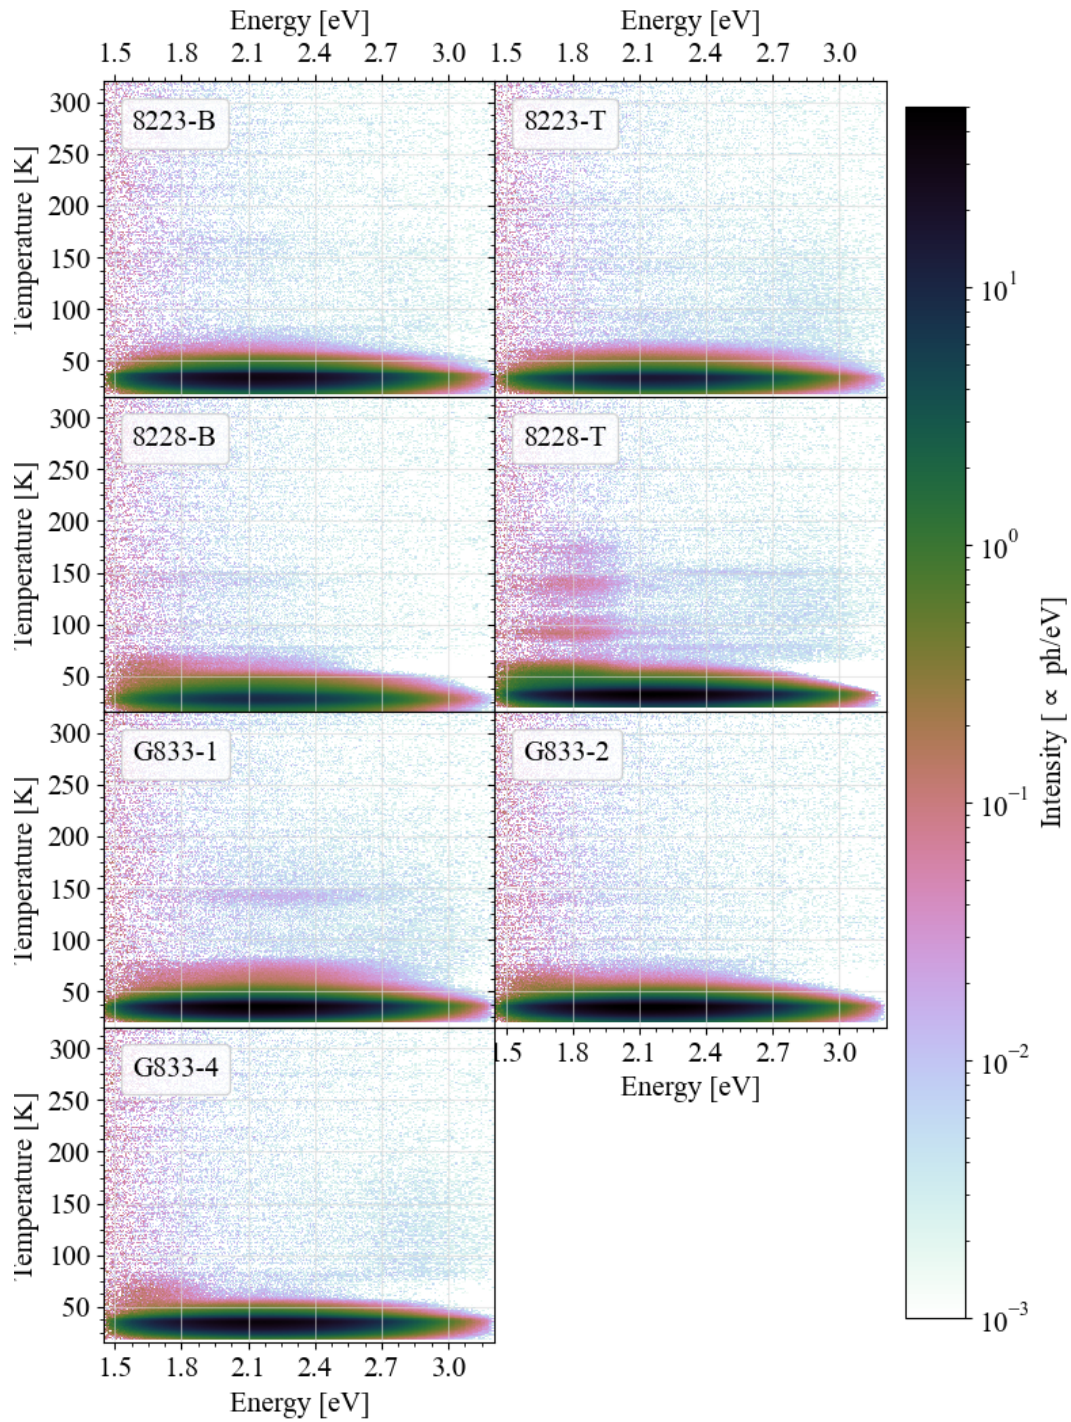


Figure 5.6: TSL contour plots for all samples.

The shape of the tails of the low-temperature peak indicate that they are composed of multiple, overlapping components with different trap concentrations [202]. Due to the strong overlap of these components, a deconvolution was not attempted, and only qualitative considerations are presented below. Accordingly, the peak positions mentioned in the discussion below should be regarded as approximate indicators for comparing different samples rather than quantitative determinations.

Low-temperature peaks in TSL glow curves have been previously observed in Li_2MoO_4 samples grown with different techniques. In the present case, only one broad low-temperature peak is visible, whereas Refs. [158, 160, 194] report multiple peaks. It should be noted that the apparent peak positions depend on the selected heating rate [158].

In the normalized plot in the right panel of Fig. 5.7, all samples except those from ingot 8228 exhibit a peak at $\sim 35\text{--}37\text{ K}$ and have similar amplitudes, indicating comparable trap concentrations. Samples G833-1/2/4, which originate from the bulk of three ingots (1/2/4) grown in parallel in the same furnace (G833) and share the same annealing program, display identical normalized glow curves. After normalization by RL, the top and bottom slices of ingot 8223 also show similar glow curves, indicating good longitudinal uniformity within this ingot.

By contrast, the samples from ingot 8228 present noticeable differences. The bottom slice (8228-B) shows a distorted peak that resolves into two components of similar amplitude at $\sim 26\text{ K}$ and $\sim 30\text{ K}$, with a near absence of traps at $\sim 35\text{--}37\text{ K}$, which are present in the other samples. The top slice (8228-T) exhibits a pronounced peak at $\sim 30\text{ K}$ and a deformed low-temperature tail, consistent with the presence of the same trap families as in 8228-B ($\sim 26\text{ K}$ and $\sim 30\text{ K}$) but at higher concentration. These observations indicate significant longitudinal non-uniformity in ingot 8228. It is noteworthy that ingot 8228 underwent a double recrystallization (CZ+BG), whereas ingot 8223 experienced at least three BG recrystallizations.

Interestingly, the CUPID crystal extracted from ingot 8228 has a lower LY measured at $\sim 15\text{ mK}$ than that from ingot 8223 (see Sec. 5.2). Although extrapolating TSL intensities above 10 K to




Figure 5.7: *Left*: TSL glow curves for all samples (as indicated in the legend), extracted from the contour plots in Fig. 5.6. *Right*: the same curves normalized to the efficiency of the luminescence centers, estimated by interpolating the RL intensity data points in Fig. 5.4.

LY at 15 mK is non trivial, the stronger non-uniformity and distinct TSL behavior of ingot 8228 suggest a possible correlation with the lower LY observed in CUPID detector modules operated at ~ 15 mK. Confirmation of such correlations requires screening a larger set of samples from multiple ingots with combined RL-TSL studies and LY measurements at ~ 15 mK.

5.6 Transmittance

The crystal requirements listed in Sec. 3.1 specify colorless crystals with an absorption length > 50 cm in the 550–700 nm spectral range. This condition is imposed to prevent self-absorption, which would reduce the LY. During the early stages of pre-production, some ingots appeared slightly yellow, suggesting that SICCAS still needs to optimize the growth parameters. In particular, ingot 8228 and the ingot used to extract sample G833-2 showed the strongest coloration, while the remaining samples were either colorless or exhibited only a weak tint.

The experimental setup used for the reflectance measurements described in Sec. 6.2.4 was adapted to measure the absorbance of the samples. In this configuration, the port previously used to host the coated crystal (see Fig. 6.2) is closed with a plug of the integrating sphere, and the sample is placed on a mask with a circular aperture positioned at the hole where the light-beam enters the sphere. The incident light passes through the sample, then through the mask hole, and enters the integrating sphere. A reference measurement is acquired without placing the sample in front of the mask hole. The absorbance A is obtained as the logarithm of the ratio between the transmitted and reference intensities, and the transmittance is then calculated as $T = 10^{-A}$.

If the crystal faces are optically polished and perfectly parallel, this configuration allows to measure of the spectrally resolved absorption coefficient. When the crystal also exhibits a flat, fully transparent window over a sufficiently wide spectral band, the refractive index can be determined as well. These quantities are relevant to CUPID, which is developing photon simulations to, among other goals, optimize the crystal surface treatment (see Section 3.1). In practice, however, surface imperfections cause scattering and refraction of the transmitted beam, resulting in a wavelength-dependent net reduction of the measured transmittance⁵. Although the integrating sphere mitigates these effects, only samples with nearly ideal optical quality permit absolute measurements. In this set, samples 8223-T/B and 8228-T/B exhibited irregular, non-parallel, and unpolished faces, while samples G833-1/2/4 had smoother and more parallel surfaces, though not of perfect optical quality due to exposure to humidity. As a result, absolute transmittance measurements could not be performed. Furthermore, given the limited sample thickness ($\ll 50$ cm), this setup is not suitable for verifying the absorption-length requirement in Section 3.1. Nonetheless, the spectral shape of the transmittance remains highly informative, providing insight into color centers and possible correlations with other properties such as trapping centers. Relative comparisons between samples are therefore meaningful and can offer valuable feedback to SICCAS.

The transmittance spectra for all seven samples are shown in Fig. 5.8. Samples from ingots 8223 and 8228, characterized by irregular surfaces, display a pronounced reduction of transmittance even in the wavelength region where Li_2MoO_4 is nominally transparent. This effect is strongly mitigated in samples G833-1/2/4, which have smoother and more parallel faces, although minor deviations remain. The residual transmittance in the transparency region is primarily due to surface reflections,

⁵Rayleigh scattering scales as λ^{-4} introducing wavelength dependent effects. Non-parallel surfaces may deviate part of the light out of the mask hole resulting in an approximately flat net reduction.




Figure 5.8: Wavelength-resolved transmittance measured for all seven samples using the setup shown in Fig. 6.2. At high wavelengths, where Li_2MoO_4 is transparent, the transmittance does not reach 100% even for the most regular samples (G833-1/2/4), in agreement with what expected from surface reflections in perfectly transparent materials. In the transparency region, the lower transmittance observed in samples 8223-T/B and 8228-T/B is an experimental effect attributed to their irregular shape and non-optical surface quality.

which depend on the refractive index⁶. The spectra can thus be interpreted relatively rather than absolutely.

All samples show an absorption band between ~ 300 nm e ~ 400 nm. Interestingly, sample 8228-T (having different TSL and RL FWHM behaviour) stands out for its distinct spectral shape. Relative to its high-wavelength reference value, it shows a marked decrease in transmittance below ~ 450 nm (~ 2.7 eV), explaining its visible yellow coloration. This absorption region does not significantly overlap with the RL emission measured at ~ 10 K (the closest RL temperature to ~ 15 mK), which peaks at 2.15 eV with a FWHM of 0.71 eV for this sample. Consequently, self-absorption is unlikely to explain the lower LY observed for the CUPID crystal extracted from ingot 8228 (see Sec. 5.2). Sample 8228-B exhibits a weaker but similar reduction, indicating longitudinal non-uniformity in this ingot.

In contrast, ingot 8223 shows a more uniform behavior, although sample 8223-T has slightly lower transmittance (relative to its high-wavelength reference value) below 400 nm compared to 8223-B, indicating minor longitudinal variation.

Samples G833-1/2/4, extracted from the bulk, cannot be used to assess longitudinal uniformity. Nevertheless, sample G833-2, which appeared more yellow, displays lower transmittance below 400 nm than G833-1/4, consistent with its observed tint. G833-1 and G833-4 show negligible differences. In all cases, the RL emission spectra at $T = 10$ K are fully contained within the

⁶In the transparency region, the plateau transmittance can be used to compute the refractive index. This procedure yields an effective $n \sim 1.6$ for samples G833-1/2/4, slightly higher than literature values for Li_2MoO_4 (birifringent, $n_o = 1.570$, $n_e = 1.433$) [203]. The discrepancy indicates residual surface nonidealities, while confirming that losses are reflection dominated and the samples are transparent above 500 nm.

transparency region, confirming that a slight yellow tint does not significantly contribute to self-absorption losses.

5.7 Summary

Overall, all samples exhibit a broad RL band whose intensity increases upon cooling. The peak emission shifts from about 2.05 eV at 320 K to 2.13 eV at 10 K, accompanied by a narrowing of the FWHM, evidences of a composite emission. TSL reveals a common low-temperature peak near 35–37 K for all samples except those from ingot 8228. Transmittance shows that slight yellow tints visible in some samples originate from absorption below about 450 nm and do not overlap the RL band at 10 K, so self-absorption is not expected to limit LY.

For ingot 8223, the two slices are largely uniform. Sample 8223-B shows typical RL intensities and a normalized TSL/RL peak near 35–37 K; its transmittance is consistent with transparency above about 450 nm. Sample 8223-T has the same RL peak and width but a systematically lower RL intensity at all temperatures; the origin of this behaviour remains unexplained. Normalized TSL/RL and transmittance of the the top and bottom samples differ only mildly below 400 nm; this points to small longitudinal variations along the ingot. No atypical trap families are detected.

Samples from ingot 8228 are atypical. Sample 8228-T exhibits a narrow TSL band at 100–150K centered near 1.8 eV, indicative of specific traps. In the same temperature range its RL FWHM evolves anomalously, consistent with the emergence of an additional defect-related emission band correlated with those traps. Normalized TSL/RL measurements reveal a higher concentration of lower temperature traps with respect to other samples. This sample also shows a clear transmittance loss below 450 nm, which explains the yellow tint of the ingot. Sample 8228-B exhibits a normalized TSL/RL peak of much lower intensity and a modest transmittance loss below 450 nm. These features indicate longitudinal disuniformity and a different and varying trap population along the growth axis. Interestingly, the CUPID crystal from this ingot (8228) has a lower LY than that from 8223.

The G833 bulk cuts behave coherently. Samples G833-1, G833-2, and G833-4 have identical normalized TSL/RL with a peak near 35–37 K and stable RL parameters in first approximation. G833-2, prepared from recycled heads, is slightly more yellow and shows reduced transmittance below 400 nm, yet its RL band remains in the transparent region. G833-4, from mixed recycled and bulk material, matches G833-1/2 after normalization, suggesting that appropriate mixing and annealing can recover optical and scintillation quality.

These measurements provide concrete guidance to SICCAS, which will investigate correlations with the specific growth parameters of each ingot.

Chapter 6

Thin Al Coatings for Li_2MoO_4 Scintillating Cryogenic Calorimeters

The scintillation light of $\text{Li}_2^{100}\text{MoO}_4$ crystals is a key tool for background mitigation in next-generation $0\nu\beta\beta$ experiments: together with the high ^{100}Mo Q -value, located in a region largely free of β/γ backgrounds, it enables effective α rejection. However, Li_2MoO_4 is only a moderate intrinsic scintillator and its light yield is relatively low. Targeted strategies to enhance scintillation performance should therefore be explored to inform the design of experiments beyond the generation of CUPID.

This chapter investigates the potential of depositing thin Al coatings on crystal surfaces as a means to increase reflectivity. This raises the light-collection efficiency at the LDs without adding passive detector elements, such as reflective foils around each crystal. However, superconducting films can modify the detector response and degrade the energy resolution. Accordingly, measurements of the calorimetric response of Al-coated crystals are presented.

6.1 Improving Light Collection Efficiency With Al Coatings

Already 15 years ago, during the CUORE R&D phase, composite cryogenic calorimeters were identified as a viable route to tag surface events and reject degraded α backgrounds [204, 205]. In that framework, thin superconducting aluminum films were proposed to discriminate surface interactions by exploiting pulse-shape differences measured with sensors faster than NTDs, e.g. NbSi [206]. More recently, the CROSS collaboration has investigated pulse-shape discrimination of surface α and β interactions by depositing Al films on Li_2MoO_4 and TeO_2 crystals with NTD readout [207, 208]. Other recent studies considered normal-metal coatings such as Pd, as well as composite Al-Pd films, tested both as continuous layers and as patterned grids [209].

In the following, a complementary use of Al coatings is investigated. Besides their potential for pulse-shape discrimination, Al films can act as reflective layers to enhance the light collection efficiency at the LD, improving the separation between the α and β/γ bands and thus the discrimination power. Reflective foils have already proven effective in demonstrators such as CUPID-Mo [150]. However, foils are passive elements: they increase the material budget and the risk of radioactive contamination, and they suppress or even prevent multi-crystal coincidences. Both effects

can worsen the backgrounds in a large-scale experiment, making reflective foils impractical for next-generation $0\nu\beta\beta$ searches.

By contrast, Al coatings directly deposited on Li_2MoO_4 crystal surfaces avoid these limitations. Maximizing the benefit requires coating multiple faces, with an optimal configuration in which five faces are coated and only the face toward the LD remains uncoated.

An additional potential advantage of coatings is mitigation of moisture-induced surface degradation, crucial for hygroscopic crystals such as Li_2MoO_4 . A continuous Al film can act as a barrier against ambient humidity, limiting haloing and opacity on coated crystal surfaces.

A superconducting Al layer on most of the surface, however, may alter the detector response and degrade the energy resolution. The measurements presented below assess the impact of Al coating on the calorimetric performance of Li_2MoO_4 crystals.

6.2 Thin Aluminum Coatings for Li_2MoO_4 Crystals

Coatings are tested on natural $2 \times 2 \times 2 \text{ cm}^3$ Li_2MoO_4 crystals. This size allows operation above ground¹, which facilitates the R&D.

6.2.1 Coating depositions

Reflective coatings are produced at INFN-LNL (Laboratori Nazionali di Legnaro) with a vacuum chamber equipped with five magnetron sputter sources, enabling simultaneous deposition on five faces of the cube (the face that will be positioned in front of the LD, is left uncoated). Initial deposition trials showed adhesion issues and partial delamination, prompting an optimization of surface preparation and deposition parameters. The adopted workflow consists of: (i) pre-cleaning with tetrachloroethylene, acetone, and isopropyl alcohol; (ii) a 48 h bake at 120°C to reduce outgassing and improve base pressure during deposition; and (iii) Al sputtering in Ar. With this workflow and the multi-source geometry, uniform coverage is obtained on all coated faces with a nominal thickness of $\sim 1.5\text{--}2.0 \mu\text{m}$ per face. Two crystals have been prepared with these specifications. Two $2 \times 2 \times 2 \text{ cm}^3$ crystals were left uncoated and will serve as reference standards. To benchmark the effect of Al reflectivity, a further crystal of the same size was coated with a non-reflective Cu film of the same thickness. Photographs of the Al- and Cu-coated crystals are shown in the left panel of Fig. 6.1.

6.2.2 Low-temperature properties of Al coatings

Al becomes superconducting below $T_c \simeq 1.2 \text{ K}$ [210]. At the operating temperature of cryogenic calorimeters ($\sim 15 \text{ mK}$), the electronic specific heat is exponentially suppressed, so only the lattice term is relevant. According to Eq. (2.6), the absorber heat capacity should be minimized to maximize the signal amplitude. For a $2 \times 2 \times 2 \text{ cm}^3$ crystal with five faces coated by a $2 \mu\text{m}$ Al film (total Al mass $\sim 10 \text{ mg}$), the added heat capacity is negligible: using Eq. (2.7) with $\Theta_D(\text{Al}) \simeq 427 \text{ K}$ gives $C_{\text{Al}}(15 \text{ mK}) \sim 0.2 \text{ MeV/K}$, while the crystal heat capacity is $\sim 183 \text{ MeV/K}$. The Al contribution on the total heat capacity of the absorber is thus at the per-mille level, leaving the intrinsic thermal

¹Operating a standard $4.5 \times 4.5 \times 4.5 \text{ cm}^3$ CUPID crystal above ground would lead to a high pile-up rate, complicating the analysis.

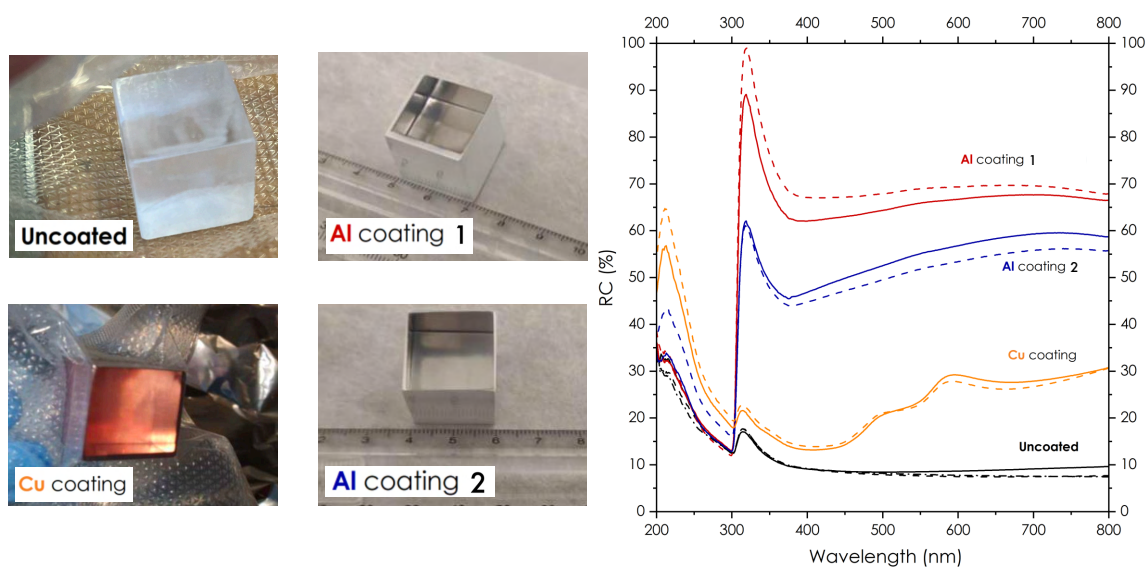


Figure 6.1: *Left:* $2 \times 2 \times 2 \text{ cm}^3$ Li_2MoO_4 samples: uncoated, Cu-coated, and two Al-coated. *Right:* Hemispherical reflectance RC% measured with the integrating sphere as explained in the text. Solid and dashed curves are repeat runs on the same face after a 90° rotation to check reproducibility.

gain essentially unchanged. By contrast, Cu is a normal metal at $\sim 15 \text{ mK}$, so its electronic specific heat does not vanish. A Cu layer of comparable thickness therefore dominates the total heat capacity degrading the intrinsic gain of the absorber.

The presence of a superconducting layer can nonetheless modify phonon transport in the absorber, which sets the equilibrium phonon population sampled by the NTD sensor. Energy deposits near an Al-coated surface generate a flux of high-energy phonons that can break Cooper pairs in the film, altering the subsequent phonon dynamics and the pulse shape [209, 207]. This effect is expected to be less relevant for bulk events, whose phonons partially thermalize before interacting with the coating, potentially introducing position dependent effects.

Finally, in the superconducting state the electronic thermal conductivity of Al is strongly suppressed at $\sim 15 \text{ mK}$. Given the small and possibly non-uniform film thickness, incorporating this and interfacial thermal resistances into a complete thermal model is non-trivial, but these effects may influence the detector response.

6.2.3 Radiopurity of Al coatings

The background-reduction benefit of reflective coatings is viable only if the coating itself does not introduce radioactive contaminants. A preliminary screening of the deposited Al layers was performed above ground with a Si surface-barrier detector. No activity was observed above the instrument sensitivity, yielding an upper limit of $< 1 \mu\text{Bq}/\text{cm}^2$. This indicates that, within the current sensitivity, the coating process does not introduce significant contamination. Should the coatings prove effective for enhancing light collection, dedicated low-background assays targeting the $\sim\text{nBq}/\text{cm}^2$ level will be carried out.

6.2.4 Reflectance of Al coating




Figure 6.2: *Left*: Photograph of the reflectance setup. The Li_2MoO_4 sample is mounted at the measurement port of the integrating sphere of a Perkin-Elmer Lambda 950. *Right*: Working principle of the reflectance measurement: a monochromatic beam illuminates the sample through the uncoated face, and the integrating sphere collects the flux reflected back into the sphere.

To quantify the optical benefit of the coatings, the hemispherical reflectance was measured for the four $2 \times 2 \times 2 \text{ cm}^3$ Li_2MoO_4 crystals described above: two Al-coated, one Cu-coated, and one uncoated. Measurements were performed with a Perkin-Elmer Lambda 950 equipped with a 60 mm integrating sphere and an 8° wedge (Fig. 6.2). The sample was mounted at the sphere port; a monochromatic beam illuminated the uncoated face and the sphere integrated the reflected light. Spectrally resolved measurements were obtained by scanning the incident wavelength.

Two clarifications apply to the results shown in the right panel of Fig. 6.1. First, solid and dashed curves correspond to two independent acquisitions of the *same* face, taken after rotating the sample by 90° . This checks reproducibility given the non-optical quality of the surface. Second, the ordinate label RC% denotes the hemispherical reflectance, expressed as a percentage, obtained when the integrating sphere is referenced with its calibrated port plugs (i.e., the sphere is closed by the plugs rather than by the crystal). In principle this yields an absolute reflectance. In practice, because the reference plugs were found to be deteriorated, the RC% values should be used as *relative* comparisons among coatings, not as absolute numbers.

Within these caveats, Fig. 6.1 show that Al-coated surfaces have a higher reflectance than uncoated Li_2MoO_4 across the scanned range, while Cu, as expected, provides a smaller enhancement. Although the two Al-coated samples differ slightly, their reflectances are both strongly enhanced relative to the uncoated crystal.

6.3 Impact of Al Coating on Calorimetric Performance

Reflectance measurements indicate that Al coatings can enhance light collection efficiency, but operation as a cryogenic calorimeter is required to assess their impact on the heat channel. This section presents measurements performed at the R&D cryogenic facility of the University of Milano–Bicocca to evaluate the effect of Al coatings on calorimetric performance.

6.3.1 Cryogenic facility

Measurements are performed in an Oxford Instruments TL 200 wet $^3\text{He}/^4\text{He}$ dilution refrigerator, the same model described in Sec. 4.2. No liquefier is present; helium is supplied by periodic refills from external dewars. To maximize the available volume, no internal lead shielding is installed. The cryostat is *not* equipped with spring-based vibration damping system.

The cryostat is equipped with a noise thermometer (NT), which provides accurate absolute temperature measurements of the detectors, provided they are well thermalized to the mixing chamber. The NT is shown in Fig. 6.3. Thermalization of the NT to the mixing chamber is ensured by fastening it to the top TSP plate.

Section 2.3.1 details the readout and biasing electronics for NTD-equipped cryogenic calorimeters. The front-end boards and JFET amplifiers are the same used in CUORE, described in Ref. [131, 132]. The back-end electronics has been recently updated to match the CUPID specifications and uses a 24-bit ADC with sampling rates up to 10 kHz.

Digitized waveforms are stored for offline analysis. Apart from the analog anti-aliasing Bessel filter, all subsequent shaping and filtering are applied digitally using the analysis pipeline described in Sec. 4.3.

6.3.2 Detector geometry and data taking

This section presents the setup used for the calorimetric characterization of Al-coated Li_2MoO_4 crystals. Figure 6.3 shows the detectors installed in the cryostat experimental volume. The paragraphs below summarize the detector geometry and the data taking.

Detector design

Four $2\times 2\times 2\text{ cm}^3$ crystals are tested: the two Al-coated samples of Fig. 6.1 (LMO-Al-1 and LMO-Al-2) and two uncoated references (LMO-A and LMO-B). Each crystal is instrumented with an NTD thermistor glued directly to the crystal surface, except for LMO-Al-1, which carries an additional NTD glued on top of the Al coating (LMO-Al-1c). This configuration allows a direct check for differences in pulse shape and intrinsic gain between an NTD on bare Li_2MoO_4 and on the Al film. A total of five sensors are therefore read out. For the coated samples, the Al layer was locally removed by mechanical scratching to expose the Li_2MoO_4 surface for NTD gluing. All the NTDs are of the same type and are glued to the crystal in the same way using a reproducible procedure.

Light detectors were installed to compare the light yield of the coated and uncoated samples. However, abnormal noise affecting the setup led to elevated LD thresholds, preventing their use for reliable LY measurements.

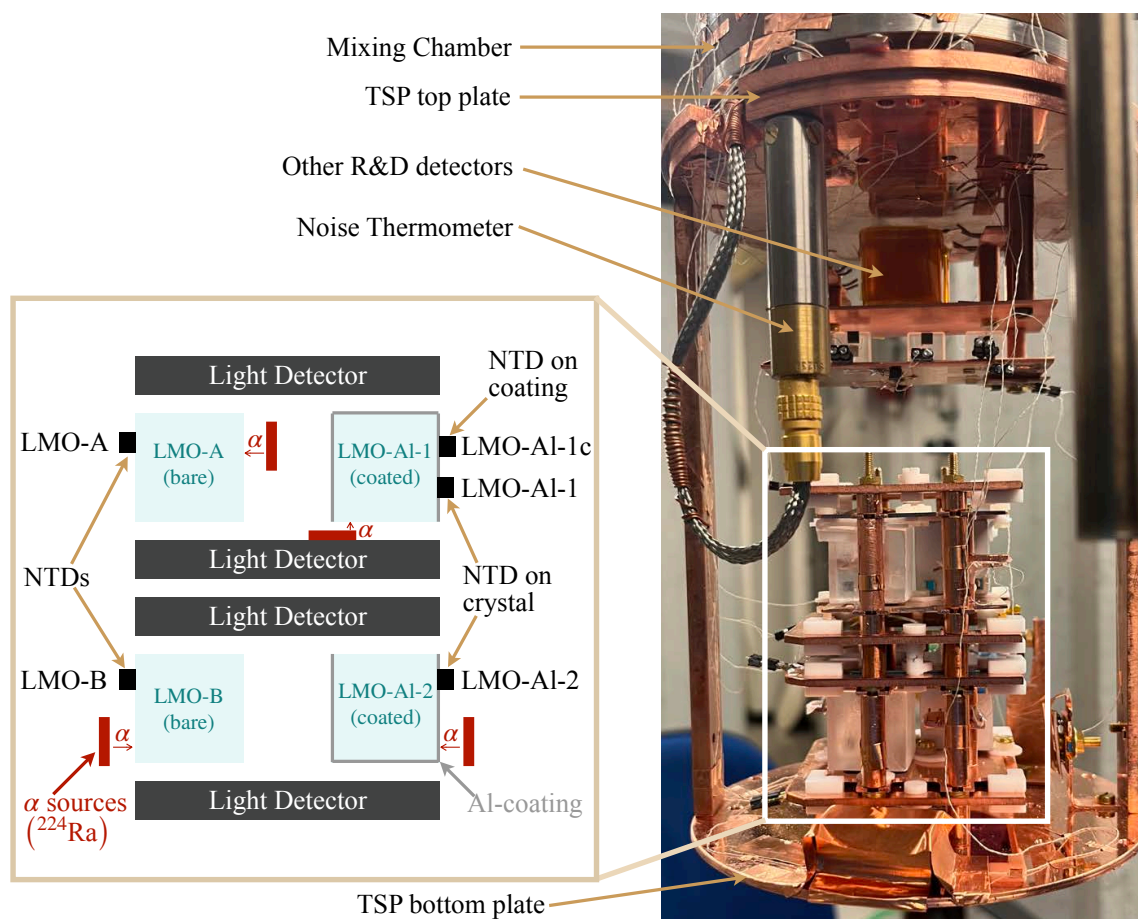


Figure 6.3: Detectors assembled inside the cryostat experimental volume. This run also includes R&D measurements for other cryogenic calorimeters installed in the upper part of the setup. *Inset:* setup used to characterize Al-coating performance. LMO-A and LMO-B are uncoated references; LMO-AI-1 and LMO-AI-2 are Al-coated on five faces (the face toward the LD is uncoated). Crystals and LDs are mounted in copper holders and secured with PTFE elements. Each crystal is read out with an NTD glued to the crystal surface; LMO-AI-1 carries a second NTD glued on top of the Al film (LMO-AI-1c). LDs face the uncoated faces. Monochromatic α sources, implemented as shallow ^{224}Ra implants (red rectangles), illuminate selected crystal sides.

Calibration sources

An external ^{232}Th source is placed outside the cryostat during data taking to increase the statistics at the ^{208}Tl 2615 keV line.

Monochromatic α lines are employed to study the detector response to surface α events, with particular focus on coated surfaces. α sources were positioned to illuminate selected crystal faces, specifically the Al-coated face of one crystal and the uncoated face of another, to probe possible

pulse-shape differences. In addition, because the crystals are not equipped with stabilization Si heaters, physics peaks must be used for thermal gain stabilization. The 2615 keV γ line alone often provides insufficient statistics, so the monochromatic α lines serve also as a robust stabilization handle.

The α sources are prepared following Ref. [211]. Adhesive copper tape is exposed to a primary ^{228}Ra source. The first α nuclide in the chain, ^{228}Th , decays to ^{224}Ra ; the recoil implants ^{224}Ra at very shallow depth in the tape. ^{224}Ra then decays (half-life 3.63 days) to ^{220}Rn , followed by ^{216}Po , ^{212}Pb , ^{212}Bi , and ^{212}Po . Because the recoiling daughters are implanted at ultra-low depth, the subsequent α emissions are effectively monochromatic. The line intensity is set by adjusting the exposure time of the tape to the primary source.

Data taking

The TSP temperature, monitored with the NT, was set to ~ 22 mK. R - P curves were acquired for all the detectors; those for the NTDs on the Al-coated and uncoated Li_2MoO_4 crystals are reported below.

The rate of the α source was higher than expected, producing severe pile-up across all measurements. Data-quality cuts were therefore applied to reject piled-up events, as described in Sec. 4.3. The resulting selection efficiency was very low, substantially reducing the available statistics. Since the source is internal, the rate could not be mitigated during the run making compromises unavoidable. Owing to the relatively short half-life of ^{224}Ra (3.63 days), the α intensity decreased over time, improving the situation progressively; nevertheless, it did not reach an acceptable level even toward the end of data taking.

Near the end of the data taking, a ~ 3 h calibration run was collected with all NTDs on the Li_2MoO_4 crystals biased at an ohmic working point, $I_{\text{bias}} \approx 0.1$ nA (0.1 V/10 G Ω). This WP does not optimize energy resolution, which is typically achieved closer to inversion (see Sec. 4.5.7). However, a WP near inversion can introduce non-linearities in the NTD response. By contrast, at an ohmic WP the detector is effectively linear, enabling a direct comparison of performance and pulse shape between coated and uncoated crystals. Data were acquired at a 5 kHz sampling frequency with a 1 kHz Bessel cutoff, and processed with the analysis pipeline described in Sec. 4.3. The results of this measurement are discussed below.

6.3.3 Calorimetric performance

The effect of a $2\ \mu\text{m}$ superconducting Al coating on the total heat capacity of the absorber is negligible. However, the low-temperature properties of Al discussed above can modify the detector thermal dynamics and thus its calorimetric performance.

R-P curve

R - P curves and I - V characteristics were acquired for all NTDs with the NT reading ~ 22 mK. The shape of the R - P curve is set by the heat capacities and thermal conductances of the full system. In principle, once all elements are laid out (including the Al coating), one can solve the thermodynamic equilibrium equations and express the NTD resistance as a function of injected power, $R(P)$, following the procedure described in Ref. [212]. In practice, building such a model requires several assumptions, and its free parameters must be constrained by R - P curves taken at multiple temperatures. In this study only a single R - P curve at 22 mK was built, which renders




Figure 6.4: R - P curve (a) and R_{dyn} - P curves (b) for Al-coated and uncoated Li₂MoO₄ crystals. Error bars are shown only in panel (a) and are dominated by the nominal 10% uncertainty on the load resistors R_L . In panel (b), the power coordinate of each point is assigned to the midpoint of the window used to compute the derivative of the I - V curve.

the fit underconstrained and potentially non-reproducible. Nevertheless, because both uncoated and Al-coated Li₂MoO₄ crystals are present in the same assembly with all other conditions equal, a differential comparison is still informative: any systematic differences between their R - P curves would indicate a measurable impact of the coating on the effective thermal model of the detector.

The R - P curves in Fig. 6.4a indicate comparable base resistances for all crystals. The small spread in R_{base} is attributed to subtle variations in NTD gluing rather than to differences in absorber temperature. Gluing affects differential thermal contraction, altering the mechanical stress on the Ge NTD and thereby shifting its R_0 parameter, and thus R_{base} .

At higher power the curves show a clear shape difference between the Al-coated samples (LMO-Al-1/2) and the uncoated references (LMO-A/B). No significant difference is observed between the NTD glued on top of the Al film (LMO-Al-1c) and the NTD glued directly to the crystal (LMO-Al-1). To better appreciate these trends, the dynamic resistance $R_{\text{dyn}} = dV/dI$ was obtained by computing the derivative of the I - V curve numerically. The resulting R_{dyn} - P curves in Fig. 6.4b confirm that all NTDs on Al-coated crystals behave similarly to each other and distinctly from the uncoated ones: in the high-power regime, the uncoated samples systematically tend to lower dynamic resistances, consistent with the trends observed in the R - P curves.

An ohmic WP at $I_{\text{bias}} \approx 0.1$ nA (0.1 V/ 10 GΩ) was selected to compare the intrinsic gain (sensitivity) and the energy resolution of coated versus uncoated crystals. This WP lies at the lowest-power end of the R - P curves in Fig. 6.4a and is well within the ohmic regime. The R_{dyn} - P curves corroborate this, showing R_{dyn} coincident with R_{NTD} (from the R - P curve) up to $\sim 10^2$ pW.

Intrinsic gain

The intrinsic gain (sensitivity) is computed as in Secs. 4.5.7: a monochromatic line in the spectrum is selected and the corresponding pulse amplitude is expressed in nV/keV. The monochromatic α

line of ^{224}Ra at 5.685 MeV was used. The results are reported in Table 6.1 and indicate that Al-coated crystals have lower sensitivities.

However, these values cannot be compared directly across channels, as the intrinsic gain also includes the NTD response and depends on the working resistance. For a meaningful comparison, the sensor contribution must first be factored out. In the ohmic regime the dynamic resistance coincides with the NTD resistance and, at the selected WP (low power), it approaches the base resistance:

$$R_{\text{dyn}} = \frac{dV}{dI} \stackrel{\text{ohmic}}{=} \frac{V_{\text{NTD}}}{I_{\text{bias}}} = R_{\text{NTD}} = R_{\text{WP}} \stackrel{\text{low power}}{\approx} R_{\text{base}}$$

because nonlinearities from electro-thermal feedback do not arise. In these conditions, Eq. (2.9) fully determines the NTD response to a given temperature pulse generated in the absorber. This implies that the intrinsic gain (in nV/keV) is directly proportional to R_{WP} , which at this WP coincides with R_{base} . Under the (strong) key assumption that all crystals share the same absorber temperature², the sensitivity can, to first order, be decoupled from the NTD response simply by dividing it by the working resistance R_{WP} . The resulting quantity has units of current per energy and is referred to as the *normalized sensitivity*; values are reported in Tab. 6.1. A clear separation is observed: the uncoated crystals (LMO-A/B) yield ~ 0.5 fA/keV, whereas the Al-coated crystals (LMO-Al-1/2) reach only ~ 0.2 fA/keV. This indicates that the superconducting coating interacts with the phonons changing their equilibrium distribution, plausibly because a fraction of the high-energy phonons generated during the particle interaction is absorbed in the Al layer via Cooper-pair breaking. The NTD glued directly on the coating (LMO-Al-1c) shows an even lower sensitivity (~ 0.1 fA/keV), consistent with this interpretation.

Ch.	$R_{\text{WP}} = R_{\text{base}}$ [M Ω]	Sens. [nV/keV]	Norm. Sens. [fA/keV]	FWHM Res. (^{224}Ra α) [%]	DT [ms]	RT [ms]
LMO-A	~ 5.5	~ 2.39	~ 0.435	~ 0.832	33 ± 2	9.1 ± 1.0
LMO-B	~ 10	~ 4.76	~ 0.476	~ 1.185	28 ± 1	7.5 ± 0.4
LMO-Al-1	~ 5.4	~ 1.09	~ 0.202	~ 1.785	18 ± 2	5.6 ± 1.0
LMO-Al-1c	~ 4.3	~ 0.42	~ 0.097	~ 5.276	13 ± 3	7.7 ± 1.8
LMO-Al-2	~ 2.9	~ 0.55	~ 0.189	~ 2.223	16 ± 3	6.7 ± 1.5

Table 6.1: Comparison of detector performance for Al-coated and uncoated Li_2MoO_4 crystals at the ohmic working point $I_{\text{bias}} \approx 0.1$ nA (0.1 V/10 G Ω). The intrinsic gain (sens.) is derived from the monochromatic α line of ^{224}Ra at 5685 keV. The norm. sens. is obtained by dividing the sensitivity by R_{WP} , which at this WP coincides with R_{base} . Reported RT and DT values are given as median \pm MAD of the distributions corresponding to all the events in the range [3,8] MeV.




Figure 6.5: Energy spectra of an Al-coated and an uncoated Li_2MoO_4 crystal, calibrated with a single-point scale using the ^{224}Ra 5.685 MeV α line.

Energy resolution

Figure 6.5 compares the energy spectra of LMO-A (uncoated) and LMO-Al-1 (Al-coated). Because the working point was set in the ohmic regime (far from the optimal WP) with the detectors being operated near R_{base} , the energy resolutions are poor; in particular, most of the β/γ population lies below threshold at this WP. Consequently, the analysis is restricted to the [3, 8] MeV interval, which is dominated by α interactions. Even within this limited window, the spectrum shows that the coated crystal exhibits visibly worse performance. A stronger deviation from the linear calibration is observed, indicating a more pronounced nonlinearity. Table 6.1 reports the FWHM energy resolution obtained by fitting the ^{224}Ra 5.685 MeV α line and confirms that Al-coated crystals have consistently worse resolution. The noise level was verified to be similar across channels by comparing the ANPS of all detectors. For comparable noise, a lower signal amplitude leads to a worse energy resolution; thus, the degraded resolutions observed for the Al-coated crystals (most notably for the NTD glued on top of the coating) are entirely attributed to their reduced intrinsic gain, as discussed above.

Time constants

At an ohmic working point the detector response is linear and electrothermal feedback does not shape its frequency response. Since all detectors are operated at the same ohmic WP, their time constants can be compared directly. Figure 6.6 shows the rise-time (RT) and decay-time (DT) distributions for pulses with reconstructed energies in the range [3, 8] MeV, dominated by α (and α +nuclear-recoil) events. The median \pm MAD values are reported in Table 6.1.

The data indicate that RT is similar for Al-coated and uncoated crystals, whereas DT separates clearly: uncoated crystals exhibit longer DTs than coated ones. In the absence of a full thermal model, a detailed interpretation is not attempted; nevertheless, the observed DT separation suggests that the Al coating affects the thermalization of the absorber to the heat sink after a particle interaction. Interestingly, the channel with the NTD glued on top of the Al film (LMO-Al-1c) shows RT/DT values comparable to those of NTDs glued directly on the crystal, indicating that any insulating effect of the superconducting layer—if present—has limited impact on the dominant thermal path (assumed through the NTD and its wiring). However, its distributions are broader, consistent with larger fluctuations in RT/DT estimation due to the lower sensitivity, which reduces

²In this case, differences in R_{base} are entirely attributed to NTD properties (e.g., small gluing-induced stress variations).

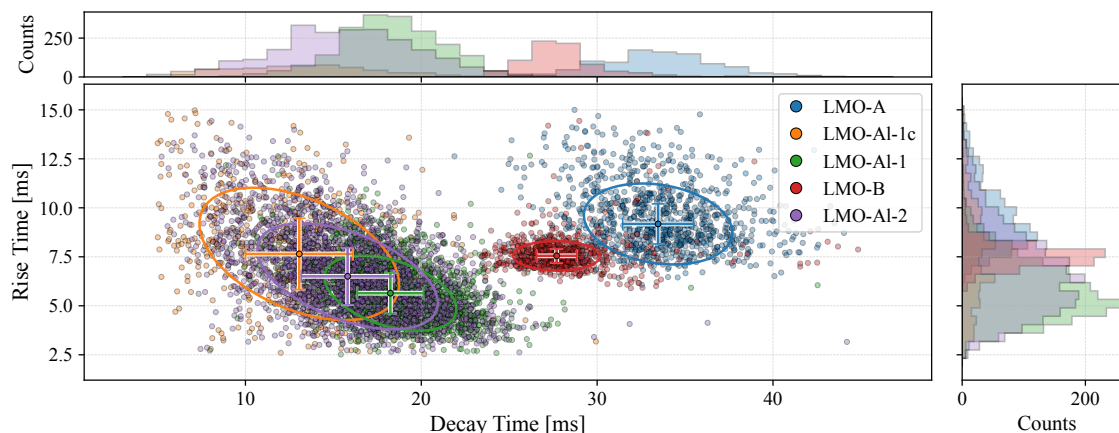


Figure 6.6: Comparison of rise time and decay time distributions for coated and uncoated crystals. The top/right panels show the distributions of DT/RT. For each channel, the solid ellipse is a robust $\sim 68\%$ (1σ) covariance contour centered at the medians, and the cross-whiskers indicate median \pm MAD along each axis.

timing precision and can mask subtler differences. As shown in Fig. 6.3, α particles illuminate the *uncoated* face of LMO-Al-1 and the *Al-coated* face of LMO-Al-2. Despite this difference, the two detectors exhibit very similar RT/DT distributions. Moreover, LMO-Al-1 (illuminated on the uncoated face) shows RT/DT distributions different to the bare crystals, indicating that the observed pulse-shape features cannot be ascribed to interactions occurring directly on a coated surface. Owing to the poor energy resolution at the chosen WP, the β/γ population lay below threshold, preventing a comparison of pulse shapes and time constants between α and β/γ events.

6.4 Outlook

This preliminary study shows a clear trade-off: while Al coatings increase hemispherical reflectance and thus the potential light collected at the LD, they also have a non-negligible impact on calorimetric performance-reducing sensitivity and, at equal noise, worsening the energy resolution. With the present statistics it is not yet possible to pinpoint the underlying mechanism. Developing a thermal model for Al-coated Li_2MoO_4 absorbers is therefore a priority to identify, and possibly mitigate, the effects introduced by the superconducting layer.

To pursue this, a new data-taking campaign will be carried out in the Milano-Bicocca cryogenic facility with an updated detector assembly. The setup will include NTL-assisted light detectors that enable the detection of scintillation signals, allowing a direct comparison of the light yield from coated versus uncoated crystals. These Ge LDs use transparent Indium Tin Oxide (ITO) electrodes deposited on the two opposite faces of Ge, creating a uniform electric field through the germanium. This differs from the NTL LDs foreseen for CUPID, which use non-transparent electrodes that generate an in-plane field. The ITO coatings are produced in collaboration with INFN-LNL (the same group depositing Al on Li_2MoO_4) as part of a broader R&D project. Preliminary measurements conducted in the past months operating ITO LDs at ~ 15 mK with artificial




Figure 6.7: *Left*: Rendering of the detector design: a two-floor stack hosting Cu-coated, Al-coated, and uncoated Li_2MoO_4 crystals, interleaved with ITO light detectors. The geometry was designed to align each crystal with the LD diameter so that the uncoated faces maximally overlap the transparent ITO electrodes. *Right*: Photograph of the assembled module. NTDs were glued to the uncoated faces, except for the Al-coated crystal on the bottom floor, which has an NTD glued directly on the coating and is visible in the picture. For R&D purposes, three additional ITO LDs were installed above the module shown in the rendering.

LED-generated light pulses indicate good amplification.

The new setup serves a dual role: (i) qualify the performance of ITO-based LDs coupled to conventional (uncoated) Li_2MoO_4 ; and (ii) exploit the NTL amplification to quantify the light-yield gain achievable with Al-coated crystals with respect to the bare ones. The detector geometry (see Fig. 6.7) has been tailored to maximize the overlap between the uncoated Li_2MoO_4 face and the circular ITO electrodes, so that most photons impinge where the electric field is strongest and most uniform. The assembly will reuse the crystals shown in Fig. 6.1, including the Cu-coated sample. Although the heat capacity of Cu is expected to suppress the intrinsic gain, R-P curves from the NTD installed on the Cu-coated crystal will allow to study how a normal-metal film perturbs the effective thermal network, providing guidance for interpreting the superconducting case. As in the present study, the geometry is designed to equalize absorber temperatures across channels, enabling meaningful, sensor-independent comparisons of intrinsic gain.

These measurements are essential to assess whether reflective Al coatings can be made compatible with the requirements of next-to-next generation $0\nu\beta\beta$ searches based on scintillating cryogenic calorimeters like Li_2MoO_4 .

Conclusions and Outlook

With CUORE nearing the end of data taking, CUPID is transitioning from R&D to construction, with demonstrators already operating to final detector specifications. At the detector’s core are 1596 isotopically enriched $\text{Li}_2^{100}\text{MoO}_4$ crystals that must simultaneously deliver excellent calorimetric performance, high light yield, and ultra-high radiopurity. CUPID has engaged SICCAS to establish a stable and reproducible industrial growth line capable of sustaining large-scale crystal production with high throughput. CUPID’s transition to construction is the natural stage at which to finalise the production process in close collaboration with the producer. This dissertation focused on crystal procurement and validation during this phase, referred to as *pre-production*.

During pre-production, SICCAS produces multiple batches of natural and enriched crystals to establish a stable industrial growth process, while CUPID operates them as scintillating cryogenic calorimeters under experiment-like conditions and provides systematic feedback. The aims are to verify convergence toward CUPID crystal requirements and to define a robust, reproducible validation protocol to be applied batch by batch once full-scale production begins, the CCVRs (CUPID Crystal Validation Runs).

The work documented in this dissertation was performed during the early stage of pre-production. Although crystal growth is fully delegated to SICCAS, it quickly became evident that a detailed understanding of the production flow was necessary to communicate effectively with the producer and to link growth parameters to detector performance. Accordingly, the SICCAS production process was documented end-to-end—from raw powders to finished crystals—so that CUPID and SICCAS could operate within a shared technical framework and common vocabulary (Chapter 3). This knowledge will also provide the basis for CUPID to define a radiopurity control protocol that SICCAS will implement throughout the entire production chain—a prerequisite for the transition to large-scale growth with enriched material.

Crystal validation was carried out through two pre-CCVR campaigns at the LNGS Hall C underground cryogenic facility (Chapter 4). The first (pre-CCVR1) established reference performance using enriched crystals provided by NIIC (the supplier of the CUPID-Mo demonstrator), which now serve as benchmarks; the second (pre-CCVR2) characterized the first batch of natural Li_2MoO_4 crystals grown by SICCAS. These campaigns set consistent procedures for assessing calorimetric and scintillation performance, including working-point optimization, signal-to-noise characterization, and energy-resolution metrics. The results show that the first batch of SICCAS crystals meet CUPID calorimetric expectations, with light yields comparable to that of reference crystals. Early pre-production growths do not prioritize radiopurity. The crystals tested in pre-CCVR2 were not grown from raw powders, ancillaries, or consumables meeting the final radiopurity specifica-

tions. Nevertheless, radiopurity was still assessed. The measurements show that re-crystallization yields substantial purification even when starting from non-radiopure powders. They also indicate how the recycling strategies adopted to minimize isotope losses affect the achieved radiopurity, providing a precious feedback for SICCAS.

At the time of writing, the first enriched $\text{Li}_2^{100}\text{MoO}_4$ crystal grown by SICCAS is being measured in a dedicated run (pre-CCVR3) that applies the methodologies established in this work. One NIIC crystal characterized here was included to provide a known reference. A second batch of enriched SICCAS crystals has arrived at LNGS and will be tested next (pre-CCVR4) under the same protocol. These measurements are the first deployment of the full validation chain defined in this dissertation. Subsequent pre-CCVR runs will place greater emphasis on radiopurity verification, as SICCAS begins implementing the radiopurity control strategy that CUPID will derive from the growth-process knowledge outlined in this work. Calorimetric and scintillation performance will continue to be monitored.

In addition to crystal validation, this dissertation examined the scintillation properties of Li_2MoO_4 , which are central to CUPID's particle-identification capability. Li_2MoO_4 is a modest scintillator, and its light yield is sensitive to growth-induced defects and impurities. Samples cut from pre-production ingots were studied over a wide temperature range using radioluminescence and thermally stimulated luminescence (Chapter 5). The measurements indicate partial longitudinal inhomogeneities along the Bridgman growth axis and suggest correlations between these observations and light yield, although the evidence remains preliminary. Further measurements on a larger sample set from later pre-production batches, ideally correlated with light-yield data from CUPID crystals grown from the same ingots, will be necessary to confirm these preliminary trends. These results complement the pre-CCVR characterizations and can provide SICCAS with practical guidance on how growth parameters influence scintillation quality.

Given the modest light yield of Li_2MoO_4 , an R&D activity on thin aluminum coatings was carried out to investigate their potential to enhance light collection (Chapter 6). Initial measurements show a clear trade-off: coatings increase hemispherical reflectance, but they alter the thermal response and degrade the signal amplitude and thus the SNR. Follow-up studies with NTL-assisted light detectors and refined thermal modeling are planned to quantify the net light-yield gain at the light detector and to clarify the origin of the coating-induced change in intrinsic gain.

In conclusion, this work represents an early yet essential step toward the realization of CUPID. It established a shared technical framework with the crystal producer, defined reproducible validation procedures, and delivered quantitative feedback that already inform the optimization of growth and detector performance. The strategies developed here will guide all upcoming pre-CCVR campaigns and future validation runs, ensuring that every crystal delivered to CUPID during large-scale production meets the requirements. Together, the measurements and methods presented in this dissertation mark a decisive advance toward constructing CUPID and advancing the search for neutrinoless double-beta decay.

Appendix A

Dilution Cryostats

To operate cryogenic calorimeters at temperatures around ~ 15 mK, a dedicated cryogenic infrastructure is required. This is achieved using dilution refrigerators, which are specifically designed to reach and maintain such ultra-low temperatures. Although the physics of dilution refrigerators is not directly related to the detection mechanism of macro-calorimeters, it is nevertheless useful to summarize their basic operation. In fact, several components of the cryogenic system can influence data-taking performance, particularly with respect to temperature stability and noise injection, for example through mechanical vibrations. A comprehensive description of the working principles of low temperature technologies can be found in Ref. [213].

Liquid helium is a commonly used refrigerant. At atmospheric pressure, its boiling point is approximately 4.2 K. A generic helium sample contains a mixture of ^4He and ^3He isotopes, with the isotopic composition determined by their natural abundances. Given that ^3He has an isotopic abundance of just 0.000137%, the thermodynamic properties of standard helium are dominated by ^4He , limiting passive cooling to around 4.2 K at standard atmospheric pressure.

To reach lower temperatures, active cooling is employed. In a system at equilibrium containing both liquid and gaseous helium, removing the vapor with a pump lowers the pressure, which in turn enhances evaporation. This process extracts heat and thereby reduces the temperature of the system. To maintain this condition, a continuous supply of liquid helium is required. This concept of *circulation* allows the removal of heat at a rate

$$\frac{dQ}{dt} = L_0 \frac{dn}{dt} \tag{A.1}$$

where L_0 is the latent heat of evaporation at 0 K, and $\frac{dn}{dt}$ is the molar flow rate of helium, related to the pressure gradient induced by the pump. Using this method, temperatures down to ~ 1 K for ^4He and ~ 300 mK for pure ^3He can be achieved.

To reach temperatures even lower than this, a different active cooling mechanism must be employed, based on the unique properties of a mixture of ^3He and ^4He at low temperatures. Dilution cryostats that utilize such mixtures are the most common systems used to maintain cryogenic calorimeters at their base temperature of ~ 15 mK.




Figure A.1: Phase diagram of the $^3\text{He}/^4\text{He}$ mixture. Below the critical temperature $T_c \simeq 870$ mK, the mixture undergoes spontaneous phase separation into a concentrated ^3He phase and a dilute phase. Reprinted from Ref. [214].

A.1 $^3\text{He}/^4\text{He}$ Dilution Cryostat: Design and Operation

A cryostat must fulfill two primary functions: insulating the experimental setup from thermal input from the surrounding ~ 300 K environment, and continuously removing the heat generated from unavoidable thermal leaks. The rate at which heat can be removed is referred to as the *cooling power* of the cryostat.

A dilution cryostat consists of a series of concentric cylindrical stages with decreasing temperature toward the center. A schematic of the dilution cryostat is shown in Fig. A.2, while Fig. 4.1 presents a photograph of the actual cryostat used in the CCVR program. The innermost stage is the coldest and is thermally coupled to the experimental region. The $^3\text{He}/^4\text{He}$ mixture circulates through these stages in a closed loop, with cooling taking place in the innermost part, known as the *mixing chamber* (MC). The outer stages serve to precool the mixture before it reaches this region.

The outermost stage of the cryostat, called the *main bath*, is a 4.2 K liquid helium reservoir, which provides the first level of insulation and precooling. The mixture, initially at room temperature, is cooled down by passing through a section thermally coupled to this main bath. It remains gaseous at this stage. To condense it into liquid, the mixture then flows to a *condenser* thermally linked to a secondary reservoir known as the *1K pot*. This pot, continuously refilled from the main bath and pumped from above, stabilizes at around 1.2 K via Eq. (A.1), enabling the condensation of the mixture.

Once liquefied, the mixture continues through heat exchangers that use the cold, outgoing mixture to further cool the incoming one before reaching the mixing chamber. As shown in Fig. A.1, below the critical temperature of $T_c \simeq 870$ mK, the mixture separates into two phases:




Figure A.2: Schematic layout of a $^3\text{He}/^4\text{He}$ dilution cryostat. The closed-loop system includes thermal stages, a condenser, a mixing chamber, a still, and heat exchangers. Reprinted from Ref. [215].

- A concentrated phase of nearly pure ^3He , which forms the upper layer.
- A dilute phase containing ^3He in solution with superfluid ^4He (approximately 6.6% ^3He and 93.4% ^4He).

The boundary between these two phases is called the *phase boundary*. A pump located outside the cryostat establishes a pressure gradient that extracts ^3He from the *still*, a component maintained at 0.6–0.7 K. At this temperature, the vapor pressure of ^3He is roughly three orders of magnitude higher than that of ^4He , so evaporation predominantly removes ^3He . This depletion lowers the ^3He concentration in the dilute phase, establishing an osmotic pressure gradient that drives ^3He atoms from the concentrated phase across the phase boundary into the dilute phase. This endothermic process is the basis of the dilution refrigerator’s cooling mechanism, as it removes heat from the MC. The cryogenic calorimeters are thermalized to the MC.

Temperatures down to a few mK can be achieved using this mechanism, which is sufficient for the operation of cryogenic calorimeters. The lowest attainable temperature depends on the available cooling power and the total heat load on the cryostat. This heat load can originate from imperfect thermal insulation, radiation from warmer stages, and thermal conduction through mechanical supports and wiring. Additional heat can also be injected by the readout electronics, as described by Eq. (2.10). Although radiation from particle interactions injects heat into the system, its net effect is negligible at the event rates typically encountered in low-background searches. The system reaches a condition of dynamic equilibrium when the cooling power balances the incoming heat, thereby defining the base temperature of the cryostat.

The circulation rate of the mixture (typically expressed in $\mu\text{mol} \cdot \text{s}^{-1}$) is a key parameter determining the cryostat’s cooling power. A higher flow rate enables greater heat extraction. The flow can be actively controlled via a resistive heater installed on the still. By increasing the still temperature, additional ^3He is evaporated, enhancing the circulation rate.

Unwanted reductions in flow can occur if impurities freeze at cryogenic temperatures and form blockages in narrow sections of the circuit like the *condenser*. To prevent this, *cold traps* and *getters* are installed to filter out these impurities. Periodic maintenance is required to clean the traps and ensure continued performance.

Some cryostats, known as *dry cryostats*, eliminate the need for liquid helium in the outer stage. Instead, they rely on pulse tube cryocoolers to reach the base temperature of helium. This is the strategy implemented in the CUORE cryostat [137], the same facility where CUPID will be installed. However, these devices operate through periodic piston motion, which introduces mechanical vibrations. Since the experimental region cannot be completely decoupled from the cryostat structure, these vibrations can introduce noise in the detector, worsening the SNR if mitigation strategies are not employed [167, 145, 168]. In contrast, liquid-helium-based refrigerators (*wet* cryostats) provide reduced vibrations during operation in steady-state conditions, offering better noise levels beneficial for cryogenic calorimeters.

Bibliography

- [1] S. Focardi and R. A. Ricci. “Historical Introduction. The Beta-Decay and the Fundamental Properties of Weak Interactions”. In: *La Rivista del Nuovo Cimento (1978-1999)* 6.11 (Nov. 1, 1983), pp. 1–40. ISSN: 1826-9850. DOI: [10.1007/BF02740918](https://doi.org/10.1007/BF02740918). URL: <https://doi.org/10.1007/BF02740918>.
- [2] J Chadwick. “Intensitätsverteilung Im Magnetischen Spectrum Der Beta-Strahlen von Radium B + C”. In: *Verhandl. Dtsch. Phys. Ges.* 16 (1914), p. 383. URL: <https://cds.cern.ch/record/262756>.
- [3] Charles Drummond Ellis and Ernest Rutherford. “Beta-Ray Spectra and Their Meaning”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 101.708 (Jan. 1997), pp. 1–17. DOI: [10.1098/rspa.1922.0022](https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1922.0022). URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1922.0022>.
- [4] Charles Drummond Ellis, W. A. Wooster, and Ernest Rutherford. “The Average Energy of Disintegration of Radium E”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 117.776 (1927), pp. 109–123. DOI: [10.1098/rspa.1927.0168](https://royalsocietypublishing.org/doi/10.1098/rspa.1927.0168). URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1927.0168>.
- [5] Enrico Fermi. “Tentativo di una Teoria Dei Raggi beta”. In: *Il Nuovo Cimento (1924-1942)* 11.1 (Jan. 1, 1934), pp. 1–19. ISSN: 1827-6121. DOI: [10.1007/BF02959820](https://doi.org/10.1007/BF02959820). URL: <https://doi.org/10.1007/BF02959820>.
- [6] F. Reines and C. L. Cowan. “Detection of the Free Neutrino”. In: *Physical Review* 92.3 (Nov. 1, 1953), pp. 830–831. DOI: [10.1103/PhysRev.92.830](https://link.aps.org/doi/10.1103/PhysRev.92.830). URL: <https://link.aps.org/doi/10.1103/PhysRev.92.830>.
- [7] S. Navas et al. “Review of Particle Physics”. In: *Physical Review D* 110.3 (Aug. 1, 2024), p. 030001. DOI: [10.1103/PhysRevD.110.030001](https://link.aps.org/doi/10.1103/PhysRevD.110.030001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.030001>.
- [8] “The Higgs Mechanism in the Glashow–Salam–Weinberg Model”. In: *Electroweak Theory*. Ed. by E. A. Paschos. Cambridge: Cambridge University Press, 2007, pp. 57–63. ISBN: 978-0-511-61104-9. DOI: [10.1017/CB09780511611049.008](https://www.cambridge.org/core/books/electroweak-theory/higgs-mechanism-in-the-glashowsalamweinberg-model/6DD7F43F57E705DAC601FEBB1F8597E7). URL: <https://www.cambridge.org/core/books/electroweak-theory/higgs-mechanism-in-the-glashowsalamweinberg-model/6DD7F43F57E705DAC601FEBB1F8597E7>.
- [9] Peter Skands. “Introduction to QCD”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Searching for New Physics at Small and Large Scales*. 2013, pp. 341–420. DOI: [10.1142/9789814525220_0008](https://doi.org/10.1142/9789814525220_0008).

- [10] Alessandro Bettini. *Introduction to Elementary Particle Physics*. 3rd ed. Cambridge University Press, June 27, 2024. ISBN: 978-1-009-44074-5 978-1-009-44073-8. DOI: [10.1017/9781009440745](https://doi.org/10.1017/9781009440745). URL: <https://doi.org/10.1017/9781009440745>.
- [11] Peter W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Physical Review Letters* 13.16 (Oct. 19, 1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://link.aps.org/doi/10.1103/PhysRevLett.13.508). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [12] P. Achard et al. “Measurement of the Mass and the Width of the W Boson at LEP”. In: *The European Physical Journal C: Particles and Fields* 45 (2006), pp. 569–587. DOI: [10.1140/epjc/s2005-02459-6](https://doi.org/10.1140/epjc/s2005-02459-6).
- [13] L. Arnaudon et al. “Measurement of the Mass of the Z Boson and the Energy Calibration of LEP”. In: *Physics Letters B* 307 (1993), pp. 187–193. DOI: [10.1016/0370-2693\(93\)90210-9](https://doi.org/10.1016/0370-2693(93)90210-9).
- [14] The CMS Collaboration. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 30–61. ISSN: 03702693. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235 \[hep-ex\]](https://arxiv.org/abs/1207.7235). URL: <http://arxiv.org/abs/1207.7235>.
- [15] The ATLAS Collaboration. “Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC”. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. ISSN: 03702693. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214 \[hep-ex\]](https://arxiv.org/abs/1207.7214). URL: <http://arxiv.org/abs/1207.7214>.
- [16] T. D. Lee and C. N. Yang. “Question of Parity Conservation in Weak Interactions”. In: *Physical Review* 104.1 (Oct. 1, 1956), pp. 254–258. DOI: [10.1103/PhysRev.104.254](https://link.aps.org/doi/10.1103/PhysRev.104.254). URL: <https://link.aps.org/doi/10.1103/PhysRev.104.254>.
- [17] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Physical Review* 105.4 (Feb. 15, 1957), pp. 1413–1415. DOI: [10.1103/PhysRev.105.1413](https://link.aps.org/doi/10.1103/PhysRev.105.1413). URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1413>.
- [18] L. Landau. “On the Conservation Laws for Weak Interactions”. In: *Nuclear Physics* 3.1 (Mar. 1, 1957), pp. 127–131. ISSN: 0029-5582. DOI: [10.1016/0029-5582\(57\)90061-5](https://www.sciencedirect.com/science/article/pii/0029558257900615). URL: <https://www.sciencedirect.com/science/article/pii/0029558257900615>.
- [19] T. D. Lee and C. N. Yang. “Parity Nonconservation and a Two-Component Theory of the Neutrino”. In: *Physical Review* 105.5 (Mar. 1, 1957), pp. 1671–1675. DOI: [10.1103/PhysRev.105.1671](https://link.aps.org/doi/10.1103/PhysRev.105.1671). URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1671>.
- [20] Abdus Salam. “On Parity Conservation and Neutrino Mass”. In: *Il Nuovo Cimento (1955-1965)* 5.1 (Jan. 1, 1957), pp. 299–301. ISSN: 1827-6121. DOI: [10.1007/BF02812841](https://doi.org/10.1007/BF02812841). URL: <https://doi.org/10.1007/BF02812841>.
- [21] Matteo Agostini et al. “Toward the Discovery of Matter Creation with Neutrinoless $\beta\beta$ Decay”. In: *Reviews of Modern Physics* 95.2 (May 30, 2023), p. 025002. DOI: [10.1103/RevModPhys.95.025002](https://link.aps.org/doi/10.1103/RevModPhys.95.025002). URL: <https://link.aps.org/doi/10.1103/RevModPhys.95.025002>.
- [22] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. In: *Physical Review* 109.3 (Feb. 1, 1958), pp. 1015–1017. DOI: [10.1103/PhysRev.109.1015](https://link.aps.org/doi/10.1103/PhysRev.109.1015). URL: <https://link.aps.org/doi/10.1103/PhysRev.109.1015>.

- [23] Francesco Vissani. “What Is Matter According to Particle Physics, and Why Try to Observe Its Creation in a Lab?” In: *Universe* 7.3 (Mar. 2021), p. 61. ISSN: 2218-1997. DOI: [10.3390/universe7030061](https://doi.org/10.3390/universe7030061). URL: <https://www.mdpi.com/2218-1997/7/3/61>.
- [24] G. Barr et al. “Weak Interactions”. In: *Particle Physics in the LHC Era*. Ed. by Giles Barr et al. Oxford University Press, Jan. 1, 2016, pp. 182–187. ISBN: 978-0-19-874855-7. DOI: [10.1093/acprof:oso/9780198748557.003.0007](https://doi.org/10.1093/acprof:oso/9780198748557.003.0007). URL: <https://doi.org/10.1093/acprof:oso/9780198748557.003.0007>.
- [25] S. Dell’Oro et al. “Neutrinoless Double Beta Decay: 2015 Review”. In: *Advances in High Energy Physics* 2016 (2016), pp. 1–37. ISSN: 1687-7357, 1687-7365. DOI: [10.1155/2016/2162659](https://doi.org/10.1155/2016/2162659). arXiv: [1601.07512 \[hep-ph\]](https://arxiv.org/abs/1601.07512). URL: <http://arxiv.org/abs/1601.07512>.
- [26] Matthew Dean Schwartz. *Quantum Field Theory and the Standard Model*. New York: Cambridge university press, 2014. ISBN: 978-1-107-03473-0.
- [27] Esteban Roulet. “Neutrinos in Physics and Astrophysics”. In: *From the Sun to the Great Attractor*. Ed. by Dany Page and Jorge G. Hirsch. Red. by R. Beig et al. Vol. 556. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 233–257. ISBN: 978-3-540-41064-5 978-3-540-45371-0. DOI: [10.1007/3-540-45371-7_6](https://doi.org/10.1007/3-540-45371-7_6). URL: <https://doi.org/10.48550/arXiv.hep-ph/9910383>.
- [28] Ian J. R. Aitchison and Anthony J. G. Hey. *Gauge Theories in Particle Physics: A Practical Introduction, Volume 2: Non-Abelian Gauge Theories: QCD and The Electroweak Theory, Fourth Edition*. Taylor & Francis, 2013. ISBN: 978-1-4665-1310-5 978-1-4665-1307-5 978-0-429-18539-7. DOI: [10.1201/9781466513105](https://doi.org/10.1201/9781466513105). URL: <https://library.oapen.org/handle/20.500.12657/50884>.
- [29] Esteban Roulet and F. Vissani. *Neutrinos in Physics and Astrophysics*. 2023rd ed. 1 online resource (xiii, 221 pages) : illustrations (some color) vols. Singapore: World Scientific Publishing Co. Pte. Ltd., 2023. ISBN: 978-981-12-6094-0 978-981-12-6095-7. URL: <https://public.ebookcentral.proquest.com/choice/PublicFullRecord.aspx?p=7132514>.
- [30] G. ’t Hooft. “Symmetry Breaking through Bell-Jackiw Anomalies”. In: *Physical Review Letters* 37.1 (July 5, 1976), pp. 8–11. DOI: [10.1103/PhysRevLett.37.8](https://doi.org/10.1103/PhysRevLett.37.8). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.37.8>.
- [31] Robert J. Noble. “Baryon- and Lepton-Number Violation by Electroweak Instantons”. In: *Physical Review D* 25.3 (Feb. 1, 1982), pp. 825–837. DOI: [10.1103/PhysRevD.25.825](https://doi.org/10.1103/PhysRevD.25.825). URL: <https://link.aps.org/doi/10.1103/PhysRevD.25.825>.
- [32] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov. “On Anomalous Electroweak Baryon-Number Non-Conservation in the Early Universe”. In: *Physics Letters B* 155.1 (May 16, 1985), pp. 36–42. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(85\)91028-7](https://doi.org/10.1016/0370-2693(85)91028-7). URL: <https://www.sciencedirect.com/science/article/pii/0370269385910287>.
- [33] Bruce T. Cleveland et al. “Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector”. In: *The Astrophysical Journal* 496.1 (Mar. 20, 1998), pp. 505–526. ISSN: 0004-637X, 1538-4357. DOI: [10.1086/305343](https://doi.org/10.1086/305343). URL: <https://iopscience.iop.org/article/10.1086/305343>.
- [34] John N. Bahcall and Roger K. Ulrich. “Solar Models, Neutrino Experiments, and Helioseismology”. In: *Reviews of Modern Physics* 60.2 (Apr. 1, 1988), pp. 297–372. DOI: [10.1103/RevModPhys.60.297](https://doi.org/10.1103/RevModPhys.60.297). URL: <https://link.aps.org/doi/10.1103/RevModPhys.60.297>.

- [35] M. Cribier et al. “Results of the Whole GALLEX Experiment”. In: *Nuclear Physics B - Proceedings Supplements*. Proceedings of the Fifth International Workshop on Topics in Astroparticle and Underground Physics 70.1 (Jan. 1, 1999), pp. 284–291. ISSN: 0920-5632. DOI: [10.1016/S0920-5632\(98\)00438-1](https://doi.org/10.1016/S0920-5632(98)00438-1). URL: <https://www.sciencedirect.com/science/article/pii/S0920563298004381>.
- [36] J. N. Abdurashitov et al. “Results from SAGE (The Russian-American Gallium Solar Neutrino Experiment)”. In: *Physics Letters B* 328.1 (May 26, 1994), pp. 234–248. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(94\)90454-5](https://doi.org/10.1016/0370-2693(94)90454-5). URL: <https://www.sciencedirect.com/science/article/pii/0370269394904545>.
- [37] K. S. Hirata et al. “Real-Time, Directional Measurement of B-8 Solar Neutrinos in the Kamiokande II Detector”. In: *Physical Review D* 44.8 (Oct. 15, 1991), pp. 2241–2260. DOI: [10.1103/PhysRevD.44.2241](https://doi.org/10.1103/PhysRevD.44.2241). URL: <https://link.aps.org/doi/10.1103/PhysRevD.44.2241>.
- [38] B. Pontecorvo. “Mesonium and Anti-Mesonium”. In: *Sov. Phys. JETP* 6 (1957), p. 429. URL: <https://inspirehep.net/literature/2884>.
- [39] S. M. Bilenky and B. Pontecorvo. “Lepton Mixing and Neutrino Oscillations”. In: *Physics Reports* 41.4 (May 1, 1978), pp. 225–261. ISSN: 0370-1573. DOI: [10.1016/0370-1573\(78\)90095-9](https://doi.org/10.1016/0370-1573(78)90095-9). URL: <https://www.sciencedirect.com/science/article/pii/0370157378900959>.
- [40] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 87.7 (July 25, 2001), p. 071301. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.87.071301](https://doi.org/10.1103/PhysRevLett.87.071301). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.87.071301>.
- [41] KamLAND Collaboration et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Physical Review Letters* 90.2 (Jan. 17, 2003), p. 021802. DOI: [10.1103/PhysRevLett.90.021802](https://doi.org/10.1103/PhysRevLett.90.021802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.90.021802>.
- [42] Super-Kamiokande Collaboration et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Physical Review Letters* 81.8 (Aug. 24, 1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562>.
- [43] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. ISSN: 0033-068X. DOI: [10.1143/PTP.28.870](https://doi.org/10.1143/PTP.28.870). URL: <https://academic.oup.com/ptp/article-lookup/doi/10.1143/PTP.28.870>.
- [44] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Soviet Journal of Experimental and Theoretical Physics* 26 (May 1968), p. 984. ISSN: 1063-7761. URL: <https://ui.adsabs.harvard.edu/abs/1968JETP...26..984P/abstract>.
- [45] Boris Kayser, Françoise Gibrat-Debu, and Frédéric Perrier. *The Physics of Massive Neutrinos*. 1 online resource (viii, 117 pages) : illustrations vols. World Scientific Lecture Notes in Physics v. 25. Singapore: World Scientific, 1989. ISBN: 978-981-279-925-8.
- [46] Ivan Esteban et al. “NuFit-6.0: Updated Global Analysis of Three-Flavor Neutrino Oscillations”. In: *Journal of High Energy Physics* 2024.12 (Dec. 30, 2024), p. 216. ISSN: 1029-8479. DOI: [10.1007/JHEP12\(2024\)216](https://doi.org/10.1007/JHEP12(2024)216). arXiv: [2410.05380](https://arxiv.org/abs/2410.05380) [hep-ph]. URL: www.nu-fit.org.

- [47] S. M. Bilenky, J. Hošek, and S. T. Petcov. “On the Oscillations of Neutrinos with Dirac and Majorana Masses”. In: *Physics Letters B* 94.4 (Aug. 25, 1980), pp. 495–498. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(80\)90927-2](https://doi.org/10.1016/0370-2693(80)90927-2). URL: <https://www.sciencedirect.com/science/article/pii/0370269380909272>.
- [48] V. Barger, K. Whisnant, and R. J. N. Phillips. “CP Nonconservation in Three-Neutrino Oscillations”. In: *Physical Review Letters* 45.26 (Dec. 29, 1980), pp. 2084–2088. DOI: [10.1103/PhysRevLett.45.2084](https://doi.org/10.1103/PhysRevLett.45.2084). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.45.2084>.
- [49] P. Langacker et al. “Implications of the Mikheyev-Smirnov-Wolfenstein (MSW) Mechanism of Amplification of Neutrino Oscillations in Matter”. In: *Nuclear Physics B* 282 (Jan. 1, 1987), pp. 589–609. ISSN: 0550-3213. DOI: [10.1016/0550-3213\(87\)90699-7](https://doi.org/10.1016/0550-3213(87)90699-7). URL: <https://www.sciencedirect.com/science/article/pii/0550321387906997>.
- [50] Claudio Giganti, Stéphane Lavignac, and Marco Zito. “Neutrino Oscillations: The Rise of the PMNS Paradigm”. In: *Progress in Particle and Nuclear Physics* 98 (Jan. 2018), pp. 1–54. ISSN: 01466410. DOI: [10.1016/j.pnpnp.2017.10.001](https://doi.org/10.1016/j.pnpnp.2017.10.001). arXiv: [1710.00715 \[hep-ex\]](https://arxiv.org/abs/1710.00715). URL: <http://arxiv.org/abs/1710.00715>.
- [51] T2K Collaboration et al. “Observation of Electron Neutrino Appearance in a Muon Neutrino Beam”. In: *Physical Review Letters* 112.6 (Feb. 10, 2014), p. 061802. DOI: [10.1103/PhysRevLett.112.061802](https://doi.org/10.1103/PhysRevLett.112.061802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.061802>.
- [52] NOvA Collaboration et al. “First Measurement of Electron Neutrino Appearance in NOvA”. In: *Physical Review Letters* 116.15 (Apr. 13, 2016), p. 151806. DOI: [10.1103/PhysRevLett.116.151806](https://doi.org/10.1103/PhysRevLett.116.151806). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.151806>.
- [53] N. Agafonova et al. “Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment”. In: *Physical Review Letters* 115.12 (Sept. 17, 2015), p. 121802. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.115.121802](https://doi.org/10.1103/PhysRevLett.115.121802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.121802>.
- [54] Z. Li et al. “Measurement of the Tau Neutrino Cross Section in Atmospheric Neutrino Oscillations with Super-Kamiokande”. In: *Physical Review D* 98.5 (Sept. 13, 2018), p. 052006. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.98.052006](https://doi.org/10.1103/PhysRevD.98.052006). URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.052006>.
- [55] P. F. de Salas et al. “Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects”. In: *Frontiers in Astronomy and Space Sciences* 5 (Oct. 9, 2018), p. 36. ISSN: 2296-987X. DOI: [10.3389/fspas.2018.00036](https://doi.org/10.3389/fspas.2018.00036). arXiv: [1806.11051 \[hep-ph\]](https://arxiv.org/abs/1806.11051). URL: <http://arxiv.org/abs/1806.11051>.
- [56] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Physical Review D* 17.9 (May 1, 1978), pp. 2369–2374. DOI: [10.1103/PhysRevD.17.2369](https://doi.org/10.1103/PhysRevD.17.2369). URL: <https://link.aps.org/doi/10.1103/PhysRevD.17.2369>.
- [57] S. P. Mikheyev and A. Yu. Smirnov. “Resonant Amplification of ν Oscillations in Matter and Solar-Neutrino Spectroscopy”. In: *Il Nuovo Cimento C* 9.1 (Jan. 1, 1986), pp. 17–26. ISSN: 0390-5551. DOI: [10.1007/BF02508049](https://doi.org/10.1007/BF02508049). URL: <https://doi.org/10.1007/BF02508049>.

- [58] Angel Abusleme et al. “Potential to Identify Neutrino Mass Ordering with Reactor Antineutrinos at JUNO*”. In: *Chinese Physics C* 49.3 (Mar. 2025), p. 033104. ISSN: 1674-1137. DOI: [10.1088/1674-1137/ad7f3e](https://doi.org/10.1088/1674-1137/ad7f3e). URL: <https://dx.doi.org/10.1088/1674-1137/ad7f3e>.
- [59] K. Abe et al. *Hyper-Kamiokande Design Report*. Nov. 28, 2018. DOI: [10.48550/arXiv.1805.04163](https://doi.org/10.48550/arXiv.1805.04163). arXiv: [1805.04163](https://arxiv.org/abs/1805.04163) [physics]. URL: <http://arxiv.org/abs/1805.04163>. Pre-published.
- [60] B. Abi et al. “Long-Baseline Neutrino Oscillation Physics Potential of the DUNE Experiment”. In: *The European Physical Journal C* 80.10 (Oct. 22, 2020), p. 978. ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-020-08456-z](https://doi.org/10.1140/epjc/s10052-020-08456-z). URL: <https://doi.org/10.1140/epjc/s10052-020-08456-z>.
- [61] Ettore Majorana. “Teoria simmetrica dell’elettrone e del positrone”. In: *Il Nuovo Cimento (1924-1942)* 14.4 (Apr. 1, 1937), pp. 171–184. ISSN: 1827-6121. DOI: [10.1007/BF02961314](https://doi.org/10.1007/BF02961314). URL: <https://doi.org/10.1007/BF02961314>.
- [62] S. F. King. “Neutrino Mass Models”. In: *Reports on Progress in Physics* 67.2 (Dec. 2003), p. 107. ISSN: 0034-4885. DOI: [10.1088/0034-4885/67/2/R01](https://dx.doi.org/10.1088/0034-4885/67/2/R01). URL: <https://dx.doi.org/10.1088/0034-4885/67/2/R01>.
- [63] KATRIN COLLABORATION et al. “Direct Neutrino-Mass Measurement Based on 259 Days of KATRIN Data”. In: *Science* 388.6743 (Apr. 11, 2025), pp. 180–185. DOI: [10.1126/science.adq9592](https://doi.org/10.1126/science.adq9592). URL: <https://www.science.org/doi/10.1126/science.adq9592>.
- [64] L. Gastaldo et al. “The Electron Capture ^{163}Ho Experiment ECHO”. In: *Journal of Low Temperature Physics* 176.5 (Sept. 1, 2014), pp. 876–884. ISSN: 1573-7357. DOI: [10.1007/s10909-014-1187-4](https://doi.org/10.1007/s10909-014-1187-4). URL: <https://doi.org/10.1007/s10909-014-1187-4>.
- [65] B. Alpert et al. “HOLMES”. In: *The European Physical Journal C* 75.3 (Mar. 10, 2015), p. 112. ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-015-3329-5](https://doi.org/10.1140/epjc/s10052-015-3329-5). URL: <https://doi.org/10.1140/epjc/s10052-015-3329-5>.
- [66] A. De Rújula and M. Lusignoli. “Calorimetric Measurements of $^{163}\text{holmium}$ Decay as Tools to Determine the Electron Neutrino Mass”. In: *Physics Letters B* 118.4 (Dec. 9, 1982), pp. 429–434. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(82\)90218-0](https://doi.org/10.1016/0370-2693(82)90218-0). URL: <https://www.sciencedirect.com/science/article/pii/0370269382902180>.
- [67] DESI Collaboration et al. “DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations”. In: *Journal of Cosmology and Astroparticle Physics* 2025.02 (Feb. 1, 2025), p. 021. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2025/02/021](https://doi.org/10.1088/1475-7516/2025/02/021). arXiv: [2404.03002](https://arxiv.org/abs/2404.03002) [astro-ph]. URL: <http://arxiv.org/abs/2404.03002>.
- [68] Daniel Naredo-Tuero et al. *Living at the Edge: A Critical Look at the Cosmological Neutrino Mass Bound*. Oct. 28, 2024. DOI: [10.48550/arXiv.2407.13831](https://doi.org/10.48550/arXiv.2407.13831). arXiv: [2407.13831](https://arxiv.org/abs/2407.13831) [astro-ph]. URL: <http://arxiv.org/abs/2407.13831>. Pre-published.
- [69] Jun-Qian Jiang et al. “Neutrino Cosmology after DESI: Tightest Mass Upper Limits, Preference for the Normal Ordering, and Tension with Terrestrial Observations”. In: *Journal of Cosmology and Astroparticle Physics* 2025.01 (Jan. 1, 2025), p. 153. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2025/01/153](https://doi.org/10.1088/1475-7516/2025/01/153). arXiv: [2407.18047](https://arxiv.org/abs/2407.18047) [astro-ph]. URL: <http://arxiv.org/abs/2407.18047>.

- [70] M. Goepfert-Mayer. “Double Beta-Disintegration”. In: *Physical Review* 48.6 (Sept. 15, 1935), pp. 512–516. DOI: [10.1103/PhysRev.48.512](https://doi.org/10.1103/PhysRev.48.512). URL: <https://link.aps.org/doi/10.1103/PhysRev.48.512>.
- [71] B. Pritychenko and V.I. Tretyak. “Comprehensive Review of 2β Decay Half-Lives”. In: *Atomic Data and Nuclear Data Tables* 161 (Jan. 2025), p. 101694. ISSN: 0092-640X. DOI: [10.1016/j.adt.2024.101694](https://doi.org/10.1016/j.adt.2024.101694). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0092640X24000597>.
- [72] Eugene Greuling and R. C Whitten. “Lepton Conservation and Double Beta-Decay”. In: *Annals of Physics* 11.4 (Dec. 1, 1960), pp. 510–533. ISSN: 0003-4916. DOI: [10.1016/0003-4916\(60\)90010-5](https://doi.org/10.1016/0003-4916(60)90010-5). URL: <https://www.sciencedirect.com/science/article/pii/S0003491660900105>.
- [73] Juris Meija et al. “Isotopic Compositions of the Elements 2013 (IUPAC Technical Report)”. In: *Pure and Applied Chemistry* 88.3 (Mar. 1, 2016), pp. 293–306. ISSN: 1365-3075. DOI: [10.1515/pac-2015-0503](https://doi.org/10.1515/pac-2015-0503). URL: <https://www.degruyterbrill.com/document/doi/10.1515/pac-2015-0503/html>.
- [74] Planck Collaboration et al. “Planck 2018 Results: VI. Cosmological Parameters”. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A6. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910). URL: <https://www.aanda.org/10.1051/0004-6361/201833910>.
- [75] Ofelia Pisanti. “Improved Nuclear Reaction Network for a Reliable Estimate of Primordial Deuterium Yield”. In: *Journal of Physics: Conference Series* 1468.1 (Feb. 2020), p. 012010. ISSN: 1742-6596. DOI: [10.1088/1742-6596/1468/1/012010](https://doi.org/10.1088/1742-6596/1468/1/012010). URL: <https://dx.doi.org/10.1088/1742-6596/1468/1/012010>.
- [76] Andrei D. Sakharov. “Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe”. In: *Soviet Physics Uspekhi* 34.5 (May 31, 1991), p. 392. ISSN: 0038-5670. DOI: [10.1070/PU1991v034n05ABEH002497](https://doi.org/10.1070/PU1991v034n05ABEH002497). URL: <https://iopscience.iop.org/article/10.1070/PU1991v034n05ABEH002497/meta>.
- [77] A.i. Bochkarev and M.e. Shaposhnikov. “Electroweak Production of Baryon Asymmetry and Upper Bounds on the Higgs and Top Masses”. In: *Modern Physics Letters A* 02.06 (June 1987), pp. 417–427. ISSN: 0217-7323. DOI: [10.1142/S0217732387000537](https://doi.org/10.1142/S0217732387000537). URL: <https://www.worldscientific.com/doi/abs/10.1142/S0217732387000537>.
- [78] M. Fukugita and T. Yanagida. “Baryogenesis without Grand Unification”. In: *Physics Letters B* 174.1 (June 26, 1986), pp. 45–47. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(86\)91126-3](https://doi.org/10.1016/0370-2693(86)91126-3). URL: <https://www.sciencedirect.com/science/article/pii/0370269386911263>.
- [79] J. Schechter and J. W. F. Valle. “Neutrinoless Double- β Decay in $SU(2)\times U(1)$ Theories”. In: *Physical Review D* 25.11 (June 1, 1982), pp. 2951–2954. DOI: [10.1103/PhysRevD.25.2951](https://doi.org/10.1103/PhysRevD.25.2951). URL: <https://link.aps.org/doi/10.1103/PhysRevD.25.2951>.
- [80] Michael Duerr, Manfred Lindner, and Alexander Merle. “On the Quantitative Impact of the Schechter-Valle Theorem”. In: *Journal of High Energy Physics* 2011.6 (June 21, 2011), p. 91. ISSN: 1029-8479. DOI: [10.1007/JHEP06\(2011\)091](https://doi.org/10.1007/JHEP06(2011)091). URL: [https://doi.org/10.1007/JHEP06\(2011\)091](https://doi.org/10.1007/JHEP06(2011)091).

- [81] Manimala Mitra, Goran Senjanović, and Francesco Vissani. “Neutrinoless Double Beta Decay and Heavy Sterile Neutrinos”. In: *Nuclear Physics B* 856.1 (Mar. 1, 2012), pp. 26–73. ISSN: 0550-3213. DOI: [10.1016/j.nuclphysb.2011.10.035](https://doi.org/10.1016/j.nuclphysb.2011.10.035). URL: <https://www.sciencedirect.com/science/article/pii/S0550321311006262>.
- [82] Gia Dvali et al. “Neutrinoless Double Beta Decay: Neutrino Mass versus New Physics”. In: *Physical Review D* 108.7 (Oct. 11, 2023), p. 075012. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.108.075012](https://doi.org/10.1103/PhysRevD.108.075012). URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.075012>.
- [83] Fedor Šimkovic, John Vergados, and Amand Faessler. “Few Active Mechanisms of the $0\nu\beta\beta$ Decay and Effective Mass of Majorana Neutrinos”. In: *Physical Review D* 82.11 (Dec. 23, 2010), p. 113015. DOI: [10.1103/PhysRevD.82.113015](https://doi.org/10.1103/PhysRevD.82.113015). URL: <https://link.aps.org/doi/10.1103/PhysRevD.82.113015>.
- [84] S. T. Petcov. “The Nature of Massive Neutrinos”. In: *Advances in High Energy Physics* 2013.1 (2013), p. 852987. ISSN: 1687-7365. DOI: [10.1155/2013/852987](https://doi.org/10.1155/2013/852987). URL: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2013/852987>.
- [85] Ferruccio Feruglio and Andrea Romanino. “Lepton Flavor Symmetries”. In: *Reviews of Modern Physics* 93.1 (Mar. 17, 2021), p. 015007. DOI: [10.1103/RevModPhys.93.015007](https://doi.org/10.1103/RevModPhys.93.015007). URL: <https://link.aps.org/doi/10.1103/RevModPhys.93.015007>.
- [86] Evgeny Akhmedov. *Non-Relativistic Neutrinos and the Question of Dirac vs. Majorana Neutrino Nature*. Dec. 23, 2024. DOI: [10.48550/arXiv.2410.11940](https://doi.org/10.48550/arXiv.2410.11940). arXiv: [2410.11940](https://arxiv.org/abs/2410.11940) [hep-ph]. URL: <http://arxiv.org/abs/2410.11940>. Pre-published.
- [87] NuPECC. *NuPECC Long Range Plan 2024 for European Nuclear Physics*. June 6, 2025. DOI: [10.48550/arXiv.2503.15575](https://doi.org/10.48550/arXiv.2503.15575). arXiv: [2503.15575](https://arxiv.org/abs/2503.15575) [nucl-ex]. URL: <http://arxiv.org/abs/2503.15575>. Pre-published.
- [88] Gail E. Dodge. “The U.S. Nuclear Science Long Range Plan”. In: *Nuclear Physics News* 34.1 (Jan. 2, 2024), pp. 3–4. ISSN: 1061-9127. DOI: [10.1080/10619127.2024.2303306](https://doi.org/10.1080/10619127.2024.2303306). URL: <https://doi.org/10.1080/10619127.2024.2303306>.
- [89] Fang Dong-Liang. “Neutrinoless double-beta decay and new physics beyond the Standard Model”. In: *PHYSICS* 53.5 (May 15, 2024), pp. 310–316. ISSN: 0379-4148. DOI: [10.7693/wl20240503](https://doi.org/10.7693/wl20240503). URL: <https://wuli.iphys.ac.cn/en/article/doi/10.7693/wl20240503.pdf>.
- [90] Oliviero Cremonesi and Maura Pavan. “Challenges in Double Beta Decay”. In: *Advances in High Energy Physics* 2014.1 (2014), p. 951432. ISSN: 1687-7365. DOI: [10.1155/2014/951432](https://doi.org/10.1155/2014/951432). URL: <https://onlinelibrary.wiley.com/doi/abs/10.1155/2014/951432>.
- [91] J. Kotila and F. Iachello. “Phase-Space Factors for Double- β Decay”. In: *Physical Review C* 85.3 (Mar. 19, 2012), p. 034316. DOI: [10.1103/PhysRevC.85.034316](https://doi.org/10.1103/PhysRevC.85.034316). URL: <https://link.aps.org/doi/10.1103/PhysRevC.85.034316>.
- [92] M. Aker et al. “Direct Neutrino-Mass Measurement with Sub-Electronvolt Sensitivity”. In: *Nature Physics* 18.2 (Feb. 2022), pp. 160–166. ISSN: 1745-2481. DOI: [10.1038/s41567-021-01463-1](https://doi.org/10.1038/s41567-021-01463-1). URL: <https://www.nature.com/articles/s41567-021-01463-1>.

- [93] KATRIN Collaboration et al. “Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN”. In: *Physical Review Letters* 123.22 (Nov. 25, 2019), p. 221802. DOI: [10.1103/PhysRevLett.123.221802](https://doi.org/10.1103/PhysRevLett.123.221802). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.221802>.
- [94] KamLAND-Zen Collaboration et al. “Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen”. In: *Physical Review Letters* 117.8 (Aug. 16, 2016), p. 082503. DOI: [10.1103/PhysRevLett.117.082503](https://doi.org/10.1103/PhysRevLett.117.082503). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.117.082503>.
- [95] Jonathan Engel and Javier Menéndez. “Status and Future of Nuclear Matrix Elements for Neutrinoless Double-Beta Decay: A Review”. In: *Reports on Progress in Physics* 80.4 (Apr. 1, 2017), p. 046301. ISSN: 0034-4885, 1361-6633. DOI: [10.1088/1361-6633/aa5bc5](https://doi.org/10.1088/1361-6633/aa5bc5). arXiv: [1610.06548](https://arxiv.org/abs/1610.06548) [nucl-th]. URL: <http://arxiv.org/abs/1610.06548>.
- [96] S. Abe et al. “Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset”. In: *Physical Review Letters* 135.26 (Dec. 29, 2025), p. 262501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/jkfk6-48j8](https://doi.org/10.1103/jkfk6-48j8). URL: <https://link.aps.org/doi/10.1103/jkfk6-48j8>.
- [97] M. Biassoni. “A New Way of Comparing DBD Experiments”. In: *Nuclear and Particle Physics Proceedings* 265–266 (Aug. 2015), pp. 34–37. ISSN: 24056014. DOI: [10.1016/j.nuclphysbps.2015.06.009](https://doi.org/10.1016/j.nuclphysbps.2015.06.009). URL: <https://linkinghub.elsevier.com/retrieve/pii/S240560141500351X>.
- [98] Curie Marie. “Notes Nature March 26 1903”. In: *Nature* 67.1743 (Mar. 1903), pp. 491–496. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/067491b0](https://doi.org/10.1038/067491b0). URL: <https://www.nature.com/articles/067491b0>.
- [99] E. Rutherford and H. T. Barnes. “Heating Effect of the Radium Emanation”. In: *Nature* 68.1774 (Oct. 29, 1903), pp. 622–622. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/068622a0](https://doi.org/10.1038/068622a0). URL: <https://www.nature.com/articles/068622a0>.
- [100] T.O. Niinikoski and F Udo. “Cryogenic Detection of Neutrinos?” In: *NP Internal Report 74-6* (Sept. 2, 1974). URL: <https://cds.cern.ch/record/200440/files/cer-000112089.pdf>.
- [101] G.V. Micelmacher, B.S. Neganov, and V.N. Trofimov. “A New Method of Neutrinoless Double-beta Decay Registration”. In: *JINR preprint P8-82-549, Dubna* (1982).
- [102] E. Fiorini and T.O. Niinikoski. “Low-Temperature Calorimetry for Rare Decays”. In: *Nuclear Instruments and Methods in Physics Research* 224.1–2 (July 1984), pp. 83–88. ISSN: 01675087. DOI: [10.1016/0167-5087\(84\)90449-6](https://doi.org/10.1016/0167-5087(84)90449-6). URL: <https://linkinghub.elsevier.com/retrieve/pii/0167508784904496>.
- [103] D.Q. Adams et al. “CUORE Opens the Door to Tonne-Scale Cryogenics Experiments”. In: *Progress in Particle and Nuclear Physics* 122 (Jan. 2022), p. 103902. ISSN: 01466410. DOI: [10.1016/j.pnnp.2021.103902](https://doi.org/10.1016/j.pnnp.2021.103902). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0146641021000612>.
- [104] D. Q. Adams et al. “Search for Majorana Neutrinos Exploiting Millikelvin Cryogenics with CUORE”. In: *Nature* 604.7904 (Apr. 1, 2022), pp. 53–58. ISSN: 1476-4687. DOI: [10.1038/s41586-022-04497-4](https://doi.org/10.1038/s41586-022-04497-4). URL: <https://doi.org/10.1038/s41586-022-04497-4>.
- [105] Charles Kittel. *Introduction to Solid-State Physics*. 5. ed. New York: Wiley, 1976. 608 pp. ISBN: 978-0-471-49024-1.

- [106] D. Alloni et al. “New Dark Matter Direct Search Based on Archaeological Pb”. In: *Physical Review D* 111.10 (May 30, 2025), p. 103050. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/wc vd-rk1f](https://doi.org/10.1103/wc vd-rk1f). URL: <https://link.aps.org/doi/10.1103/wc vd-rk1f>.
- [107] 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](https://doi.org/10.1103/PhysRevD.100.102002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.102002>.
- [108] G. Angloher et al. “Deep-Underground Dark Matter Search with a COSINUS Detector Prototype”. In: *Physical Review D* 110.4 (Aug. 5, 2024), p. 043010. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.110.043010](https://doi.org/10.1103/PhysRevD.110.043010). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.043010>.
- [109] Denys Poda. “Scintillation in Low-Temperature Particle Detectors”. In: *Physics* 3.3 (July 1, 2021), pp. 473–535. ISSN: 2624-8174. DOI: [10.3390/physics3030032](https://doi.org/10.3390/physics3030032). URL: <https://www.mdpi.com/2624-8174/3/3/32>.
- [110] M. Barucci et al. “Measurement of Low Temperature Specific Heat of Crystalline TeO₂ for the Optimization of Bolometric Detectors”. In: *Journal of Low Temperature Physics* 123.5–6 (June 2001), pp. 303–314. ISSN: 0022-2291, 1573-7357. DOI: [10.1023/A:1017555615150](https://doi.org/10.1023/A:1017555615150). URL: <https://link.springer.com/10.1023/A:1017555615150>.
- [111] A.E. Musikhin et al. “The Heat Capacity of Li₂MoO₄ in the Temperature Range 6–310 K”. In: *Journal of Alloys and Compounds* 639 (Aug. 2015), pp. 145–148. ISSN: 09258388. DOI: [10.1016/j.jallcom.2015.03.159](https://doi.org/10.1016/j.jallcom.2015.03.159). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925838815008816>.
- [112] D.Q. Adams et al. “An Energy-Dependent Electro-Thermal Response Model of CUORE Cryogenic Calorimeter”. In: *Journal of Instrumentation* 17.11 (Nov. 1, 2022), P11023. ISSN: 1748-0221. DOI: [10.1088/1748-0221/17/11/P11023](https://doi.org/10.1088/1748-0221/17/11/P11023). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/17/11/P11023>.
- [113] A. Alessandrello et al. “Signal Modelling for TeO₂ Bolometric Detectors”. In: *Journal of Low Temperature Physics* 93.3–4 (Nov. 1993), pp. 207–212. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/BF00693421](https://doi.org/10.1007/BF00693421). URL: <http://link.springer.com/10.1007/BF00693421>.
- [114] A. Alessandrello et al. “An Electrothermal Model for Large Mass Bolometric Detectors”. In: *IEEE Transactions on Nuclear Science* 40.4 (Aug. 1993), pp. 649–656. ISSN: 0018-9499, 1558-1578. DOI: [10.1109/23.256634](https://doi.org/10.1109/23.256634). URL: <https://ieeexplore.ieee.org/document/256634/>.
- [115] I. Nutini et al. “Latest Results from the CUORE Experiment”. In: *Journal of Low Temperature Physics* 209.5–6 (Dec. 2022), pp. 927–935. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-022-02873-y](https://doi.org/10.1007/s10909-022-02873-y). URL: <https://link.springer.com/10.1007/s10909-022-02873-y>.
- [116] CUORE Collaboration et al. “Constraints on Lepton Number Violation with the 2 Tonne · Year CUORE Dataset”. In: *Science* (Oct. 16, 2025), eadp6474. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.adp6474](https://doi.org/10.1126/science.adp6474). URL: <https://www.science.org/doi/10.1126/science.adp6474>.

- [117] C. Arnaboldi et al. “A Front-End Electronic System for Large Arrays of Bolometers”. In: *Journal of Instrumentation* 13.02 (Feb. 23, 2018), P02026–P02026. ISSN: 1748-0221. DOI: [10.1088/1748-0221/13/02/P02026](https://doi.org/10.1088/1748-0221/13/02/P02026). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/13/02/P02026>.
- [118] Mario De Lucia et al. “Transition Edge Sensors: Physics and Applications”. In: *Instruments* 8.4 (Oct. 31, 2024), p. 47. ISSN: 2410-390X. DOI: [10.3390/instruments8040047](https://doi.org/10.3390/instruments8040047). URL: <https://www.mdpi.com/2410-390X/8/4/47>.
- [119] S. Kempf et al. “Physics and Applications of Metallic Magnetic Calorimeters”. In: *Journal of Low Temperature Physics* 193.3–4 (Nov. 2018), pp. 365–379. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-018-1891-6](https://doi.org/10.1007/s10909-018-1891-6). URL: <http://link.springer.com/10.1007/s10909-018-1891-6>.
- [120] Jonas Zmuidzinias. “Superconducting Microresonators: Physics and Applications”. In: *Annual Review of Condensed Matter Physics* 3.1 (Mar. 1, 2012), pp. 169–214. ISSN: 1947-5454, 1947-5462. DOI: [10.1146/annurev-conmatphys-020911-125022](https://doi.org/10.1146/annurev-conmatphys-020911-125022). URL: <https://www.annualreviews.org/doi/10.1146/annurev-conmatphys-020911-125022>.
- [121] V. Zema et al. “Description and Performance of the COSINUS remoTES Design”. In: *Journal of Low Temperature Physics* 217.3–4 (Nov. 2024), pp. 393–400. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-024-03201-2](https://doi.org/10.1007/s10909-024-03201-2). URL: <https://link.springer.com/10.1007/s10909-024-03201-2>.
- [122] Allen Miller and Elihu Abrahams. “Impurity Conduction at Low Concentrations”. In: *Physical Review* 120.3 (Nov. 1, 1960), pp. 745–755. ISSN: 0031-899X. DOI: [10.1103/PhysRev.120.745](https://doi.org/10.1103/PhysRev.120.745). URL: <https://link.aps.org/doi/10.1103/PhysRev.120.745>.
- [123] S. Luryi. *Electronic Properties of Doped Semiconductors*. Springer Series in Solid-State Sciences Ser v.45. Berlin, Heidelberg: Springer Berlin / Heidelberg, 1984. 1 p. ISBN: 978-3-540-12995-0 978-3-662-02403-4.
- [124] K. Alfonso et al. “An Automated System to Define the Optimal Operating Settings of Cryogenic Calorimeters”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1008 (Aug. 2021), p. 165451. ISSN: 01689002. DOI: [10.1016/j.nima.2021.165451](https://doi.org/10.1016/j.nima.2021.165451). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900221004368>.
- [125] Davide Trotta. “The CUPID Neutrinoless Double-Beta Decay Experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1066 (Sept. 2024), p. 169657. ISSN: 0168-9002. DOI: [10.1016/j.nima.2024.169657](https://doi.org/10.1016/j.nima.2024.169657). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900224005837>.
- [126] C. Enss and D. McCammon. “Physical Principles of Low Temperature Detectors: Ultimate Performance Limits and Current Detector Capabilities”. In: *Journal of Low Temperature Physics* 151.1–2 (Apr. 2008), pp. 5–24. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-007-9611-7](https://doi.org/10.1007/s10909-007-9611-7). URL: <http://link.springer.com/10.1007/s10909-007-9611-7>.
- [127] S. H. Moseley, J. C. Mather, and D. McCammon. “Thermal Detectors as X-Ray Spectrometers”. In: *Journal of Applied Physics* 56.5 (Sept. 1, 1984), pp. 1257–1262. ISSN: 0021-8979, 1089-7550. DOI: [10.1063/1.334129](https://doi.org/10.1063/1.334129). URL: <https://pubs.aip.org/jap/article/56/5/1257/170683/Thermal-detectors-as-x-ray-spectrometers>.

- [128] D. McCammon. “Thermal Equilibrium Calorimeters – An Introduction”. In: *Cryogenic Particle Detection*. Ed. by Christian Enss. Red. by Claus E. Ascheron, Hans J. Kölsch, and Werner Skolaut. Vol. 99. Berlin, Heidelberg: Springer Berlin Heidelberg, July 9, 2005, pp. 1–34. ISBN: 978-3-540-20113-7 978-3-540-31478-3. DOI: [10.1007/10933596_1](https://doi.org/10.1007/10933596_1). URL: http://link.springer.com/10.1007/10933596_1.
- [129] Norman E. Booth, Blas Cabrera, and Ettore Fiorini. “LOW-TEMPERATURE PARTICLE DETECTORS”. In: *Annual Review of Nuclear and Particle Science* 46.1 (Dec. 1996), pp. 471–532. ISSN: 0163-8998, 1545-4134. DOI: [10.1146/annurev.nucl.46.1.471](https://doi.org/10.1146/annurev.nucl.46.1.471). URL: <https://www.annualreviews.org/doi/10.1146/annurev.nucl.46.1.471>.
- [130] A. Nucciotti. “The Use of Low Temperature Detectors for Direct Measurements of the Mass of the Electron Neutrino”. In: *Advances in High Energy Physics* 2016 (2016), pp. 1–41. ISSN: 1687-7357, 1687-7365. DOI: [10.1155/2016/9153024](https://doi.org/10.1155/2016/9153024). URL: <http://www.hindawi.com/journals/ahep/2016/9153024/>.
- [131] G Pessina. “Room Temperature Differential Voltage Sensitive Preamplifier for Large Mass Bolometric Detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 444.1–2 (Apr. 2000), pp. 132–135. ISSN: 0168-9002. DOI: [10.1016/S0168-9002\(99\)01345-5](https://doi.org/10.1016/S0168-9002(99)01345-5). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900299013455>.
- [132] A Alessandrello et al. “A Programmable Front-End System for Arrays of Bolometers”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 444.1–2 (Apr. 2000), pp. 111–114. ISSN: 0168-9002. DOI: [10.1016/S0168-9002\(99\)01340-6](https://doi.org/10.1016/S0168-9002(99)01340-6). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900299013406>.
- [133] J B Birks. “Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations”. In: *Proceedings of the Physical Society. Section A* 64.10 (Oct. 1, 1951), pp. 874–877. ISSN: 0370-1298. DOI: [10.1088/0370-1298/64/10/303](https://doi.org/10.1088/0370-1298/64/10/303). URL: <https://iopscience.iop.org/article/10.1088/0370-1298/64/10/303>.
- [134] D. Q. Adams et al. *Exploring the keV-scale Physics Potential of CUORE*. May 29, 2025. DOI: [10.48550/arXiv.2505.23955](https://doi.org/10.48550/arXiv.2505.23955). arXiv: [2505.23955](https://arxiv.org/abs/2505.23955) [hep-ex]. URL: <http://arxiv.org/abs/2505.23955>. Pre-published.
- [135] D. Q. Adams et al. “Data-Driven Background Model for the CUORE Experiment”. In: *Physical Review D* 110.5 (Sept. 11, 2024), p. 052003. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.110.052003](https://doi.org/10.1103/PhysRevD.110.052003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.052003>.
- [136] S. Rahaman et al. “Double-Beta Decay Q Values of ^{116}Cd and ^{130}Te ”. In: *Physics Letters B* 703.4 (Sept. 2011), pp. 412–416. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2011.07.078](https://doi.org/10.1016/j.physletb.2011.07.078). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0370269311008975>.
- [137] C. Alduino et al. “The CUORE Cryostat: An Infrastructure for Rare Event Searches at Millikelvin Temperatures”. In: *Cryogenics* 102 (Sept. 2019), pp. 9–21. ISSN: 0011-2275. DOI: [10.1016/j.cryogenics.2019.06.011](https://doi.org/10.1016/j.cryogenics.2019.06.011). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0011227519301031>.

- [138] G Bellini et al. “Cosmic-Muon Flux and Annual Modulation in Borexino at 3800 m Water-Equivalent Depth”. In: *Journal of Cosmology and Astroparticle Physics* 2012.05 (May 15, 2012), pp. 015–015. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2012/05/015](https://doi.org/10.1088/1475-7516/2012/05/015). URL: <https://iopscience.iop.org/article/10.1088/1475-7516/2012/05/015>.
- [139] H. Wulandari et al. “Neutron Flux at the Gran Sasso Underground Laboratory Revisited”. In: *Astroparticle Physics* 22.3–4 (Nov. 2004), pp. 313–322. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2004.07.005](https://doi.org/10.1016/j.astropartphys.2004.07.005). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650504001276>.
- [140] A. Alessandrello et al. “Measurements on Radioactivity of Ancient Roman Lead to Be Used as Shield in Searches for Rare Events”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 61.1 (July 1991), pp. 106–117. ISSN: 0168-583X. DOI: [10.1016/0168-583x\(91\)95569-y](https://doi.org/10.1016/0168-583x(91)95569-y). URL: <https://linkinghub.elsevier.com/retrieve/pii/0168583X9195569Y>.
- [141] F. Alessandria et al. “Validation of Techniques to Mitigate Copper Surface Contamination in CUORE”. In: *Astroparticle Physics* 45 (May 2013), pp. 13–22. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2013.02.005](https://doi.org/10.1016/j.astropartphys.2013.02.005). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650513000443>.
- [142] E. Buccheri et al. “An Assembly Line for the Construction of Ultra-Radio-Pure Detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 768 (Dec. 2014), pp. 130–140. ISSN: 0168-9002. DOI: [10.1016/j.nima.2014.09.046](https://doi.org/10.1016/j.nima.2014.09.046). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900214010742>.
- [143] G. Benato et al. “Radon Mitigation during the Installation of the CUORE $0\nu\beta\beta$ Decay Detector”. In: *Journal of Instrumentation* 13.01 (Jan. 12, 2018), P01010–P01010. ISSN: 1748-0221. DOI: [10.1088/1748-0221/13/01/p01010](https://doi.org/10.1088/1748-0221/13/01/p01010). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/13/01/P01010>.
- [144] Antonio Branca. “Performance of the Low Threshold Optimum Trigger on CUORE Data”. In: *Journal of Physics: Conference Series* 1468.1 (Feb. 1, 2020), p. 012118. ISSN: 1742-6588, 1742-6596. DOI: [10.1088/1742-6596/1468/1/012118](https://doi.org/10.1088/1742-6596/1468/1/012118). URL: <https://iopscience.iop.org/article/10.1088/1742-6596/1468/1/012118>.
- [145] Kenneth J. Vetter et al. *Improving the Performance of Cryogenic Calorimeters with Non-linear Multivariate Noise Cancellation Algorithms*. Version 2. 2023. DOI: [10.48550/ARXIV.2311.01131](https://doi.org/10.48550/ARXIV.2311.01131). URL: <https://arxiv.org/abs/2311.01131>. Pre-published.
- [146] E. Gatti and P. F. Manfredi. “Processing the Signals from Solid-State Detectors in Elementary-Particle Physics”. In: *La Rivista del Nuovo Cimento* 9.1 (Jan. 1986), pp. 1–146. ISSN: 1826-9850. DOI: [10.1007/bf02822156](https://doi.org/10.1007/bf02822156). URL: <http://link.springer.com/10.1007/BF02822156>.
- [147] D. Q. Adams et al. *End-to-End Data Analysis Methods for the CUORE Experiment*. Version 1. 2025. DOI: [10.48550/ARXIV.2510.25720](https://doi.org/10.48550/ARXIV.2510.25720). URL: <https://arxiv.org/abs/2510.25720>. Pre-published.
- [148] M. Agostini et al. “Final Results of GERDA on the Search for Neutrinoless Double- β Decay”. In: *Physical Review Letters* 125.25 (Dec. 17, 2020). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/physrevlett.125.252502](https://doi.org/10.1103/physrevlett.125.252502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.252502>.

- [149] O. Azzolini et al. “Final Result on the Neutrinoless Double Beta Decay of Se82 with CUPID-0”. In: *Physical Review Letters* 129.11 (Sept. 6, 2022). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.129.111801](https://doi.org/10.1103/PhysRevLett.129.111801). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.111801>.
- [150] C. Augier et al. “Final Results on the $0\nu\beta\beta$ Decay Half-Life Limit of ^{100}Mo from the CUPID-Mo Experiment”. In: *The European Physical Journal C* 82.11 (Nov. 15, 2022). ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-022-10942-5](https://doi.org/10.1140/epjc/s10052-022-10942-5). URL: <https://link.springer.com/10.1140/epjc/s10052-022-10942-5>.
- [151] KamLAND-Zen Collaboration et al. “Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen”. In: *Physical Review Letters* 130.5 (Jan. 30, 2023), p. 051801. DOI: [10.1103/PhysRevLett.130.051801](https://doi.org/10.1103/PhysRevLett.130.051801). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.130.051801>.
- [152] C. Arnaboldi et al. “Production of High Purity TeO₂ Single Crystals for the Study of Neutrinoless Double Beta Decay”. In: *Journal of Crystal Growth* 312.20 (Oct. 2010), pp. 2999–3008. ISSN: 0022-0248. DOI: [10.1016/j.jcrysgro.2010.06.034](https://doi.org/10.1016/j.jcrysgro.2010.06.034). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0022024810004343>.
- [153] F. Alessandria et al. “CUORE Crystal Validation Runs: Results on Radioactive Contamination and Extrapolation to CUORE Background”. In: *Astroparticle Physics* 35.12 (July 2012), pp. 839–849. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2012.02.008](https://doi.org/10.1016/j.astropartphys.2012.02.008). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650512000473>.
- [154] K. Alfonso et al. “CUPID, the CUORE Upgrade with Particle Identification”. In: *The European Physical Journal C* 85.7 (July 3, 2025), p. 737. ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-025-14352-1](https://doi.org/10.1140/epjc/s10052-025-14352-1). URL: <https://link.springer.com/10.1140/epjc/s10052-025-14352-1>.
- [155] CUPID Collaboration et al. “A Gravity-Based Mounting Approach for Large-Scale Cryogenic Calorimeter Arrays”. In: *The European Physical Journal C* 85.9 (Sept. 2, 2025), p. 935. ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-025-14613-z](https://doi.org/10.1140/epjc/s10052-025-14613-z). URL: <https://link.springer.com/10.1140/epjc/s10052-025-14613-z>.
- [156] C. Arnaboldi et al. “CdWO₄ Scintillating Bolometer for Double Beta Decay: Light and Heat Anticorrelation, Light Yield and Quenching Factors”. In: *Astroparticle Physics* 34.3 (Oct. 2010), pp. 143–150. ISSN: 0927-6505. DOI: [10.1016/j.astropartphys.2010.06.009](https://doi.org/10.1016/j.astropartphys.2010.06.009). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650510001210>.
- [157] J. Amaré et al. “Light Yield of Undoped Sapphire at Low Temperature under Particle Excitation”. In: *Applied Physics Letters* 87.26 (Dec. 26, 2005). ISSN: 0003-6951, 1077-3118. DOI: [10.1063/1.2158518](https://doi.org/10.1063/1.2158518). URL: <https://pubs.aip.org/apl/article/87/26/264102/330503/Light-yield-of-undoped-sapphire-at-low-temperature>.
- [158] Francesca Cova et al. “Effect of Point Defects on the Performances of Li₂MoO₄ and Li₂WO₄ Crystals as Scintillating Cryogenic Calorimeters”. In: *Journal of Alloys and Compounds* 1022 (Apr. 2025), p. 179848. ISSN: 09258388. DOI: [10.1016/j.jallcom.2025.179848](https://doi.org/10.1016/j.jallcom.2025.179848). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925838825014069>.
- [159] D.A. Spassky et al. “Trap Centers in Molybdates”. In: *Optical Materials* 35.12 (Oct. 2013), pp. 2465–2472. ISSN: 0925-3467. DOI: [10.1016/j.optmat.2013.06.054](https://doi.org/10.1016/j.optmat.2013.06.054). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925346713003662>.

- [160] D.A. Spassky et al. “Low Temperature Luminescence and Charge Carrier Trapping in a Cryogenic Scintillator Li_2MoO_4 ”. In: *Journal of Luminescence* 166 (Oct. 2015), pp. 195–202. ISSN: 0022-2313. DOI: [10.1016/j.jlumin.2015.05.042](https://doi.org/10.1016/j.jlumin.2015.05.042). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0022231315002896>.
- [161] B. S. Neganov and V. N. Trofimov. “Calorimetric Method Measuring Ionizing Radiation”. In: *Otkryt. Izobret.* 146.USSR Patent No 1037771 (1985), p. 215.
- [162] P. N. Luke. “Voltage-Assisted Calorimetric Ionization Detector”. In: *Journal of Applied Physics* 64.12 (Dec. 15, 1988), pp. 6858–6860. ISSN: 0021-8979, 1089-7550. DOI: [10.1063/1.341976](https://doi.org/10.1063/1.341976). URL: <https://pubs.aip.org/jap/article/64/12/6858/174257/Voltage-assisted-calorimetric-ionization-detector>.
- [163] M. Biassoni et al. “Large Area Si Low-Temperature Light Detectors with Neganov–Luke Effect”. In: *The European Physical Journal C* 75.10 (Oct. 2015), p. 480. ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-015-3712-2](https://doi.org/10.1140/epjc/s10052-015-3712-2). URL: <http://link.springer.com/10.1140/epjc/s10052-015-3712-2>.
- [164] V. Novati et al. “Charge-to-Heat Transducers Exploiting the Neganov-Trofimov-Luke Effect for Light Detection in Rare-Event Searches”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 940 (Oct. 2019), pp. 320–327. ISSN: 0168-9002. DOI: [10.1016/j.nima.2019.06.044](https://doi.org/10.1016/j.nima.2019.06.044). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900219308897>.
- [165] C. Augier et al. “The Background Model of the CUPID-Mo $0\nu\beta\beta$ Experiment”. In: *The European Physical Journal C* 83.7 (July 28, 2023). ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-023-11830-2](https://doi.org/10.1140/epjc/s10052-023-11830-2). URL: <https://link.springer.com/10.1140/epjc/s10052-023-11830-2>.
- [166] D. M. Chernyak et al. “Random Coincidence of $2\nu 2\beta$ Decay Events as a Background Source in Bolometric $0\nu 2\beta$ Decay Experiments”. In: *The European Physical Journal C* 72.4 (Apr. 2012). ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-012-1989-y](https://doi.org/10.1140/epjc/s10052-012-1989-y). URL: <http://link.springer.com/10.1140/epjc/s10052-012-1989-y>.
- [167] D. Q. Adams et al. *First-Ever Detection of Microseismic Activity with a Tonne-Scale Cryogenic Experiment*. Version 1. 2025. DOI: [10.48550/ARXIV.2505.09652](https://doi.org/10.48550/ARXIV.2505.09652). URL: <https://arxiv.org/abs/2505.09652>. Pre-published.
- [168] A. D’Addabbo et al. “An Active Noise Cancellation Technique for the CUORE Pulse Tube Cryocoolers”. In: *Cryogenics* 93 (July 2018), pp. 56–65. ISSN: 0011-2275. DOI: [10.1016/j.cryogenics.2018.05.001](https://doi.org/10.1016/j.cryogenics.2018.05.001). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0011227517304460>.
- [169] C. Augier et al. “Measurement of the $2\nu\beta\beta$ Decay Rate and Spectral Shape of Mo100 from the CUPID-Mo Experiment”. In: *Physical Review Letters* 131.16 (Oct. 18, 2023). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/physrevlett.131.162501](https://doi.org/10.1103/physrevlett.131.162501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.162501>.
- [170] A. Ahmine et al. “Enhanced Light Signal for the Suppression of Pile-up Events in Mo-based Bolometers for the $0\nu\beta\beta$ Decay Search.” In: *The European Physical Journal C* 83.5 (May 6, 2023). ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-023-11519-6](https://doi.org/10.1140/epjc/s10052-023-11519-6). URL: <https://link.springer.com/10.1140/epjc/s10052-023-11519-6>.

- [171] A. Agrawal et al. “Improved Limit on Neutrinoless Double Beta Decay of Mo100 from AMoRE-I”. In: *Physical Review Letters* 134.8 (Feb. 27, 2025). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/physrevlett.134.082501](https://doi.org/10.1103/PhysRevLett.134.082501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.134.082501>.
- [172] K. Alfonso et al. *Sensitivity of the CUPID Experiment to $0\nu\beta\beta$ Decay of ^{100}Mo* . Apr. 19, 2025. URL: <http://arxiv.org/abs/2504.14369>. Pre-published.
- [173] E. Armengaud et al. “Development of ^{100}Mo -containing Scintillating Bolometers for a High-Sensitivity Neutrinoless Double-Beta Decay Search”. In: *The European Physical Journal C* 77.11 (Nov. 2017). ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-017-5343-2](https://doi.org/10.1140/epjc/s10052-017-5343-2). URL: <http://link.springer.com/10.1140/epjc/s10052-017-5343-2>.
- [174] Veronika Grigorieva et al. “Li₂MoO₄ Crystals Grown by Low-Thermal-Gradient Czochralski Technique”. In: *Journal of Materials Science and Engineering B* 7.2 (Apr. 28, 2017). ISSN: 2161-6221, 2161-6221. DOI: [10.17265/2161-6221/2017.3-4.002](https://doi.org/10.17265/2161-6221/2017.3-4.002). URL: <http://www.davidpublisher.org/index.php/Home/Article/index?id=31529.html>.
- [175] P. Gorla, C. Bucci, and S. Pirro. “Complete Elimination of 1K Pot Vibrations in Dilution Refrigerators”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 520.1–3 (Mar. 2004), pp. 641–643. ISSN: 0168-9002. DOI: [10.1016/j.nima.2003.11.365](https://doi.org/10.1016/j.nima.2003.11.365). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900203032856>.
- [176] S. Pirro. “Further Developments in Mechanical Decoupling of Large Thermal Detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 559.2 (Apr. 2006), pp. 672–674. ISSN: 0168-9002. DOI: [10.1016/j.nima.2005.12.197](https://doi.org/10.1016/j.nima.2005.12.197). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900205025350>.
- [177] Agilent Technologies and Keysight Technologies. *Agilent/HP 6627A Quad-Output DC Power Supply Datasheet*. 2002. URL: <https://www.testunlimited.com/pdf/Keysight-6627A-Specs.pdf>.
- [178] A.M. Vignati et al. “First Data from CUORE-0”. In: *Physics Procedia* 61 (2015), pp. 289–294. ISSN: 1875-3892. DOI: [10.1016/j.phpro.2014.12.047](https://doi.org/10.1016/j.phpro.2014.12.047). URL: <https://linkinghub.elsevier.com/retrieve/pii/S1875389214006609>.
- [179] D. Q. Adams et al. “Improved Limit on Neutrinoless Double-Beta Decay in Te130 with CUORE”. In: *Physical Review Letters* 124.12 (Mar. 26, 2020). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/physrevlett.124.122501](https://doi.org/10.1103/PhysRevLett.124.122501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.122501>.
- [180] J. W. Beeman et al. “Double-Beta Decay Investigation with Highly Pure Enriched ^{82}Se for the LUCIFER Experiment”. In: *The European Physical Journal C* 75.12 (Dec. 2015). ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-015-3822-x](https://doi.org/10.1140/epjc/s10052-015-3822-x). URL: <http://link.springer.com/10.1140/epjc/s10052-015-3822-x>.
- [181] E. Armengaud et al. “The CUPID-Mo Experiment for Neutrinoless Double-Beta Decay: Performance and Prospects”. In: *The European Physical Journal C* 80.1 (Jan. 2020). ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-019-7578-6](https://doi.org/10.1140/epjc/s10052-019-7578-6). URL: <http://link.springer.com/10.1140/epjc/s10052-019-7578-6>.

- [182] I. Antcheva et al. “ROOT — A C++ Framework for Petabyte Data Storage, Statistical Analysis and Visualization”. In: *Computer Physics Communications* 180.12 (Dec. 2009), pp. 2499–2512. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2009.08.005](https://doi.org/10.1016/j.cpc.2009.08.005). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010465509002550>.
- [183] Michael Stonebraker and Lawrence A. Rowe. “The Design of POSTGRES”. In: *ACM SIGMOD Record* 15.2 (June 15, 1986), pp. 340–355. ISSN: 0163-5808. DOI: [10.1145/16856.16888](https://doi.org/10.1145/16856.16888). URL: <https://dl.acm.org/doi/10.1145/16856.16888>.
- [184] James W. Cooley and John W. Tukey. “An Algorithm for the Machine Calculation of Complex Fourier Series”. In: *Mathematics of Computation* 19.90 (Apr. 1965), p. 297. ISSN: 0025-5718. DOI: [10.2307/2003354](https://doi.org/10.2307/2003354). JSTOR: [2003354](https://www.jstor.org/stable/2003354). URL: <https://www.jstor.org/stable/2003354?origin=crossref>.
- [185] Ulrich Schötzgig. “Half-Life and X-ray Emission Probabilities of ^{55}Fe ”. In: *Applied Radiation and Isotopes* 53.3 (Sept. 2000), pp. 469–472. ISSN: 0969-8043. DOI: [10.1016/S0969-8043\(00\)00166-4](https://doi.org/10.1016/S0969-8043(00)00166-4). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0969804300001664>.
- [186] J. W. Beeman et al. “New Experimental Limits on the α Decays of Lead Isotopes”. In: *The European Physical Journal A* 49.4 (Apr. 2013). ISSN: 1434-6001, 1434-601X. DOI: [10.1140/epja/i2013-13050-7](https://doi.org/10.1140/epja/i2013-13050-7). URL: <http://link.springer.com/10.1140/epja/i2013-13050-7>.
- [187] J.W. Beeman et al. “A Next-Generation Neutrinoless Double Beta Decay Experiment Based on ZnMoO₄ Scintillating Bolometers”. In: *Physics Letters B* 710.2 (Apr. 2012), pp. 318–323. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.03.009](https://doi.org/10.1016/j.physletb.2012.03.009). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0370269312002602>.
- [188] F. Bellini et al. “Response of a TeO₂ Bolometer to α Particles”. In: *Journal of Instrumentation* 5.12 (2010). ISSN: 17480221. DOI: [10.1088/1748-0221/5/12/P12005](https://doi.org/10.1088/1748-0221/5/12/P12005). URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-78751512052&doi=10.1088%2f1748-0221%2f5%2f12%2fP12005&partnerID=40&md5=84694518d47a00c53ea14ec367bedcb3>.
- [189] A. Alessandrello et al. “The Thermal Detection Efficiency for Recoils Induced by Low Energy Nuclear Reactions, Neutrinos or Weakly Interacting Massive Particles”. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 408.1–4 (1997), pp. 465–468. ISSN: 03702693. DOI: [10.1016/S0370-2693\(97\)00765-X](https://doi.org/10.1016/S0370-2693(97)00765-X). URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0002478999&doi=10.1016%2fS0370-2693%2897%2900765-X&partnerID=40&md5=a79f46280a9d2229dfa29a13d10f1031>.
- [190] L Cardani et al. “Development of a Li₂MoO₄scintillating Bolometer for Low Background Physics”. In: *Journal of Instrumentation* 8.10 (Oct. 2, 2013), P10002–P10002. ISSN: 1748-0221. DOI: [10.1088/1748-0221/8/10/p10002](https://doi.org/10.1088/1748-0221/8/10/p10002). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/8/10/P10002>.
- [191] A. Armatol et al. “A CUPID Li₂¹⁰⁰MoO₄ Scintillating Bolometer Tested in the CROSS Underground Facility”. In: *Journal of Instrumentation* 16.02 (Feb. 24, 2021), P02037–P02037. ISSN: 1748-0221. DOI: [10.1088/1748-0221/16/02/p02037](https://doi.org/10.1088/1748-0221/16/02/p02037). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/16/02/P02037>.
- [192] Howard W. Johnson and Martin Graham. *High Speed Digital Design: A Handbook of Black Magic*. 30. pr. Upper Saddle River, NJ: Prentice Hall PTR, 1993. 447 pp. ISBN: 978-0-13-395724-2.

- [193] Gary J. Feldman and Robert D. Cousins. “Unified Approach to the Classical Statistical Analysis of Small Signals”. In: *Physical Review D* 57.7 (Apr. 1, 1998), pp. 3873–3889. ISSN: 0556-2821, 1089-4918. DOI: [10.1103/PhysRevD.57.3873](https://doi.org/10.1103/PhysRevD.57.3873). URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.3873>.
- [194] T.B. Bekker et al. “Aboveground Test of an Advanced Li_2MoO_4 Scintillating Bolometer to Search for Neutrinoless Double Beta Decay of ^{100}Mo ”. In: *Astroparticle Physics* 72 (Jan. 2016), pp. 38–45. ISSN: 09276505. DOI: [10.1016/j.astropartphys.2015.06.002](https://doi.org/10.1016/j.astropartphys.2015.06.002). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650515000973>.
- [195] O.P. Barinova et al. “First Test of Li_2MoO_4 Crystal as a Cryogenic Scintillating Bolometer”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 613.1 (Jan. 2010), pp. 54–57. ISSN: 01689002. DOI: [10.1016/j.nima.2009.11.059](https://doi.org/10.1016/j.nima.2009.11.059). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900209022712>.
- [196] Yu. Hizhnyi et al. “Role of Native and Impurity Defects in Optical Absorption and Luminescence of Li_2MoO_4 Scintillation Crystals”. In: *Journal of Alloys and Compounds* 867 (June 2021), p. 159148. ISSN: 09258388. DOI: [10.1016/j.jallcom.2021.159148](https://doi.org/10.1016/j.jallcom.2021.159148). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925838821005557>.
- [197] M. Buryi et al. “Electron and Hole Trapping in Li_2MoO_4 Cryogenic Scintillator”. In: *Optical Materials* 114 (Apr. 2021), p. 110971. ISSN: 09253467. DOI: [10.1016/j.optmat.2021.110971](https://doi.org/10.1016/j.optmat.2021.110971). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0925346721001725>.
- [198] Wenyu Liu et al. “ Li_2MoO_4 晶体的坩埚下降法生长及其发光性能 (Growth of Li_2MoO_4 Crystals by Crucible Descent Method and Their Luminescence Properties)”. In: *Journal Of Synthetic Crystals* 54.5 (May 2025). DOI: [10.16553/j.cnki.issn1000-985x.2024.0270](https://doi.org/10.16553/j.cnki.issn1000-985x.2024.0270).
- [199] *Symphony II Detection System: Operation Manual*. Version Part No. J810010 (Rev. C). Edison, NJ, USA: HORIBA Jobin Yvon, Optical Spectroscopy Division, 2003. URL: https://www.artisantg.com/info/Horiba_Jobin_Yvon_Symphony_Solo_Manual_201911812246.pdf.
- [200] D. A. Spasskii et al. “Luminescence Peculiarities and Optical Properties of MgMoO_4 and $\text{MgMoO}_4:\text{Yb}$ Crystals”. In: *Optics and Spectroscopy* 106.4 (Apr. 2009), pp. 556–563. ISSN: 0030-400X, 1562-6911. DOI: [10.1134/S0030400X09040171](https://doi.org/10.1134/S0030400X09040171). URL: <http://link.springer.com/10.1134/S0030400X09040171>.
- [201] Dmitry Spassky et al. “Luminescence Investigation of Zinc Molybdate Single Crystals”. In: *physica status solidi (a)* 206.7 (July 2009), pp. 1579–1583. ISSN: 1862-6300, 1862-6319. DOI: [10.1002/pssa.200824311](https://doi.org/10.1002/pssa.200824311). URL: <https://onlinelibrary.wiley.com/doi/10.1002/pssa.200824311>.
- [202] Stephen W. S. McKeever. *A Course in Luminescence Measurements and Analyses for Radiation Dosimetry*. Hoboken: Wiley, 2022. 390 pp. ISBN: 978-1-119-64689-1.
- [203] Howard E Swanson et al. *Standard X-Ray Diffraction Powder Patterns*. NBS MONO 25-1. Gaithersburg, MD: National Bureau of Standards, 1962, NBS MONO 25-1. DOI: [10.6028/NBS.MONO.25-1](https://doi.org/10.6028/NBS.MONO.25-1). URL: <https://nvlpubs.nist.gov/nistpubs/Legacy/MONO/nbsmonograph25-1.pdf>.

- [204] Luca Foggetta et al. “Surface-Sensitive Macrobolometers for the Identification of External Charged Particles”. In: *Applied Physics Letters* 86.13 (Mar. 28, 2005), p. 134106. ISSN: 0003-6951, 1077-3118. DOI: [10.1063/1.1894608](https://doi.org/10.1063/1.1894608). URL: <https://pubs.aip.org/apl/article/86/13/134106/329676/Surface-sensitive-macrobolometers-for-the>.
- [205] Luca Foggetta et al. “Composite Macro-Bolometers for the Rejection of Surface Radioactive Background in Rare-Event Experiments”. In: *Astroparticle Physics* 34.11 (June 2011), pp. 809–821. ISSN: 09276505. DOI: [10.1016/j.astropartphys.2011.02.004](https://doi.org/10.1016/j.astropartphys.2011.02.004). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650511000399>.
- [206] C. Nones et al. “Superconducting Aluminum Layers as Pulse Shape Modifiers: An Innovative Solution to Fight Against Surface Background in Neutrinoless Double Beta Decay Experiments”. In: *Journal of Low Temperature Physics* 167.5–6 (June 2012), pp. 1029–1034. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-012-0558-y](https://doi.org/10.1007/s10909-012-0558-y). URL: <http://link.springer.com/10.1007/s10909-012-0558-y>.
- [207] H. Khalife et al. “The CROSS Experiment: Rejecting Surface Events by PSD Induced by Superconducting Films”. In: *Journal of Low Temperature Physics* 199.1–2 (Apr. 2020), pp. 19–26. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-020-02369-7](https://doi.org/10.1007/s10909-020-02369-7). URL: <http://link.springer.com/10.1007/s10909-020-02369-7>.
- [208] D. Auguste et al. “A Novel Mechanical Design of a Bolometric Array for the CROSS Double-Beta Decay Experiment”. In: *Journal of Instrumentation* 19.09 (Sept. 1, 2024), P09014. ISSN: 1748-0221. DOI: [10.1088/1748-0221/19/09/P09014](https://doi.org/10.1088/1748-0221/19/09/P09014). URL: <https://iopscience.iop.org/article/10.1088/1748-0221/19/09/P09014>.
- [209] I. C. Bandac et al. “Phonon-Mediated Crystal Detectors with Metallic Film Coating Capable of Rejecting α and β Events Induced by Surface Radioactivity”. In: *Applied Physics Letters* 118.18 (May 3, 2021), p. 184105. ISSN: 0003-6951, 1077-3118. DOI: [10.1063/5.0050124](https://doi.org/10.1063/5.0050124). URL: <https://pubs.aip.org/apl/article/118/18/184105/40157/Phonon-mediated-crystal-detectors-with-metallic>.
- [210] John F. Cochran and D. E. Mapother. “Superconducting Transition in Aluminum”. In: *Physical Review* 111.1 (July 1, 1958), pp. 132–142. ISSN: 0031-899X. DOI: [10.1103/PhysRev.111.132](https://doi.org/10.1103/PhysRev.111.132). URL: <https://link.aps.org/doi/10.1103/PhysRev.111.132>.
- [211] M. Biassoni et al. “Production of Monochromatic ^{228}Ra α -Sources for Detector Characterization”. In: *The European Physical Journal Plus* 137.12 (Dec. 5, 2022), p. 1309. ISSN: 2190-5444. DOI: [10.1140/epjp/s13360-022-03519-4](https://doi.org/10.1140/epjp/s13360-022-03519-4). URL: <https://link.springer.com/10.1140/epjp/s13360-022-03519-4>.
- [212] M. Biassoni et al. “A Thermal Model for Low Temperature TeO_2 Calorimeters”. In: *Journal of Low Temperature Physics* 206.1–2 (Jan. 2022), pp. 80–96. ISSN: 0022-2291, 1573-7357. DOI: [10.1007/s10909-021-02639-y](https://doi.org/10.1007/s10909-021-02639-y). URL: <https://link.springer.com/10.1007/s10909-021-02639-y>.
- [213] Frank Pobell. *Matter and Methods at Low Temperatures*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007. ISBN: 978-3-540-46356-6 978-3-540-46360-3. DOI: [10.1007/978-3-540-46360-3](https://doi.org/10.1007/978-3-540-46360-3). URL: <http://link.springer.com/10.1007/978-3-540-46360-3>.
- [214] Mets501. *Helium Phase Diagram*. Mar. 23, 2013. URL: https://commons.wikimedia.org/wiki/File:Helium_phase_diagram.svg.

- [215] N. H. Balshaw. *Practical Cryogenics: An Introduction to Laboratory Cryogenics*. Eynsham: Oxford Instruments, Scientific Research Division, 1996. ISBN: 978-0-9527594-0-9. URL: <https://research.physics.illinois.edu/bezryadin/links/practical%20cryogenics.pdf>.

Acknowledgements

This work would not have been possible without the support, guidance, and constructive feedback of several people, to whom sincere gratitude is expressed.

First, I would like to thank my tutor, Chiara Brofferio, who has been and will remain a fundamental reference in this field for me. I am equally grateful to my co-tutor, Luca Gironi, for the valuable discussions and for teaching me how to properly guide physical intuition, as well as for the continuous support shown throughout all three years of my PhD. I would also like to thank my supervisor, Matteo Biassoni, for his daily guidance, critical perspective, and for providing clear scientific direction during the development of this thesis, while consistently maintaining a stimulating and engaging research environment.

I would like to express a sincere thank you to Monica Sisti and Gianluigi Pessina, who provided me with two very different but complementary perspectives on experimental particle physics, while sharing part of their extensive, solid, and highly specialized expertise. Beyond being patient and excellent teachers, they have been a lasting source of inspiration in approaching complex scientific and technical decisions within a demanding research environment.

A special thank you also goes to Laura Marini, who shared with me her extensive expertise in the laboratory, together with a pragmatic approach that I did not have before meeting her and that I have since successfully integrated into my daily work and personality. Her guidance and systematic approach to problem solving significantly contributed to my development as a more effective researcher. I would also like to thank Mattia Beretta, who dedicated a significant fraction of his time to sharing with me his thoughts and clever considerations on the experimental operation of cryogenic calorimeters.

I would like to thank my friend Andrea Nava, with whom I spent many hours discussing experimental neutrino physics. These discussions were instrumental in clarifying my understanding of the role of neutrino experiments in the current experimental landscape.

I would like to thank the reviewers, Michelle Dolinski and Riccardo Brugnera, for the time and effort dedicated to the review of this work, and for their constructive comments and suggestions, which contributed to improving its clarity and overall quality.

Finally, I would like to express a general acknowledgment to all the colleagues and collaborators with whom I had the opportunity to work during this research activity. Their contributions, discussions, and shared experience have been a constant source of inspiration and guidance. I am grateful to everyone who, in different ways, supported this work and contributed to my scientific and professional growth.