



**Universidad**  
Zaragoza



Facultad de Ciencias  
**Universidad** Zaragoza

MEMORIA DE TESIS DOCTORAL  
PROGRAMA DE DOCTORADO EN FÍSICA  
2021-2024

---

**Simulation Infrastructure and Cosmic-Ray Background Modeling for  
BabyIAXO Micromegas Detectors**

*REST-for-Physics/restG4 workflows, cosmic-neutron studies, and active-veto  
validation*

---

Autor

**Luis Antonio Obis Aparicio**

Directores

**Dra. Gloria Luzón Marco**

Zaragoza, 1 de enero de 2025



# Abstract

BabyIAXO is the intermediate stage of the International Axion Observatory program and will operate as a solar axion helioscope with low-background X-ray detectors installed on a moving, surface-level apparatus. For the Micromegas detector line, this operating scenario makes the control of cosmic-ray-induced backgrounds a central requirement. The expected axion signal consists of a small excess of keV X-rays focused onto the detector plane, so the detector response, event reconstruction, background normalization, and veto strategy must be understood at the level of the same observables used in data analysis.

This thesis develops simulation infrastructure and background-modeling tools for IAXO-D0 and BabyIAXO Micromegas detectors, with emphasis on detector-response-level Monte Carlo workflows, cosmic-ray and cosmic-neutron-induced backgrounds, and active-veto validation. The software work is based on `REST-for-Physics` and its `Geant4` interface `restG4`, including workflow organization, source generation, detector-response chains, and large-scale production strategies. These tools are used to construct a background-model methodology in which source-specific simulations are propagated through common reconstruction and selection observables.

A major part of the work addresses the shielding and veto system required for surface operation. The studies show that passive shielding alone cannot fully suppress the high-energy neutron component, because cosmic neutrons can generate secondary cascades inside the lead shield. The resulting active-veto strategy combines prompt scintillator signals, delayed neutron-capture-related activity, and veto multiplicity in a multilayer plastic-scintillator and cadmium design. Waveform-level simulations predict strong rejection of muon-induced backgrounds and partial rejection of neutron- and proton-induced residuals after Micromegas cuts. Prototype data taken with the IAXO-D0 detector and veto system validate the prompt/delayed/multiplicity discrimination strategy, while also showing that an absolute neutron-veto efficiency measurement still requires a prototype-matched simulation and a more complete treatment of thresholds, timing, calibration, and source normalization.

The thesis therefore contributes to the transition from idealized energy-deposition simulations to reconstructed observables that can be compared with experimental data.

It provides the software and analysis basis for a quantitative BabyIAXO detector-background model and identifies the main remaining uncertainties: cosmic-neutron normalization, DESY site dependence, hadronic modeling, detector-response calibration, veto thresholds, timing alignment, geometry details, and finite Monte Carlo statistics.

# Resumen

BabyIAXO es la etapa intermedia del programa del International Axion Observatory y funcionará como un helioscopio solar de axiones con detectores de rayos X de bajo fondo instalados en una infraestructura móvil a nivel de superficie. Para la línea de detectores Micromegas, estas condiciones hacen que el control del fondo inducido por rayos cósmicos sea un requisito central. La señal esperada consiste en un pequeño exceso de rayos X de energía keV focalizados sobre el plano del detector, por lo que la respuesta del detector, la reconstrucción de sucesos, la normalización del fondo y la estrategia de veto deben entenderse al nivel de los mismos observables utilizados en el análisis de datos.

Esta tesis desarrolla infraestructura de simulación y herramientas de modelado de fondo para los detectores Micromegas de IAXO-D0 y BabyIAXO, con énfasis en flujos de trabajo Monte Carlo a nivel de respuesta del detector, fondos inducidos por rayos cósmicos y neutrones cósmicos, y validación del veto activo. El trabajo de software se basa en **REST-for-Physics** y en su interfaz **restG4** con **Geant4**, incluyendo la organización de flujos de trabajo, la generación de fuentes, las cadenas de respuesta del detector y las estrategias de producción a gran escala. Estas herramientas se utilizan para construir una metodología de modelo de fondo en la que simulaciones específicas de cada fuente se procesan mediante observables comunes de reconstrucción y selección.

Una parte importante del trabajo se centra en el sistema de blindaje y veto necesario para la operación en superficie. Los estudios muestran que el blindaje pasivo por sí solo no puede suprimir completamente la componente de neutrones de alta energía, ya que los neutrones cósmicos pueden generar cascadas secundarias dentro del blindaje de plomo. La estrategia de veto activo resultante combina señales rápidas en centelleadores, actividad retardada asociada a capturas neutrónicas y multiplicidad de canales en un diseño multicapa de centelleador plástico y cadmio. Las simulaciones a nivel de forma de onda predicen un fuerte rechazo del fondo inducido por muones y un rechazo parcial de los residuos inducidos por neutrones y protones tras los cortes Micromegas. Los datos del prototipo IAXO-D0 con sistema de veto validan la estrategia de discriminación basada en señales rápidas, retardadas y de multiplicidad, aunque una medida absoluta de la eficiencia del veto de neutrones requiere todavía una simulación

ajustada al prototipo y un tratamiento más completo de umbrales, tiempos, calibración y normalización de fuentes.

La tesis contribuye así a la transición desde simulaciones idealizadas de depósito de energía hacia observables reconstruidos comparables con datos experimentales. Proporciona la base de software y análisis para un modelo cuantitativo de fondo de los detectores BabyIAXO e identifica las principales incertidumbres restantes: normalización de neutrones cósmicos, dependencia del emplazamiento en DESY, modelado hadrónico, calibración de la respuesta del detector, umbrales del veto, alineamiento temporal, detalles geométricos y estadística finita de Monte Carlo.

# Agradecimientos

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetur id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.



# Contents

<b>Abstract</b>	<b>3</b>
<b>Resumen</b>	<b>5</b>
<b>Introduction</b>	<b>13</b>
<b>1 Axions and Axion-Like Particles</b>	<b>17</b>
1.1 The Strong CP Problem . . . . .	18
1.2 Axion and Axion-Like Particles (ALPs) . . . . .	19
1.3 Axion Detection . . . . .	22
1.3.1 Helioscopes . . . . .	22
1.3.2 Light Shining Through Wall . . . . .	22
1.3.3 Helioscopes . . . . .	23
<b>2 The IAXO Collaboration and the BabyIAXO Experiment</b>	<b>25</b>
2.1 The Helioscope Technique . . . . .	27
2.1.1 Solar Axions . . . . .	27
2.1.2 Axion-Photon Conversion in a Helioscope . . . . .	28
2.1.3 Experimental Figure of Merit . . . . .	29
2.2 From CAST to IAXO . . . . .	29
2.2.1 CAST as the Reference Helioscope . . . . .	29
2.2.2 Why a Dedicated Helioscope Is Required . . . . .	30
2.3 The IAXO Experiment . . . . .	31
2.3.1 Physics Reach and Design Philosophy . . . . .	31
2.3.2 Magnet . . . . .	32
2.3.3 X-ray Optics . . . . .	32
2.3.4 Detector Technologies . . . . .	33
2.3.5 Tracking and Observatory Operation . . . . .	33
2.4 The IAXO Collaboration and Project Context . . . . .	34
2.5 BabyIAXO as an Intermediate Stage . . . . .	34

---

2.5.1	Role Within the IAXO Program . . . . .	34
2.5.2	Evolution of the DESY Site Scenario . . . . .	35
2.5.3	Main Experimental Features . . . . .	37
2.6	Connection to the Present Thesis . . . . .	38
<b>3</b>	<b>Micromegas X-ray Detectors for BabyIAXO</b>	<b>39</b>
3.1	The Time Projection Chamber (TPC) . . . . .	40
3.1.1	Ionization in Gases . . . . .	41
3.1.2	Electron Transport . . . . .	47
3.1.3	Micropattern Gaseous Detectors: Micromegas . . . . .	53
3.2	The BabyIAXO Micromegas Detector Prototypes . . . . .	55
3.2.1	IAXO-D0 . . . . .	56
3.2.2	IAXO-D1 . . . . .	57
3.3	Detector Ancillary Systems . . . . .	62
3.3.1	Gas System . . . . .	62
3.3.2	High-Voltage System . . . . .	64
3.3.3	Slow-Control System . . . . .	65
3.4	Shielding and Veto Systems . . . . .	65
3.5	Data Acquisition System . . . . .	65
3.5.1	Data Acquisition Software . . . . .	69
3.6	Detector Calibration . . . . .	74
3.6.1	Standard x-ray calibration . . . . .	75
3.6.2	UV-light calibration R&D . . . . .	78
<b>4</b>	<b>Computational Framework</b>	<b>79</b>
4.1	Scope and software contributions of this thesis . . . . .	80
4.2	The ROOT Data Analysis Framework and Python interoperability . . . . .	81
4.2.1	Uproot and fsspec integration . . . . .	81
4.3	REST-for-Physics as the common event-processing framework . . . . .	82
4.3.1	Event-based processing and the EventTree/AnalysisTree model . . . . .	83
4.3.2	Reproducibility and maintainability . . . . .	85
4.4	Detector readout metadata and channel mapping . . . . .	85
4.5	Geant4 simulation infrastructure . . . . .	87
4.5.1	restG4 and Geant4Lib: the interface to Geant4 . . . . .	87
4.5.2	Geometry-generation infrastructure for IAXO simulations . . . . .	89
4.5.3	Python interface to REST-for-Physics . . . . .	90
4.5.4	Probability Distribution Projection cosmic-ray generator . . . . .	90
4.5.5	Cosmic-ray source generators . . . . .	92
4.5.6	Radiation transport in gases: Garfield++ . . . . .	93
4.6	Visualization and online diagnostics . . . . .	94

---

4.6.1	Browser-based <b>Geant4</b> event viewer . . . . .	94
4.6.2	DAQ event viewer: <b>feminos-viewer</b> . . . . .	96
4.7	Monte Carlo production at scale . . . . .	97
4.7.1	End-to-end simulation and analysis chain . . . . .	97
4.7.2	Production workflow with <b>HTCondor</b> . . . . .	98
4.8	<b>geant4-python-application</b> : Python prototyping with <b>Geant4</b> . . .	102
4.9	Summary and role in the thesis . . . . .	104
<b>5</b>	<b>Shielding and Veto System</b> . . . . .	<b>107</b>
5.1	Motivation: surface backgrounds and veto requirements . . . . .	108
5.1.1	Design requirements and constraints . . . . .	108
5.1.2	Surface cosmic-ray source term . . . . .	110
5.1.3	Relevant interaction processes in the shielding and veto materials	116
5.1.4	Interaction of high-energy neutrons and photons with the detector	124
5.1.5	Effect of detector inclination during solar tracking . . . . .	129
5.2	Limits of passive shielding . . . . .	131
5.2.1	Lead-thickness scan: electromagnetic suppression versus neutron multiplication . . . . .	131
5.2.2	Borated-HDPE moderation studies . . . . .	133
5.3	Neutron-sensitive active veto concept . . . . .	136
5.3.1	Plastic scintillators . . . . .	136
5.3.2	From prompt muon veto to neutron-sensitive stages . . . . .	139
5.4	Multilayer veto optimization . . . . .	145
5.5	Final veto geometry and readout . . . . .	148
5.5.1	Prototype construction and commissioning . . . . .	151
5.5.2	Electronics and readout . . . . .	152
5.6	Waveform-level veto observables . . . . .	154
5.6.1	Boosted veto BSD selection . . . . .	157
5.7	Simulation performance . . . . .	161
5.8	Experimental validation with IAXO-D0 data . . . . .	167
5.9	Systematic limitations and remaining work . . . . .	183
5.9.1	Design-level systematics . . . . .	183
5.9.2	Limitations of the neutron-enriched comparison . . . . .	184
5.10	Summary and outlook . . . . .	185
<b>6</b>	<b>Background Model</b> . . . . .	<b>187</b>
6.1	Radiopurity Measurements . . . . .	189
6.2	Background Measurements . . . . .	192
6.3	Simulation and Analysis Methodology . . . . .	192
6.3.1	Workflow overview . . . . .	192

6.3.2	Source-specific simulations . . . . .	194
6.3.3	Background-model component catalogue . . . . .	197
6.3.4	Event types and detector-response chain . . . . .	199
6.3.5	Analysis processes and observables . . . . .	204
6.3.6	Reference samples for cuts and normalization . . . . .	211
6.4	External Background . . . . .	221
6.4.1	Cosmic Rays . . . . .	222
6.4.2	Cosmic-ray event classes and veto survival . . . . .	223
6.4.3	Environmental Background . . . . .	224
6.5	Intrinsic Background . . . . .	234
6.5.1	Gas Contamination . . . . .	234
6.5.2	Radon Contamination . . . . .	235
6.5.3	Detector Materials . . . . .	239
6.5.4	Front-end Electronics . . . . .	243
6.5.5	Shielding . . . . .	248
6.6	Background-model status and uncertainty roadmap . . . . .	253
<b>Summary and Conclusions</b>		<b>257</b>
<b>Appendix</b>		<b>261</b>
6.7	Supplementary interaction plots for detector and veto design . . . . .	262
6.8	Supplementary veto simulation campaign metadata . . . . .	265
6.9	Supplementary active-veto design diagnostics . . . . .	267
6.10	Supplementary HENSA veto layer-scan diagnostics . . . . .	270
6.11	Supplementary neutron-tagging diagnostics . . . . .	272
6.12	Supplementary neutron-enriched data-simulation diagnostics . . . . .	272
6.13	Supplementary passive-shielding scans . . . . .	273
6.14	Supplementary environmental-radioactivity plots . . . . .	276
6.15	Supplementary Micromegas system diagrams . . . . .	281
6.16	Supplementary UV-light calibration R&D . . . . .	283
<b>Bibliography</b>		<b>287</b>

# Introduction

The Standard Model of particle physics provides an exceptionally successful description of the known elementary particles and their interactions. Nevertheless, several observations and theoretical questions point to physics beyond this framework. Among them, the nature of dark matter and the absence of observed CP violation in the strong interaction remain two of the most compelling open problems. The axion was originally proposed as a dynamical solution to the strong CP problem, but it also emerged as a well-motivated dark-matter candidate. More generally, axion-like particles appear naturally in many extensions of the Standard Model and can be searched for through their weak couplings to photons, electrons, and nucleons.

Solar axion helioscopes exploit one of the most direct experimental signatures of these particles. If axions are produced in the solar interior, they can traverse the Sun and the interplanetary medium essentially unattenuated. Inside a strong laboratory magnetic field, a small fraction of them can convert into X-ray photons through the inverse Primakoff effect. The experimental task is therefore conceptually simple but technically demanding: point a powerful magnet toward the Sun, focus any converted photons onto a small detector area, and identify a possible excess of keV X-rays above an extremely low background.

The International Axion Observatory (IAXO) is designed as the next major step in this technique, building on the experience of previous helioscopes and especially on the CERN Axion Solar Telescope (CAST). BabyIAXO is the intermediate stage of this program. It will validate the main technologies required for IAXO while also operating as a competitive helioscope in its own right. For the detector line studied in this thesis, BabyIAXO is not only a scaled-down version of a future experiment. It is a realistic environment in which low-background X-ray detection, solar tracking, mechanical integration, and surface-level operation must be made compatible.

This thesis focuses on the Micromegas detector line developed for IAXO-D0 and BabyIAXO. Microbulk Micromegas detectors are well suited to helioscope searches because they combine low intrinsic radioactivity, good energy response in the keV range, topological discrimination, and compatibility with compact shielding and focusing optics. However, the expected signal rate is extremely small. The physics reach of

the experiment therefore depends not only on detector performance, but also on the reliability of the background model, the realism of the detector-response simulation, and the effectiveness of the shielding and active veto strategy.

The work presented here addresses these requirements from three connected directions. First, it describes contributions to the software and simulation infrastructure used by the collaboration, with particular emphasis on `REST-for-Physics`, its `Geant4` interface `restG4`, and the production workflows needed for large Monte Carlo campaigns. Second, it develops a background-model methodology for IAXO-D0 and BabyIAXO, combining radiopurity information, environmental measurements, cosmic-ray source terms, detector-response emulation, and common reconstruction observables. Third, it studies the surface-level cosmic-ray-induced background and the corresponding active veto system, including the optimization of a multilayer plastic-scintillator and cadmium design, waveform-level veto observables, construction and commissioning aspects, and comparison with experimental data.

This thesis develops the software and detector-response simulation infrastructure required to build a quantitative background model for BabyIAXO Micromegas detectors, with emphasis on surface cosmic-ray backgrounds, cosmic-neutron-induced residuals, and active-veto observables comparable to experimental data.

A central theme of the thesis is the transition from idealized background estimates to analysis objects that can be compared with real detector data. The relevant question is not only whether a simulated particle deposits energy in the detector volume, but whether the resulting event would pass the same energy, topology, timing, and veto selections applied to the experimental data. For this reason, the simulations are propagated through a detector-response chain whenever possible, and the veto studies are expressed in terms of prompt signals, delayed activity, channel multiplicity, and reconstructed observables. This approach is especially important for surface operation, where muons, high-energy neutrons, and secondary particles produced in the shielding can generate backgrounds that are not adequately described by passive shielding arguments alone.

Contribution	Location	Original role in the thesis	Validation / status
REST-for-Physics/restG4 detector-response workflow	Chs. 4 and 6	Common simulation and reconstruction path for source-specific Monte Carlo samples and detector-like observables.	Used as the basis of the background-model methodology.
Geometry and GDML infrastructure	Chs. 4-6	High-level, version-controlled geometry generation for IAXO-D0, shielding, and veto configurations.	Used in passive-shield scans, cosmic-ray simulations, and veto-optimization studies.
Cosmic-ray source generation and injection	Chs. 4 and 5	Efficient generation and transport strategies for surface cosmic-ray backgrounds reaching the detector geometry.	Applied to cosmic-background and veto studies; normalization remains an input uncertainty.
Cosmic-neutron background studies	Chs. 5 and 6	Identification of high-energy cosmic neutrons as a residual component not solved by lead shielding alone.	Spectral shape checked against reference calculations and HENSA data; final normalization pending.
Waveform-level veto observables	Ch. 5	Prompt, delayed, and multiplicity-based veto logic connected to simulated and measured waveforms.	Prototype data validate the discrimination strategy, not yet an absolute neutron efficiency.
Background-model integration	Ch. 6	Source-specific simulations organized through a common detector-response and selection framework.	Requires final source normalizations and a complete master-rate table before final review.
Experimental-data reanalysis	Chs. 5 and 6	Analysis of existing IAXO-D0 calibration and background data with the same REST-for-Physics observable model used for simulations.	Used to validate veto observables, accidental-coincidence models, and the applicability of topology selections to real data.

Table 1: Summary of the main original contributions and their status within the thesis.

The structure of the thesis follows this logic. Chapter 1 introduces the axion and axion-like-particle motivation, the strong CP problem, and the main experimental approaches used in axion searches. Chapter 2 describes the IAXO program, the helioscope figure of merit, the role of BabyIAXO, and the experimental context in which the detector work is carried out. Chapter 3 presents the Micromegas detector technology, the IAXO-D0 and BabyIAXO detector prototypes, and the associated gas, high-voltage, slow-control, data-acquisition, and calibration systems. Chapter 4 describes the computational framework used in the thesis, including ROOT, REST-for-Physics, restG4, data production, visualization, and related software developments. Chapter 5 studies the shielding and veto system, with emphasis on cosmic-ray-induced backgrounds, passive-shielding limitations, the active scintillator-cadmium veto concept, and the comparison between simulations and prototype data. Chapter 6 presents the broader background model for IAXO-D0 and BabyIAXO, including external and intrinsic background contributions and the simulation methodology used to estimate their impact in the signal region. The final chapter summarizes the results and outlines the next steps toward a complete BabyIAXO detector-background model.

The common objective of these chapters is to show how the detector, software,

simulation, and shielding developments fit together into a single experimental program. In a helioscope, sensitivity is ultimately limited by the ability to convert a rare solar-axion signal into a small, well-characterized X-ray excess. This thesis contributes to that goal by developing the tools and background-rejection strategy needed to make the BabyIAXO Micromegas detector line a quantitatively understood low-background instrument.

---

# Chapter 1

## Axions and Axion-Like Particles

1.1	The Strong CP Problem . . . . .	18
1.2	Axion and Axion-Like Particles (ALPs) . . . . .	19
1.3	Axion Detection . . . . .	22
1.3.1	Haloscopes . . . . .	22
1.3.2	Light Shining Through Wall . . . . .	22
1.3.3	Helioscopes . . . . .	23

### Introduction

The Standard Model of particle physics (SM) is the quantum field theory that describes the fundamental structure of matter and its interactions. This theory exhibits great predictive power, as nearly all experimental measurements of the interactions within the model satisfactorily match the predictions. Currently, evidence exists for the existence of all the fundamental particles described in the model, with the Higgs boson being the last to be found, its discovery announced in 2012 at the LHC.

Despite the model's success, it is far from being a theory of everything. Firstly, the Standard Model does not describe gravitational interaction. It is also unable to explain the nature of dark matter or dark energy. The particles of the Standard Model constitute only 5% of the total matter and energy of the known universe. Additionally, there are multiple unresolved questions, among which we highlight the fine-tuning or naturalness problems.

The Standard Model requires a large number of parameters, apparently set arbitrarily, which must be measured experimentally. A strategy to extend the Standard

Model consists of proposing new physics theories where some of these constants appear naturally, since naturalness is a desirable property in a theory.

One of these fine-tuning problems is known as the *strong CP problem*, due to the extremely small value of the dimensionless parameter  $\theta$  in quantum chromodynamics (QCD), the theory that describes strong interactions. The Peccei-Quinn mechanism is proposed as a solution to this problem and is currently the most accepted theory, although, for the moment, no experimental evidence supports it. An important consequence of the theory is the existence of a new pseudo-scalar particle: the axion.

The axion is a very low-mass particle that must have extremely suppressed interactions with the Standard Model particles to be compatible with observations. One property that makes this particle especially interesting is that it could be produced non-thermally in the early universe and be one of the constituents of dark matter.

It is possible to perform a direct search for the axion and pseudoscalar particles with similar properties, known as axion-like particles (ALPs), through experiments that involve the most experimentally accessible interaction of axions: the axion-photon conversion via the Primakoff effect.

The IAXO experiment has been conceived as a fourth-generation helioscope that will use a superconducting magnet situated on a mobile platform capable of tracking the sun. A small percentage of solar axions would be converted to low-energy X-ray photons during their transit through the intense magnetic field. Subsequently, these photons would be detected in a very low-background X-ray detector located at the other end of the experiment.

As in any rare-event experiment, the number of expected signals is very low and may be hidden by background events. Therefore, it is crucial to understand these backgrounds to work in two directions: designing appropriate shielding to reduce the background, and data analysis techniques to discriminate signal from background.

## 1.1 The Strong CP Problem

The *strong CP problem* arises in the theory of quantum chromodynamics (QCD), the gauge theory that describes one of the three fundamental interactions of the Standard Model, the strong interaction.

It is possible to add an additional term,  $\mathcal{L}_{\bar{\theta}}$ , to the QCD Lagrangian while still obtaining a renormalizable theory

$$\mathcal{L}_{\bar{\theta}} = \bar{\theta} \frac{\alpha_S}{8\pi} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}, \quad (1.1)$$

where  $G_{\mu\nu}^a$  is the gluon field strength tensor, and  $\tilde{G}^{a,\mu\nu} = \varepsilon^{\mu\nu\lambda\rho} G_{\lambda\rho}^a / 2$  is its dual.

Adding this term has measurable physical consequences. In particular, it is the

only term in the Lagrangian that violates CP symmetry. However, no experiment has detected CP violation in the strong interaction to date.

Experimental measurements can set limits on the parameter  $\bar{\theta}$ . The most stringent constraint on  $\bar{\theta}$  comes from the measurement of the neutron electric dipole moment (nEDM). If CP symmetry is not conserved in the strong interaction, a neutron electric dipole moment is predicted [1, 2]

$$|d_n| \sim \bar{\theta} \cdot (6 \times 10^{-17}) e \text{ cm}, \quad (1.2)$$

which is proportional to the parameter  $\bar{\theta}$  that controls CP violation.

The most recent experiments [3] yield

$$d_n = (0.0 \pm 1.1_{\text{stat}} \pm 0.2_{\text{sys}}) \times 10^{-26} e \text{ cm}, \quad (1.3)$$

corresponding to a limit of

$$|\bar{\theta}| \lesssim 1 \times 10^{-10}, \quad (1.4)$$

for an adimensional parameter, an angle, which, by naturalness arguments, would be expected to have a value of order unity.

This fine-tuning, known in the literature as the *strong CP problem*, suggests the existence of new physics, where CP symmetry conservation in the strong interaction would emerge more naturally.

## 1.2 Axion and Axion-Like Particles (ALPs)

In the literature [4, 5], different solutions to the *strong CP problem* are proposed. Some of the most important are the massless quark solution [5], which is ruled out by observations [6]; the soft CP solution [7, 8]; and the axion solution via the Peccei-Quinn mechanism.

### The Peccei-Quinn mechanism

The solution proposed in 1977 by Roberto Peccei and Helen Quinn [9, 10] postulates the existence of an additional global  $U(1)_{\text{PQ}}$  symmetry. This symmetry is spontaneously broken at an energy scale  $f_A$ . The breaking of this symmetry results in the emergence of a new pseudoscalar particle (pseudo-Goldstone boson) called the axion [11, 12], which acquires a very small mass due to the fact that the  $U(1)_{\text{PQ}}$  symmetry is not exact, being explicitly broken by the so-called QCD axial anomaly. This gives rise to the coupling between gluons and the axion field  $\phi_A$ ,

$$\mathcal{L}_{\bar{\theta}} = \left( \frac{\phi_A}{f_A} - \bar{\theta} \right) \frac{g^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a,\mu\nu}. \quad (1.5)$$

Quantum fluctuations of QCD, which are responsible for the  $\theta$  vacuum structure of the theory, ensure that the axion potential is minimized for  $\phi_A = f_A \bar{\theta}$ , thereby acquiring a vacuum expectation value and preserving CP symmetry in the QCD Lagrangian.

The axion mass  $m_A$  can be derived from the QCD topological susceptibility  $\chi$ , since its origin lies in the axial anomaly.

An approximate expression for the axion mass can be obtained as  $m_A f_A = \sqrt{\chi} \approx m_\pi f_\pi$  [13, 14]. Substituting the susceptibility [15], we find

$$m_A = 5.691(51) \left( \frac{10^9 \text{ GeV}}{f_A} \right) \text{ meV}. \quad (1.6)$$

The energy scale parameter  $f_A$  is unknown but must be consistent with experimental results.

### Axion models and axion-like particles (ALPs)

The first proposed model was the Peccei-Quinn-Weinberg-Wilczek (PQWW) axion [9, 11, 12]. In this model, the symmetry breaking occurs at  $f_A \approx v_{\text{EW}} \approx 250 \text{ GeV}$ , on the order of the electroweak scale, which caused the model to be quickly ruled out, for instance, by experiments measuring the branching ratio (BR) of the process  $K^+ \rightarrow \pi^+ + A$  [16].

Observations necessarily impose that  $f_A \gg v_{\text{EW}}$ , leading to extremely small axion masses and couplings, which explains why it has not yet been detected [17, 18]. This is known in the literature as the *invisible axion*.

The two most notable invisible-axion models are the KSVZ model [17, 19], which introduces a new heavy, electrically neutral quark carrying the PQ charge, and the DFSZ model [20, 21], in which two Higgs doublets appear and the PQ charge is carried by Standard Model quarks and leptons. Additionally, more complex models exist that combine ideas from both.

Depending on the model, the axion will have couplings to various Standard Model particles such as electrons, nucleons, or photons. Even if there is no explicit coupling to photons, the axion will inherit such a coupling from the  $\pi^0$  due to mixing states with the neutral mesons  $\pi^0$  and  $\eta$  (as a consequence of 1.5). The axion-to-two-photon conversion ( $A \rightarrow \gamma\gamma$ ) is known as the Primakoff effect and plays a crucial role from an experimental perspective. The interaction term is given by

$$\mathcal{L}_{A\gamma\gamma} = -\frac{1}{4} g_{A\gamma\gamma} F_{\mu\nu} \tilde{F}^{\mu\nu} \phi_A = g_{A\gamma\gamma} \vec{E} \cdot \vec{B} \phi_A, \quad (1.7)$$

where  $F_{\mu\nu}$  is the electromagnetic field tensor and  $\tilde{F}^{\mu\nu}$  its dual. The coupling constant is

$$g_{A\gamma\gamma} = \frac{\alpha_S}{2\pi f_A} \left( \frac{E}{N} - 1.92(4) \right) = \left( 0.203(3) \frac{E}{N} - 0.39(1) \right) \frac{m_A}{\text{GeV}^2}, \quad (1.8)$$

which is proportional to the axion mass and depends on the model through the constants  $E$  and  $N$ , taking  $E/N = 8/3$  for the DFSZ model and  $E/N = 0$  for KSVZ.

In addition to the aforementioned models, many extensions of the Standard Model, such as string theories or supersymmetric theories, predict the existence of very light and weakly interacting particles (WISPs). Some of these particles arise from the spontaneous breaking of an approximate global  $U(1)$  symmetry and exhibit phenomenology similar to the axion described earlier. These are known as axion-like particles (ALPs).

Examples of symmetries producing ALPs include lepton number symmetry [22], R-symmetry [23], or certain symmetries appearing in string theories [24, 25].




Figure 1.1: Parameter space mass-coupling adapted from [26]. The plot shows colored areas (green, blue, and gray) corresponding to different experimental techniques. Solid colors represent existing results, while translucent colors indicate future predictions. The QCD axions are restricted to a diagonal band (yellow) due to the relation between their coupling and mass (equation 1.8). The KSVZ line is highlighted in red.

In the literature, QCD axions (a solution to the *strong CP problem*) and ALPs are often referred to collectively as axions. From a detection perspective, the main difference is that for QCD axions, the mass and interactions are determined by the

energy scale  $f_A$ , leading to a relationship between the coupling and mass (equation 1.8), whereas for ALPs, this relationship does not exist. Axion detection experiments explore regions of the parameter space (figure 1.1), providing limits for both QCD axions and ALPs. In addition to experimental limits, there are other limits from astrophysical or cosmological observations that will not be mentioned in this work.

## 1.3 Axion Detection

According to the properties described above, various natural sources of axions are predicted, as well as the possibility of creating axions in the laboratory. All major experiments share a common goal: detecting the photon produced in the Primakoff conversion of the axion, facilitated by the presence of an intense magnetic field (Equation 1.8), although different techniques employ very different strategies. A comprehensive summary of the main techniques can be found in [26].

The entirety or an arbitrary proportion of the universe's dark matter could be composed of axions. It is believed that 25.9% of the total matter and energy of the universe corresponds to cold dark matter [27], which is distributed in galactic halos. Although axions are particles with extremely small mass, through the non-thermal production mechanism known as the vacuum realignment mechanism, oscillations of the axion field produced at an early stage of the universe would behave on a large scale like cold dark matter, suggesting the presence of axions in our galactic halo.

### 1.3.1 Haloscopes

The haloscope technique [28] uses resonant cavities with a high quality factor placed in very intense magnetic fields to detect axions from the galactic halo. They have the highest sensitivity among all experiments but are limited to measuring, by tuning the cavity frequency, narrow axion mass ranges where the coherence condition is met, as shown in Figure 1.1, which displays results from the ADMX experiment [29]. These experiments are highly model-dependent, as they assume that the entirety, or at least a large proportion, of the dark matter in the galactic halo consists of axions. Additionally, there is significant uncertainty associated with their distribution.

### 1.3.2 Light Shining Through Wall

Through the Primakoff conversion, axion production is possible in the laboratory. Using a source of electromagnetic radiation (primarily lasers but also microwaves) along with very intense magnetic fields, it is possible to produce axions. These axions would pass through a barrier that blocks unconverted photons to reach a new conversion region and produce photons that would be detected. This strategy is known as light shining

through a wall (LSTW) and is the most model-independent since it only uses the fundamental properties of axions. However, it is currently unable to explore parameter space regions not covered by other experiments (see Figure 1.1) due to the power required to produce the necessary photon flux. Two collaborations are working on such experiments: ALPS [30], at DESY, and OSQAR [31], at CERN.

### 1.3.3 Helioscopes

Axions are also emitted by astrophysical objects such as supernovae or stars due to the intense electromagnetic fields within them. This means the Sun is the most intense astrophysical source of axions within reach. Helioscopes use the Primakoff effect to detect axions of solar origin.



---

## Chapter 2

# The IAXO Collaboration and the BabyIAXO Experiment

2.1	The Helioscope Technique . . . . .	27
2.1.1	Solar Axions . . . . .	27
2.1.2	Axion-Photon Conversion in a Helioscope . . . . .	28
2.1.3	Experimental Figure of Merit . . . . .	29
2.2	From CAST to IAXO . . . . .	29
2.2.1	CAST as the Reference Helioscope . . . . .	29
2.2.2	Why a Dedicated Helioscope Is Required . . . . .	30
2.3	The IAXO Experiment . . . . .	31
2.3.1	Physics Reach and Design Philosophy . . . . .	31
2.3.2	Magnet . . . . .	32
2.3.3	X-ray Optics . . . . .	32
2.3.4	Detector Technologies . . . . .	33
2.3.5	Tracking and Observatory Operation . . . . .	33
2.4	The IAXO Collaboration and Project Context . . . . .	34
2.5	BabyIAXO as an Intermediate Stage . . . . .	34
2.5.1	Role Within the IAXO Program . . . . .	34
2.5.2	Evolution of the DESY Site Scenario . . . . .	35
2.5.3	Main Experimental Features . . . . .	37
2.6	Connection to the Present Thesis . . . . .	38

## Introduction

The International Axion Observatory (IAXO) is the next major step in the development of the helioscope technique for the search for solar axions and axion-like particles [32–34]. It builds on the experience accumulated in CAST and replaces the logic of a repurposed installation with a fully purpose-built experiment in which the magnet, the X-ray optics, the detectors, and the solar-tracking system are optimized as parts of a single instrument.

BabyIAXO is the intermediate stage toward that final observatory [35, 36]. It is designed both as a competitive helioscope in its own right and as a technological demonstrator in which the key IAXO subsystems can be validated at a relevant scale. This dual role is especially important for the detector line studied in this thesis, where background reduction, realistic simulation, and system integration must all be demonstrated under surface-level operating conditions.

This chapter first summarizes the helioscope technique and the experimental figure of merit that drives the design. It then describes the transition from CAST to IAXO, the main design choices of the full observatory, and the specific role of BabyIAXO within the broader IAXO program.

## 2.1 The Helioscope Technique

### 2.1.1 Solar Axions




Figure 2.1: Differential solar axion flux at Earth for representative production mechanisms [37]. The Primakoff component depends on the axion-photon coupling, while the ABC components (atomic processes, bremsstrahlung, and Compton-like production) depend on the axion-electron coupling. The normalization of each component is coupling-dependent; the plotted curves correspond to the benchmark normalizations used in the cited solar-flux calculation.

Solar axions are produced in the hot and dense plasma of the Sun through several processes. The best known is Primakoff conversion, in which thermal photons convert into axions in the electromagnetic fields of the plasma constituents. Additional contributions arise from processes involving the axion-electron coupling, notably axio-recombination, axio-deexcitation, bremsstrahlung, and Compton-like scattering, often grouped together as the ABC channels [38]. Figure 2.1 shows the corresponding spectral components in the keV range relevant for helioscope searches.

This solar source is particularly attractive from an experimental perspective. First, the expected axion energies lie in the soft X-ray range, where efficient focusing optics and highly specialized low-background detectors are available. The Primakoff spectrum peaks around a few keV, while the electron-coupling channels have a relatively softer contribution, which motivates low-threshold detector options in addition to Micromegas, such as GridPix and cryogenic sensors. Second, the solar axion flux is large enough that helioscopes can probe parameter space beyond previous laboratory searches. Finally,

helioscopes are directly sensitive to the axion-photon coupling through the inverse Primakoff process and can also probe scenarios involving the axion-electron coupling through the product  $g_{ay}g_{ae}$ .

## 2.1.2 Axion-Photon Conversion in a Helioscope




Figure 2.2: Conceptual diagram of a helioscope. Solar axions enter a magnetic conversion region mounted on a mobile platform that tracks the Sun. The converted X-ray photons are then focused by X-ray optics onto a low-background detector.

The basic working principle of a helioscope is illustrated in Figure 2.2. Solar axions enter a strong transverse magnetic field, where they can convert into X-ray photons through the inverse Primakoff effect. Those photons are subsequently focused by grazing-incidence optics onto a detector optimized for the 1–10 keV energy range. The entire system is mounted on a platform capable of following the Sun during the daily observation window.

For a homogeneous magnetic field of length  $L$ , the conversion probability can be written as

$$P_{a \rightarrow \gamma}(L) = \left( \frac{g_{ay}BL}{2} \right)^2 \mathcal{F}(qL), \quad (2.1)$$

where  $B$  is the magnetic field intensity and

$$\mathcal{F}(qL) = \left( \frac{2}{qL} \right)^2 \sin^2 \left( \frac{qL}{2} \right) \quad (2.2)$$

is the coherence factor. In vacuum,  $q \simeq m_a^2/2E$  for relativistic axions of energy  $E$ . The coherent regime corresponds to  $qL \ll 1$ , in which case  $\mathcal{F} \simeq 1$  and the conversion

probability grows as  $B^2L^2$ .

This coherence condition determines the mass range that can be explored in vacuum. At sufficiently large axion mass, the momentum mismatch between the axion and the photon suppresses the conversion probability. As in CAST, this coherence can be partially restored by introducing a buffer gas in the magnet bores, which gives the photon an effective mass and extends the accessible mass range. In a buffer gas, the photon acquires an effective mass  $m_\gamma$ , so  $q \simeq |m_a^2 - m_\gamma^2|/(2E)$ , with additional damping from photon absorption. Equations 2.1 and 2.2 are therefore the vacuum, negligible-absorption limit [39–41].

### 2.1.3 Experimental Figure of Merit

The helioscope signal rate is not determined by the magnetic conversion region alone. It also depends on the focusing efficiency of the optics, the background level and efficiency of the detector, and the fraction of time during which the instrument can track the Sun. These dependencies are conveniently summarized by the helioscope figure of merit [32, 34]

$$f = f_M f_{DO} f_T, \quad (2.3)$$

with

$$f_M = B^2 L^2 A, \quad f_{DO} = \frac{\epsilon_d \epsilon_o}{\sqrt{ba}}, \quad f_T = \sqrt{\epsilon_t} t. \quad (2.4)$$

Here  $A$  is the bore cross-sectional area,  $\epsilon_d$  is the detector efficiency,  $\epsilon_o$  is the optics efficiency,  $b$  is the normalized detector background,  $a$  is the focal spot area on the detector,  $\epsilon_t$  is the tracking efficiency, and  $t$  is the total data-taking time.

This decomposition is particularly useful because it makes explicit how a next-generation helioscope should be designed. The magnetic term favors large aperture in addition to field strength and length, while the detector-optics term rewards efficient focusing onto a very small spot and an ultra-low background readout. The tracking term shows that a mobile platform with long daily observation time is an integral part of the physics reach. This FOM is a background-dominated sensitivity proxy; since the Primakoff signal scales as  $g_{a\gamma}^4$ , improvements in  $f$  translate into coupling reach approximately as  $g_{a\gamma} \propto f^{-1/4}$ . These are precisely the principles that motivated the transition from CAST to IAXO.

## 2.2 From CAST to IAXO

### 2.2.1 CAST as the Reference Helioscope

CAST established the helioscope technique as the leading laboratory search for solar axions [42]. It used a repurposed LHC test dipole magnet together with several

low-background X-ray detector systems on its four magnet bores. One line was equipped with focusing X-ray optics and a pn-CCD detector, and later IAXO-pathfinder configurations explored Micromegas operation with focused optics, passive shielding, active vetoing, and upgraded readout [43–45]. Over its successive vacuum,  $^4\text{He}$ , and  $^3\text{He}$  running periods, CAST demonstrated both the maturity of the helioscope concept and the practicality of restoring coherence with a buffer gas to extend the accessible axion-mass range [39–42].

The importance of CAST for IAXO is not limited to the exclusion limits it produced. CAST showed that very low detector backgrounds are achievable in helioscope conditions, that X-ray optics materially improve the signal-to-background ratio, and that long-term operation of a Sun-tracking magnet-detector system is feasible. In that sense, CAST provided the full experimental foundation on which the IAXO program was built.

## 2.2.2 Why a Dedicated Helioscope Is Required

At the same time, CAST made the main limitations of a repurposed installation evident. Its magnet was optimized for accelerator use rather than helioscope physics, which implies a small aperture, restricted geometrical freedom for optics and detectors, and limited tracking time. These are not secondary limitations. Because the helioscope figure of merit scales linearly with bore cross-sectional area  $A$  and strongly with the detector-optics term, a step change in sensitivity requires a dedicated experimental design rather than incremental upgrades of an existing accelerator component.

Parameter	CAST	BabyIAXO	IAXO
Magnet concept	repurposed LHC dipole	purpose-built dipole	purpose-built toroid
Field at bores $B$	9 T	$\sim 2$ T	$\sim 2.5$ T
Magnetic length $L$	9.26 m	$\sim 10$ m	$\sim 20$ m
Bore cross-section $A$	$\sim 0.003\text{ m}^2$	$\sim 0.77\text{ m}^2$	$\sim 2.3\text{ m}^2$
Representative $f_M$	$\sim 21\text{ T}^2\text{ m}^4$	$\sim 230\text{ T}^2\text{ m}^4$	$\sim 6000\text{ T}^2\text{ m}^4$
Detector-background target $b$	$\sim 10^{-6}$	$\sim 10^{-7}$	$\sim 10^{-8}$
Focal-spot scale $a$	focused line only	$\sim 0.2\text{ cm}^2$	$\sim 0.2\text{ cm}^2$
Tracking efficiency $\epsilon_t$	$\sim 0.12$	$\sim 0.5$	$\sim 0.5$

Table 2.1: Representative comparison of the main helioscope design parameters for CAST, BabyIAXO, and IAXO, compiled from CAST publications, the IAXO and BabyIAXO conceptual designs, and collaboration summary material [34, 35, 42, 46]. The normalized background  $b$  is given in counts  $\text{keV}^{-1}\text{ cm}^{-2}\text{ s}^{-1}$ ; the CAST value is indicative of the low-background Micromegas/pathfinder scale rather than a single detector configuration.

The comparison in Table 2.1 clarifies the design shift from CAST to the IAXO program. CAST maximized field strength by reusing a high-field accelerator magnet, whereas BabyIAXO and IAXO trade some field intensity for a much larger aperture,

systematic use of focusing optics, and substantially longer daily tracking. Because  $f_M$  grows linearly with bore cross-sectional area  $A$ , this change is decisive. The strategy of IAXO is therefore not to reproduce CAST at larger scale, but to optimize the full helioscope figure of merit in a balanced and purpose-built way [32–34].

## 2.3 The IAXO Experiment

### 2.3.1 Physics Reach and Design Philosophy

IAXO is conceived as a fourth-generation axion helioscope optimized for solar axions and axion-like particles [33, 34]. In the low-mass region, the full observatory is designed to improve the CAST signal-to-background ratio by approximately four to five orders of magnitude and to reach sensitivity to axion-photon couplings down to a few  $\times 10^{-12} \text{ GeV}^{-1}$ . In addition to Primakoff solar axions, it is also expected to probe scenarios involving the axion-electron coupling with sensitivity beyond previous laboratory experiments.




Figure 2.3: Axion-photon parameter space adapted from [26]. Existing limits and projected future sensitivities are shown for several experimental techniques. Helioscopes occupy the low-mass region of this plane; BabyIAXO begins to enter the region populated by benchmark QCD axion models, while IAXO extends the projected helioscope reach substantially deeper into unexplored parameter space.

This projected reach is illustrated in Figure 2.3. In the language of the axion-photon parameter space, BabyIAXO is expected to begin probing the region where benchmark QCD axion models appear, while IAXO aims at a much broader advance through the unexplored low-mass domain. The essential design choice that makes this possible is that the entire experiment is optimized for the helioscope figure of merit. Instead of maximizing only the field strength, IAXO emphasizes aperture, simultaneous instrumentation of multiple bores, small focal spots, long tracking time, and detector backgrounds at the level required for rare-event searches. In this sense, IAXO should not be understood simply as a larger CAST, but as a dedicated observatory built around the needs of helioscope physics.

### 2.3.2 Magnet

The core of IAXO is a purpose-built superconducting toroidal magnet approximately 20 m long, formed by eight coils and providing eight bores of 60 cm diameter each [34]. The field at the bores is of order a few tesla, with peak values above 5 T, but the central performance gain with respect to CAST comes from the large aperture rather than from field strength alone. This geometry is one of the defining differences with respect to previous helioscopes.

The toroidal layout has several advantages for a helioscope. It provides a large total conversion volume, leaves open and accessible bores for the insertion of optics and detectors, and naturally defines multiple parallel detection lines. In addition, the bores can be operated in vacuum or filled with a buffer gas when coherence recovery at higher axion mass is required. The magnet is therefore not simply a source of field, but the structural element that organizes the entire observatory.

### 2.3.3 X-ray Optics

The second key ingredient is the systematic use of focusing X-ray optics on every helioscope line [34, 47]. In IAXO, each bore is foreseen to be equipped with dedicated grazing-incidence optics that focus the converted photons onto spots of order  $0.2 \text{ cm}^2$ . This is a decisive improvement over CAST, where focusing optics were available only on part of the setup.

The function of the optics is not merely to image the source, but to compress the signal onto a very small detector area. Since the detector background roughly scales with the active area exposed to it, focusing the expected axion signal onto a compact spot directly improves the detector-optics figure of merit. This coupling between optics and low-background detectors is one of the defining features of the IAXO concept.

### 2.3.4 Detector Technologies

The low-background detectors located at the focal planes must combine high efficiency in the 1–10 keV region with stage-dependent background targets. BabyIAXO aims at the  $10^{-7}$  counts keV $^{-1}$  cm $^{-2}$  s $^{-1}$  scale, while the full IAXO concept requires a further reduction toward  $10^{-8}$  counts keV $^{-1}$  cm $^{-2}$  s $^{-1}$  [34, 35]. For a surface helioscope, reaching even the BabyIAXO target is nontrivial because the detector must reject cosmic-ray-induced backgrounds without relying on underground overburden. This requirement has driven a broad detector-development program within the collaboration [34, 35]. Several technologies are under study, including Micromegas, GridPix, metallic magnetic calorimeters, transition-edge sensors, and silicon drift detectors.

Among these options, microbulk Micromegas constitute the baseline gaseous detector technology because of their demonstrated performance in CAST and their combination of radiopurity, topological discrimination power, stable operation, and compatibility with focused soft X-ray signals. At the same time, the BabyIAXO detector program has broadened to include other specialized concepts, in particular GridPix-based detectors, which aim at very fine-grained event reconstruction under the same low-background constraints [48, 49]. The detailed implementation of the Micromegas detector line addressed in this thesis is presented in the following chapters.

### 2.3.5 Tracking and Observatory Operation

IAXO is designed as a true observatory rather than a static test setup. The magnet, optics, and detector assembly are mounted on elevation and azimuth drives that allow solar tracking for up to about half of each day, corresponding to a tracking efficiency  $\epsilon_t \simeq 0.5$  in the conceptual design [34]. This is a very substantial gain with respect to CAST and is one of the reasons why the tracking system enters explicitly into the helioscope figure of merit.

The corresponding mechanical and cryogenic requirements are nontrivial. The structure must support a large cold mass, preserve alignment between bores, optics, and detectors during motion, and allow a reliable transition between Sun-tracking periods and off-Sun background measurements. These operational constraints are part of the experimental design, not merely engineering afterthoughts.

This has immediate consequences for the detector program. Unlike many rare-event experiments, a helioscope of this scale cannot simply be placed deep underground to suppress cosmic radiation. The expected installation is at surface level, and the required background suppression must therefore be achieved through detector design, passive shielding, topology-based analysis, and active veto systems. This constraint is central to the present thesis.

## 2.4 The IAXO Collaboration and Project Context

IAXO is inherently a collaboration-driven experiment. Its realization requires the integration of expertise from axion phenomenology, superconducting magnet engineering, X-ray optics, cryogenics, detector development, low-background techniques, and data acquisition. This breadth is already evident in the Letter of Intent and the conceptual-design documents, which frame IAXO as an observatory assembled from mature but previously separate technological lines [33, 34].

The staged realization through BabyIAXO reflects both physics logic and project logic. A medium-scale helioscope allows the magnet, optics, detector infrastructure, gas handling, alignment, and tracking systems to be commissioned together before scaling to the full toroidal observatory. In this sense, BabyIAXO is the bridge between subsystem R&D and the final experiment, and not simply a reduced-scale prototype [35, 49]. For the detector work of this thesis, the collaboration structure is not incidental: magnet geometry, optical focal-spot assumptions, readout constraints, radiopurity screening, and veto integration are defined across different working groups. The simulation and background-model tasks therefore act as an interface between detector development, mechanical integration, and physics sensitivity.

## 2.5 BabyIAXO as an Intermediate Stage

### 2.5.1 Role Within the IAXO Program

BabyIAXO was proposed as the intermediate stage between CAST and the full IAXO observatory [35, 36]. Its purpose is twofold. On the one hand, it provides a realistic environment in which the main IAXO subsystems can be integrated and validated: magnet, optics, detectors, cryogenics, tracking, alignment, and data acquisition. On the other hand, it is itself a fully-fledged helioscope with nontrivial discovery potential.

In the baseline vacuum phase, BabyIAXO is expected to reach sensitivities of order  $g_{a\gamma} \sim 1.5 \times 10^{-11} \text{ GeV}^{-1}$  for masses up to  $m_a \sim 2 \times 10^{-2} \text{ eV}$ . With a buffer-gas phase, the accessible mass range can be extended, allowing BabyIAXO to probe the KSVZ benchmark region approximately in the 0.06–0.25 eV interval [35]. In practical terms, this places BabyIAXO well beyond a mere engineering prototype. As indicated in Figure 2.3, it already begins to explore physically relevant parameter space associated with benchmark QCD axion models while simultaneously de-risking the construction of the full observatory.




Figure 2.4: Conceptual view of the BabyIAXO platform. The experiment combines a superconducting magnet, X-ray optics, low-background detectors, and a movable support structure able to track the Sun.

## 2.5.2 Evolution of the DESY Site Scenario

The site assumptions for BabyIAXO evolved during the period in which the work presented in this thesis was being developed. At the beginning of the background-model and veto-simulation program, the reference implementation was the HERA South Hall at DESY. This was an underground accelerator hall rather than a deep-underground low-background laboratory, but it nevertheless implied a different environmental boundary condition from a fully outdoor surface installation: the hall structure, access shafts, and surrounding infrastructure could modify the cosmic-ray field seen by the detector. For this reason, early simulation work treated the cosmic-ray background as an important but site-dependent contribution, with attention to overburden, openings, and local shielding details.

During 2025 and 2026 the working site assumptions shifted in recent internal collaboration discussions toward an outside, on-surface location at DESY. The main arguments reported in the collaboration meetings were practical and programmatic: reduced civil-infrastructure effort, less dependence on overstretched DESY infrastructure groups, lower expected cost, and a faster route to site activation [50]. The change also has a physics and engineering advantage because it makes BabyIAXO closer to the expected surface conditions of the full IAXO observatory. By the 23rd IAXO Collaboration Meeting, the surface scenario had become the current internal working

baseline for technical planning, and the DESY directorate had encouraged further site exploration and cost/schedule assessment, although formal full project and site approval still depended on the complete financial and construction review [51].




Figure 2.5: Adapted view, based on collaboration-meeting material, comparing the previous HERA South Hall context with the newer on-surface BabyIAXO site scenario at DESY [52]. The surface option changes the environmental assumptions relevant for detector-background simulations: the experiment should be treated as a surface helioscope rather than as an installation benefiting from underground-laboratory overburden.

This change is directly relevant for the interpretation of the simulations in this thesis. The initial motivation included the question of whether the HERA South environment would provide enough passive reduction of cosmic-ray-induced backgrounds for a low-background Micromegas line. The shift toward an on-surface working baseline strongly motivates treating BabyIAXO background studies as surface-level studies, even when earlier simulations were developed with HERA South Hall boundary conditions in mind. The detector system must therefore be robust against sea-level muons, high-energy neutrons, and the secondary showers produced in the passive shielding. Consequently, the active scintillator–cadmium veto described in Chapter 5 becomes a central design element for the surface-detector concept rather than a secondary upgrade.

The outdoor site also introduces engineering constraints that feed back into detector-background studies. Recent site studies emphasize the larger outdoor turning circle, magnetic stray-field limits, the need for a fenced exclusion region of order 35 m diameter for BabyIAXO, and checks of possible interference with nearby DESY magnet-test activities [53]. In addition, the support and drive system must be validated for outdoor

load cases such as wind, snow, ice, and temperature gradients. These developments do not invalidate the earlier simulation program; rather, they make the surface-cosmic component and the active-veto strategy the conservative and experimentally relevant reference for BabyIAXO.

### 2.5.3 Main Experimental Features

The BabyIAXO design follows the same experimental logic as IAXO, but at reduced scale [35]. Its magnet comprises two 10 m long bores of 70 cm diameter, each intended to host a complete detection line with optics and detector dimensions representative of the final observatory. The superconducting system is based on two parallel flat coils and conventional NbTi/Cu Rutherford cable technology, with the cold mass operated at about 4.5 K. In this way, the magnet aperture, cryogenic integration, and detector interfaces are all tested under realistic conditions before the transition to the larger toroidal concept.




Figure 2.6: Conceptual rendering of the BabyIAXO magnet. The two large bores allow the simultaneous deployment of two helioscope detection lines, making BabyIAXO both a subsystem prototype for IAXO and a stand-alone physics instrument.

The optics and detector strategy of BabyIAXO is also deliberately close to that of IAXO. Dedicated X-ray optics based on multilayer-coated segmented-glass Wolter-I concepts focus the expected signal from the full bore onto a spot of order  $0.2 \text{ cm}^2$  [35]. The present BabyIAXO plan includes two telescope solutions with different focal

lengths, one custom design and one based on an available XMM-Newton spare optic, both serving as technology and integration demonstrators for the final observatory.

On the detector side, Micromegas remain the baseline low-background technology for the helioscope line studied in this thesis, but the broader BabyIAXO program also includes GridPix and cryogenic detector developments [35, 49]. This diversification is scientifically valuable because different detector technologies probe complementary aspects of performance, such as energy threshold, topology, time structure, or energy resolution. The detector-side roadmap is also informed by the IAXO pathfinder program carried out at CAST, where an IAXO-oriented Micromegas line was already operated together with focused optics, passive shielding, an active muon veto, and AGET-based readout in a surface installation [45]. In parallel, the BabyIAXO infrastructure has motivated auxiliary axion-search concepts, including haloscope programs based on radiofrequency cavities integrated with the available magnet geometry [54].

The conceptual design places BabyIAXO at DESY in Hamburg [35]. The recent shift toward an outside surface working baseline, discussed in Section 2.5.2, makes the operating conditions especially relevant for the present work: a large moving magnet, full detector infrastructure, and no possibility of relying on underground shielding against cosmic rays. As a consequence, background modeling and active rejection become enabling technologies rather than secondary refinements. This point is emphasized repeatedly in recent BabyIAXO detector-development work, where veto design, material selection, and integration constraints are treated as central elements of the experiment rather than add-ons [49].

## 2.6 Connection to the Present Thesis

The chapters that follow focus on one detector line within the broader IAXO program: the Micromegas-based BabyIAXO line, together with its software and background-reduction strategy. The need for realistic simulation, a quantitative background model, and an optimized active veto follows directly from the IAXO and BabyIAXO design choices described above.

The central challenge is the combination of focused keV X-ray detection, ultra-low background performance, and operation in a large surface-level moving apparatus. This thesis addresses that challenge through the Micromegas detector line, the `REST-for-Physics/restG4` simulation chain, the background model, and the scintillator-cadmium veto strategy.

---

## Chapter 3

# Micromegas X-ray Detectors for BabyIAXO

3.1	The Time Projection Chamber (TPC) . . . . .	40
3.1.1	Ionization in Gases . . . . .	41
3.1.2	Electron Transport . . . . .	47
3.1.3	Micropattern Gaseous Detectors: Micromegas . . . . .	53
3.2	The BabyIAXO Micromegas Detector Prototypes . . . . .	55
3.2.1	IAXO-D0 . . . . .	56
3.2.2	IAXO-D1 . . . . .	57
3.3	Detector Ancillary Systems . . . . .	62
3.3.1	Gas System . . . . .	62
3.3.2	High-Voltage System . . . . .	64
3.3.3	Slow-Control System . . . . .	65
3.4	Shielding and Veto Systems . . . . .	65
3.5	Data Acquisition System . . . . .	65
3.5.1	Data Acquisition Software . . . . .	69
3.6	Detector Calibration . . . . .	74
3.6.1	Standard x-ray calibration . . . . .	75
3.6.2	UV-light calibration R&D . . . . .	78

## Introduction

Particle detectors are devices used for the detection of ionizing particles. The information from the detector can be used to obtain physical properties of the particle such as energy, momentum, charge, particle identification, etc.

An ionizing particle interacts with the target media which depending on the type of media and incoming particle can range from ionization to scintillation light, Cherenkov radiation, or phonon production.

Gaseous particle detectors primarily rely on the ionization of the gas medium for detection. Following ionization, electrons generated are guided through electric fields, often amplified, and subsequently, the resulting charge induces a measurable current.

### 3.1 The Time Projection Chamber (TPC)

Time Projection Chambers (TPCs) are advanced gaseous detectors used extensively in particle physics and nuclear experiments for their exceptional ability to measure the three-dimensional trajectories of ionizing particles. Since their inception in the late 1970s, TPCs have become invaluable in tracking and identifying particles, thanks to their high spatial resolution, large sensitive volumes, and capacity to reconstruct particle momentum. A TPC works by detecting the drift of ionized electrons generated by charged particles traversing a gaseous medium, enabling precise position measurements. The design and versatility of TPCs make them particularly suitable for experiments seeking detailed tracking over extensive volumes, such as those studying rare events or exotic particles.

A TPC typically consists of a gas-filled chamber subjected to a uniform electric field, which guides ionized electrons from particle interactions toward a readout plane. The chamber walls are often covered with high-precision sensors, enabling the reconstruction of particle trajectories. As a charged particle passes through the gas, it ionizes the molecules along its path, creating electron-ion pairs. The electrons drift under the influence of the electric field toward a readout plane, where their positions and arrival times are measured. This drift allows the reconstruction of both the position and timing information, which is crucial for 3D tracking of particle trajectories.

The readout mechanism, crucial to a TPC's performance, traditionally employs multi-wire proportional chambers (MWPCs) but has evolved to incorporate newer technologies, such as micromegas (MICRO-MEsh Gaseous Structure) detectors. Micromegas technology has brought significant improvements to TPCs, offering higher resolution and greater stability under high particle fluxes.




Figure 3.1: Schematic representation of a Time Projection Chamber (TPC). The ionization of the gas medium by an incoming particle results in the generation of electrons. The electrons drift toward the anode under the influence of an electric field  $E_d$  and are amplified upon reaching the amplification region where a higher electric field  $E_a$  is applied.

### 3.1.1 Ionization in Gases

#### Photons




Figure 3.2: Mass attenuation coefficient for argon as a function of energy. Data obtained from [55].

Photons interact with matter via electromagnetic processes. Unlike charged particles the interaction of photons with matter is not continuous, but discrete. The probability of each process depends on the photon energy and the density and atomic number of the target material. The main processes are: Rayleigh scattering, photoelectric effect, Compton scattering, and pair production.

The intensity  $I$  of a beam of photons after crossing a distance  $x$  in a material is given by

$$I(x) = I_0 e^{-\mu x} \quad (3.1)$$

where  $\mu$  is the linear attenuation coefficient with units of  $[L^{-1}]$ .  $\mu$  may be also expressed as  $\mu_m \rho$  with  $\mu_m$  the mass attenuation coefficient with units of  $[L^2/M]$  and  $\rho$  the density of the material. Figure 3.2 shows  $\mu_m$  as a function of energy for argon, a commonly used gaseous detector medium. At low energies the photoelectric effect is dominant, at intermediate energies Compton scattering is dominant and at high energies pair production dominates. The energy range at which each process is dominant depends on the material but this pattern remains for all materials.

### **Rayleigh scattering**

Rayleigh scattering, also known as coherent scattering, is the elastic scattering of photons by bound electrons. The probability of Rayleigh scattering decreases rapidly with the photon energy as can be seen in figure 3.2. This process is more probably at low energies, however, it does not have a direct effect on the energy spectrum since the scattered photon has the same energy as the incident photon.

## Photoelectric effect




Figure 3.3: Photoelectric attenuation coefficient for noble gases as a function of energy (left) and total attenuation coefficient for some saturated hydrocarbons, commonly used in gaseous mixtures as quenchers, with argon for comparison (right). Data from [55] evaluated at a temperature of 25 °C and a pressure of 1 bar.

A photon in or above the X-ray region can be absorbed by an atom or molecule, transferring its energy to one or more transitions in the electron shells. This process can only happen if the energy of the photon is greater than the binding energy of a given shell. The probability of a given transition is highest just above the binding energy of the shell and then decreases as the energy of the photon increases. As the photon energy increases, the probability of the photoelectric effect decreases; however, new transitions become available, leading to a characteristic pattern in the cross-section as a function of energy with the appearance of the so-called shell edges.

The photoelectric cross section  $\sigma_{\text{phot}}$  as a function of photon energy can be analytically approximated [56]

$$\sigma_{\text{phot}} \propto \frac{Z^5}{E_\gamma^{3.5}} \quad (3.2)$$

to be proportional to a power of the atomic number  $Z$  and inversely proportional to a power of the photon energy  $E_\gamma$ . Gaseous detectors tend to use high  $Z$  noble gases such as Ne, Ar or Xe to increase the cross section. The increase in cross section with  $Z$  from equation 3.2 can be seen in figure 3.3 (left). Gaseous mixtures usually incorporate a small fraction of some quencher gas which is commonly a saturated hydrocarbon. Its contribution to the mixture's cross section is usually small due to the small fraction

and low  $Z$  of the quencher as can be seen in figure 3.3 (right) for some commonly used quenchers.

After the photoelectric absorption, the atom is left in an excited state and will eventually decay to its ground state via one of two mechanisms: fluorescence or Auger effect.

Fluorescence occurs when an electron from an outer shell fills the vacancy left by the ejected electron. An X-ray photon is emitted in the process, with energy equal to the difference in energy between the two shells. The emitted photon can interact with the gaseous medium and be completely absorbed, leading to a full photon energy deposition in the detector. It can also escape the detector, leading to a partial energy deposition. Since the energy of the emitted photon is characteristic of the atom, an escape event produces a peak in the energy spectrum known as the escape peak.

The Auger effect occurs when an electron from an outer shell fills the vacancy left by the ejected electron, and the energy difference is transferred to another electron in the atom. This electron, known as Auger electron, is then ejected from the atom, which can cause further electrons to be ejected by the same process.

The fraction of de-excitation events that lead to the emission of a photon is known as the fluorescence yield and on average increases with  $Z$ , ranging from values close to zero for low atomic numbers to values close to one for high  $Z$  [57]. For reference, the fluorescence yield of argon is approximately between 0.07 and 0.14 [58] while in xenon it is close to 1 [58].

### Compton scattering




Figure 3.4: Compton scattering diagram depicting the scattering of a photon by an electron.

Compton scattering is the scattering of a photon via the interaction with a charged particle, usually a bound electron in an atom. A diagram of the process is shown in figure 3.4. In contrast with Rayleigh scattering, Compton scattering does not conserve the energy of the photon and is therefore a form of inelastic scattering. Compton scattering is most probable at energies much higher than the typical electron bound energy (figure 3.2) so the bound electron can be approximated as a free particle.

The energy of the scattered photon  $E'_\gamma$  can be calculated from the energy of the incident photon  $E_\gamma$  and the scattering angle  $\theta_\gamma$  using the formula

$$E'_\gamma = \frac{1}{1 + \frac{E_\gamma}{m_e c^2} (1 - \cos \theta_\gamma)}. \quad (3.3)$$

The energy ( $E_e$ ) and scattering angle ( $\theta_e$ ) of the scattered electron can be calculated from the conservation of energy and momentum.

The differential cross section is given by the Klein–Nishina formula [59, 60] and for unpolarized photons is given by

$$\frac{d\sigma_{\text{compton}}}{d\theta_\gamma} = \pi r_e^2 \left(\frac{E'_\gamma}{E_\gamma}\right)^2 \left[ \frac{E'_\gamma}{E_\gamma} + \frac{E_\gamma}{E'_\gamma} - \sin^2 \theta_\gamma \right] \quad (3.4)$$

where  $r_e \approx 2.82$  fm is the classical electron radius. Polarized photons also have a dependence to the azimuth angle  $\phi$ .

The energy transfer to the electron becomes smaller as the photon energy decreases. In the limit of low photon energy equation 3.4 becomes

$$\frac{d\sigma_{\text{compton}}}{d\theta_\gamma} \approx \pi r_e^2 (1 + \cos^2 \theta_\gamma) \quad \text{when } E_\gamma \ll m_e c^2 \quad (3.5)$$

known as Thomson scattering, a form of coherent scattering.




Figure 3.5:  $E'_\gamma/E_\gamma$  (equation 3.3) for different photon scattering angles and energies (left) and the differential cross section (equation 3.4) (right). In the low-photon-energy limit, the energy transfer to the electron approaches zero and the differential cross section becomes the Thomson scattering formula (equation 3.5).

## Pair production

In the present context, pair production is the conversion of a photon into an electron-positron pair, although technically speaking it may also refer to the creation of a particle and its antiparticle from a neutral boson.

Due to conservation of energy the photon energy must be greater than the combined rest mass energy of the electron-positron pair ( $2m_e c^2 = 1.022$  MeV). Conservation of momentum prevents this process from taking place in free space; another body, usually an atomic nucleus, must receive the excess momentum.

## Charged Particles

A charged particle traversing a medium will continuously lose some of its energy resulting in ionization of the medium. The main mechanism for energy loss is the interaction of the Coulomb field of the particle with the fields of the molecules in the medium [57] known as Coulomb scattering. The rate at which the particle loses energy in a medium is known as the stopping power and depends on the nature of the particle and the medium, as well as the kinetic energy of the particle.

Besides Coulomb scattering, another relevant mechanism for energy loss is bremsstrahlung, where the charged particle emits electromagnetic radiation as its slowed down by the medium. This process becomes significant when the particle kinetic energy is much higher than its rest mass energy and when the when the particle is slowed down rapidly like in a high density and high atomic number medium.

For the present discussion, bremsstrahlung is only relevant for electrons. In contrast to Coulomb scattering, the ionization may take place far away from the particle trajectory, for example, an electron reaching the copper gas vessel might produce a bremsstrahlung photon that can ionize the gas.




Figure 3.6: Stopping power of muons in copper as a function of the muon momentum. From [61]

## Neutral Particles

TODO.

### 3.1.2 Electron Transport

#### Primary Charge Production

The ionization of the gas medium by an incoming particle results in the generation of electrons and ions. The number of electron-ion pairs generated is proportional to the energy deposited by the particle in the medium. These electrons, referred to here as *primary charge*, are then drifted toward the anode under the influence of an electric field  $E_d$  and are amplified upon reaching the amplification region, where a higher electric field  $E_a$  is applied. After amplification, the charge is collected by the anode and the signal is read out. The ions drift toward the cathode, where their charge is neutralized by electrons from the cathode.

The average number of electrons  $N_e$  produced by a ionization event of energy  $E$  is given by

$$N_e = \frac{E}{W} \quad (3.6)$$

where  $W$  is the work function of the gas medium, which is the average energy required to produce an electron-ion pair. The work function is a material property and depends on the properties of the gas medium, such as the composition, pressure, temperature, and electric field.

The number of electrons produced  $N_e$  follows a probability distribution with variance

$$\sigma_e^2 = FN_e \quad (3.7)$$

where  $F$  is the Fano factor [62], a dimensionless constant that depends on the gas medium and the energy of the ionizing particle. The Fano factor accounts for the statistical fluctuations in the number of electrons produced. The Fano factor has values typically between 0.1 and 0.3 for noble gases.

### Electron Drift

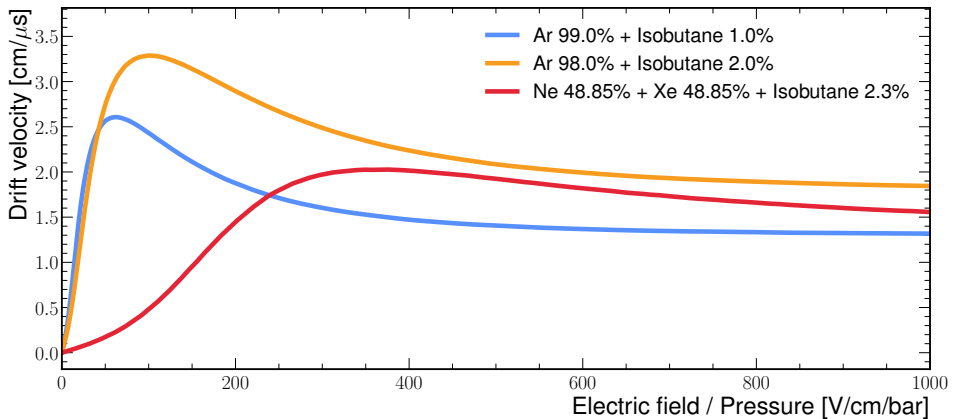


Figure 3.7: Electron drift velocity for commonly used gas mixtures as computed by `Garfield++` [63]. Gas temperature is 20 °C.

The electrons produced by the ionization are drifted toward the anode under the influence of an electric field  $E_d$ . The electrons are accelerated by the electric field and suffer collisions with the gas molecules, which results in a random walk of the electrons. The drift velocity of the electrons  $\vec{v}_d$  is given by

$$\vec{v}_d = \mu \frac{\vec{E}}{p} \quad (3.8)$$

where  $\vec{E}$  is the electric field,  $p$  is the pressure of the gas, and  $\mu$  is the electron mobility. The electron mobility is a material property that depends on the gas medium

and the electric field. The typical values of drift velocity for electrons in noble gases TPCs are in the range of  $1 \text{ cm}/\mu\text{s}$ .




Figure 3.8: Diagrams illustrating electron drift, diffusion (a), and amplification (b).

On average the electrons drift toward the anode at  $\vec{v}_d$ , however, due to the random walk of the electrons, the electrons will spread out in a process known as electron diffusion. The standard deviation  $\sigma_{\hat{u}}$  of the electron cloud in a given direction  $\hat{u}$  after a given drift distance  $d$  is given by

$$\sigma_{\hat{u}} = D_{\hat{u}} \sqrt{d} \quad (3.9)$$

where  $D_{\hat{u}}$  [ $\text{L}^{1/2}$ ] is the diffusion coefficient in the direction  $\hat{u}$ . Due to the anisotropy produced by the electric field, the diffusion coefficient is different in the direction parallel to the electric field  $D_l$  (longitudinal diffusion) and perpendicular to the electric field  $D_t$  (transverse diffusion). Typically the longitudinal diffusion is much smaller than the transverse diffusion as shown in figure 3.8a.

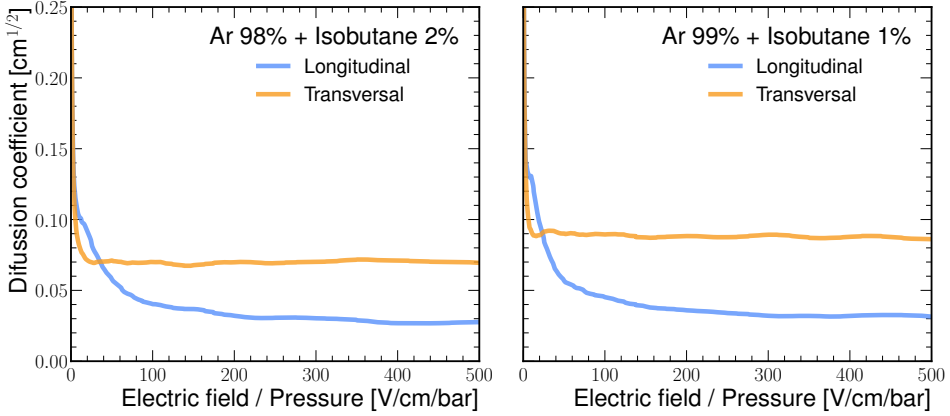


Figure 3.9: Diffusion coefficients for commonly used gas mixtures as computed by `Garfield++` [63]. Gas temperature is 20 °C.

The drift electric field should be strong enough to ensure that the electrons produced do not recombine with the ions, in a process called recombination. These electrons can also be captured by electronegative molecules in the gas medium, in a process called electron attachment, also resulting in a loss of charge. Trace amounts of impurities, most notably oxygen and water vapor, can significantly reduce the number of electrons that reach the anode. The resulting ions, such as  $O_2 + e^- \rightarrow O_2^-$ , can potentially drift toward the anode and produce a signal; however, these ions are usually short-lived and drift significantly slower than the electrons, which effectively means that the signal produced by the negative ions is lost.

Besides the problem of attachment, impurities also affect the properties of the gas medium and the lifetime of the detector. It is therefore critical to maintain the purity of the gas medium by sourcing high-purity gases, using detector materials with low outgassing rates, and continuously circulating clean gas through the detector either with an open-loop system, where the gas is continuously replaced, or a closed-loop system, where the gas is continuously purified by a filter.

### Charge Amplification

The primary charge produced by the ionization of the gas medium is usually not enough to produce a measurable signal. The charge is therefore amplified in a region of the detector where a higher electric field  $E_a$  is applied. The amplification is usually done by avalanche multiplication.

Avalanche multiplication is a process where the primary charge produced by the ionization of the gas medium is accelerated by the electric field and produces secondary charge by ionizing the gas molecules. The secondary charge is then accelerated by the

electric field and produces tertiary charge, and so on. The process is repeated until the charge reaches the anode and is collected.

The number of secondary electrons  $N_a$  produced by the amplification of  $N_e$  primary electrons in a region of length  $L$  is given by

$$N_a = N_e e^{\int_0^L \alpha(x) dx} \quad (3.10)$$

where  $\alpha$  is the first Townsend coefficient, which depends on the gas medium and the electric field. The number of electrons after amplification of the primary charge  $N_a$  can be expressed as

$$N_a = GN_e \quad (3.11)$$

Applying error propagation to equation 3.11 and using expression 3.7, the variance of the number of secondary electrons after amplification,  $\sigma_a^2$ , can be written as

$$\sigma_a^2 = G\sigma_e^2 + N_e\sigma_G^2 = N_a \left( \frac{\sigma_e^2}{N_e} + \frac{\sigma_G^2}{G} \right) = N_a (F + b) \quad (3.12)$$

where  $b$  is defined as

$$b = \frac{\sigma_G^2}{G} \quad (3.13)$$

The resolution  $R = 2\sqrt{2 \ln 2} \sigma_a / N_a \approx 2.35 \cdot \sigma_a / N_a$  is more commonly used instead of the relative variance and is defined as the full width at half maximum (FWHM) of the distribution of the number of secondary electrons after amplification  $N_a$  divided by the mean number of secondary electrons after amplification  $N_a$ .

$$R \approx 2.35 \sqrt{\frac{1}{N_a} (F + b) + f(N_a)} \quad (3.14)$$

where  $f(N_a)$  is an additional term that accounts for other sources of noise such as electronic noise. The resolution  $R$  is commonly expressed as a percentage.

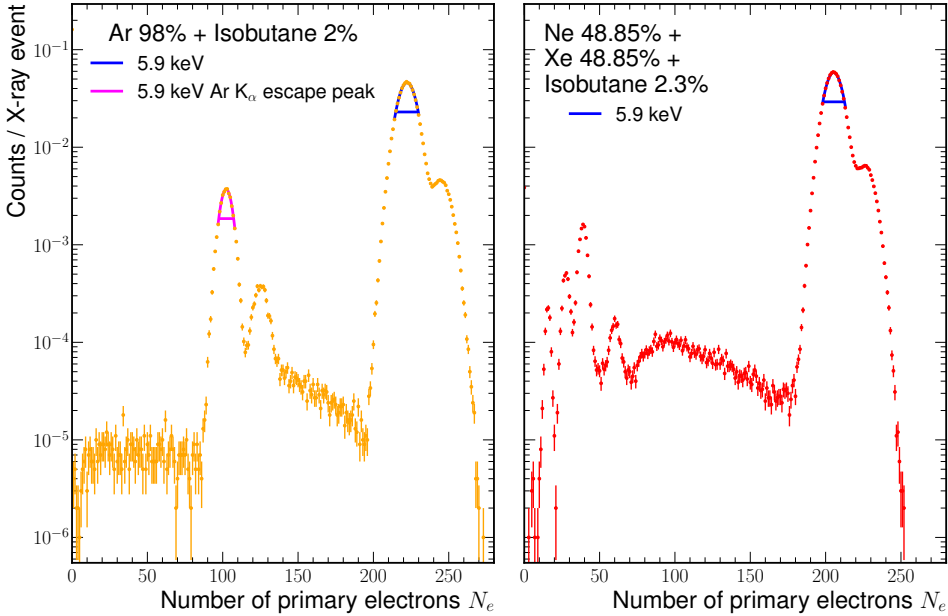


Figure 3.10: Simulation of the primary charge produced by the ionization of the gas medium by the X-rays from a  $^{55}\text{Fe}$  source. A volume of  $6 \times 6 \times 3 \text{ cm}^3$  is considered. The photons are generated at the center of the square face and directed toward the center. For the argon mixture the pressure is 1.4 bar and the drift field is 135 V/cm. For the xenon mixture the pressure is 1.05 bar and the drift field is 110 V/cm. The temperature is 20 °C for both mixtures. The 5.9 keV peak is fitted with a Gaussian function to obtain the FWHM resolution  $R$ , which is 15.0% for the argon mixture and 14.4% for the xenon mixture. The argon escape peak for 5.9 keV is clearly visible and is around 12.4 times less intense than the main peak.

Expression 3.14 shows that there is a lower limit to the detector resolution at relatively low energies, such as those of soft X-rays, and that this limit is dominated by the Fano factor. Simulations of the calibration of the BabyIAXO Micromegas detector with a  $^{55}\text{Fe}$  source 3.10 indicate that a resolution of around 15% is achieved for the 5.9 keV peak. These simulations only take into account the primary-charge generation; diffusion, amplification, and readout effects are not included. In practice, the detector resolution is therefore worse than the resolution obtained from primary-charge generation alone.

Besides avalanche multiplication, there is another significant charge-amplification process: electroluminescence. Electrons produced by the ionization of the gas medium can excite the gas molecules to higher energy levels without directly ionizing them. The excited molecules will eventually decay to their ground state by emitting a photon. The photon can ionize the gas medium and produce secondary charge. In the present

detector context, this process can have a negative impact on the resolution of the detector, both in the energy and spatial domains, since the photons produced can travel long distances and ionize the gas medium far away from the primary charge production site.

This unwanted ionization is mitigated by adding a small fraction of quencher gas to the noble gas. The quencher gas absorbs the photons produced by the electroluminescence process without ionization and is usually a saturated hydrocarbon, such as isobutane, which has a high probability of absorbing the photons produced by the electroluminescence process. The addition of the quencher gas also has the effect of increasing the drift velocity and reducing the diffusion of the electrons, which results in a better spatial and temporal resolution of the detector.

TODO: graph showing how diffusion is reduced by the addition of the quencher gas and drift velocity increases???

### 3.1.3 Micropattern Gaseous Detectors: Micromegas

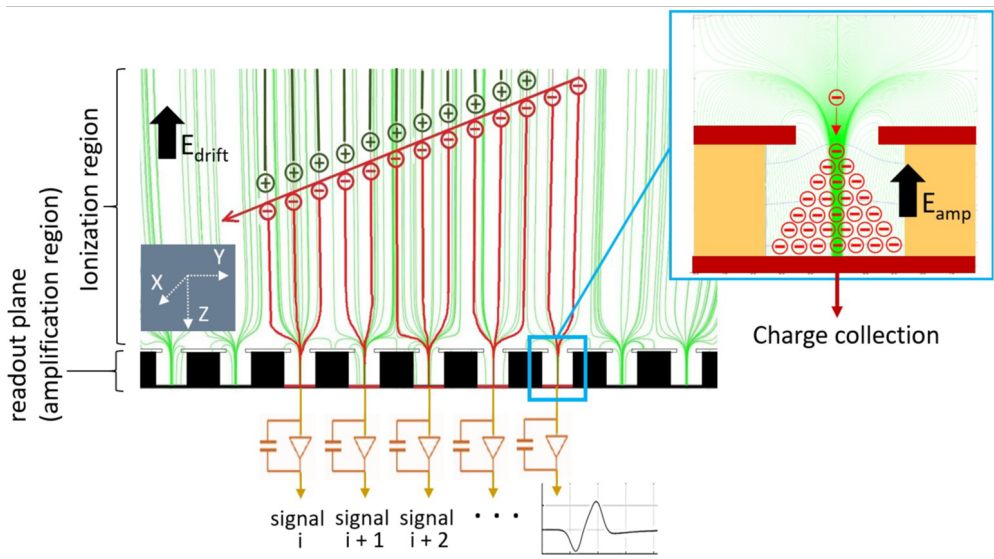


Figure 3.11: Diagram depicting the operation of a microbulk micromegas Micromegas detector. A charged particle such as a muon produces continuous ionization that is drifted into the readout which is then amplified in the region between the mesh and readout strips. Image from [64].

Micromegas (MICRO-MESH Gaseous Structure) detectors are a type of micropattern gaseous detector (MPGD) designed to achieve high spatial and temporal resolution, along with excellent energy resolution. Initially proposed in 1996 by Ioannis Giomataris

and colleagues at CEA Saclay[65], micromegas detectors have proven to be a robust and versatile technology, particularly well-suited for particle physics experiments, astrophysics, and medical imaging.

These detectors combine the advantages of gaseous detectors with microfabrication techniques, enabling them to achieve fine-grained spatial resolution while maintaining fast signal response times. Micromegas have some distinct advantages over other detector technologies, such as the ability to handle high radiation environments with minimal aging effects, making them a good choice for experiments with high particle fluxes, such as those at the Large Hadron Collider (LHC) at CERN. ATLAS, one of the LHC's primary detectors, uses micromegas technology in its muon spectrometer [66].

Another distinct advantage of the Micromegas detectors is its radiopurity. The detectors are made of materials with low intrinsic radioactivity, mostly copper and kapton. The gas medium is usually a noble gas, which can be either radiopure or has a low intrinsic radioactivity. The quencher gas is usually a saturated hydrocarbon, which is also radiopure. The radiopurity of this technology makes it an excellent choice for experiments searching for rare events, such as dark matter detection experiments. Some notable examples include solar axion searches in the CAST experiment [42, 67], WIMP direct searches in TREX-DM [68] or neutrinoless double beta decay searches in PANDA-X [69].

All these advantages, together with the successful history of operation in CAST, make the Micromegas technology the baseline detector technology for the IAXO experiment.

TODO: Talk about technologies, microbulk, gain and transparency curves...




Figure 3.12: Simplified design of a Micromegas detector. Showing only the strips and a cut-out of the mesh. The design is based on the BabyIAXO Micromegas detector (figure 3.17) but with  $8 + 8$  strips instead of  $120 + 120$  strips. The strips have been colored orange (X) or yellow (Y) according to their direction. The X strips pixels are connected at the same layer of the PCB while the Y strips pixels are connected at the bottom layer of the PCB. Only the three conducting layers of the PCB (mesh, pixels, traces) are shown. In reality these are supported by a dielectric layer in between.

## 3.2 The BabyIAXO Micromegas Detector Prototypes

The baseline detector technology for IAXO and BabyIAXO is the microbulk Micromegas. Its successful track record in the CAST experiment, where it was used to detect solar axions, makes it the ideal choice for the IAXO experiment.

The TPC design of the BabyIAXO Micromegas detector remains similar to the CAST Micromegas detectors, since both experiments are designed with the same goal of detecting solar axions. The gas volume is a cylinder with a diameter of 10 cm and a height of 3 cm, with a total volume of  $235.62 \text{ cm}^3$ . The microbulk micromegas readout has a square shape with a side of 6 cm. The readout has 120 strips for each direction, for a total of 240 channels.

There are two distinct generations of detector prototypes: IAXO-D0 and IAXO-D1. Both prototypes have the same core design specifications: TPC dimensions, readout size and number of strips, lead shielding thickness, etc. The main differences are in the

design of the chamber and pipe, lead shielding, the micromegas detector PCB and the electronics.

### 3.2.1 IAXO-D0




Figure 3.13: Simplified geometry of the IAXO-D0 detector and shielding as implemented in the simulations. Showing the TPC (a), the copper pipe (b), the 20 cm thick lead shielding (c) and the copper backplate protruding out of the shielding through the lateral opening (d).

The IAXO-D0 prototype is in essence the latest CAST design. The chamber is cylindrical with a large copper backplate that protrudes laterally from the shielding. The copper pipe is thinner in the middle and the shielding is made of lead bricks. In order to access or move the detector, the lead shielding must be partially removed brick by brick. The front-end electronic chips used are the AGET chips (4 chips in total) and the back-end electronics is a Feminos card [70]. The front-end and back-end electronics are physically connected into a single unit (FEC-Feminos) which is connected to the readout via flat cables.

A significant data taking campaign was performed with the IAXO-D0 prototype in Zaragoza using the exact same Micromegas detector as the CAST experiment, right after it was decommissioned. This campaign was used to validate the detector and to test the prototype veto system focusing on the reduction of the background caused by cosmic

rays. Other measurements campaigns pre-dating this work were also performed with the IAXO-D0 prototype using other Micromegas detectors from the same generation [48]. The results from these campaigns were used to validate the simulations and to improve the design of the IAXO-D1 prototype.

### 3.2.2 IAXO-D1




Figure 3.14: Simplified geometry of the IAXO-D1 detector and shielding as implemented in the simulations. Showing the TPC (a), the copper pipe (b), the 20 cm thick lead shielding (c) and the inner copper liner for the shielding (d). This design allows the shielding to be built as a single piece and the detector to be serviced by removing the shielding.

The IAXO-D1 prototype has some significant improvements with respect to IAXO-D0. The shielding is now a lead box with a square shaft in the middle. The shielding will sit on linear rails that will allow it to be moved laterally in order to access the detector. This will significantly reduce the time required to service the detector. The chamber design has a square footprint and the copper pipe is now significantly thicker. The Micromegas PCB can be bent in order to fit the hole and the readout strips exit the shielding hole in parallel to the pipe.

The front-end electronics are now the STAGE chips (4 chips, same as with IAXO-D0) and the back-end electronics is based in the ARC design, which is compatible

with the Feminos card. (TODO: CITE ARC, STAGE). The front-end electronics (figure 3.20a) are now placed inside the shielding next to the pipe in order to minimize the length of the flat cables. The back-end electronics are placed outside of the shielding (figure 3.20b) and are connected to the front-end electronics via another set of flat cables.




Figure 3.15: IAXO-D1 detector chamber. The left image shows the simplified geometry as used in the simulations and the right image shows a picture of the actual chamber. The electric field shaper ring can be seen in the outside of the PTFE liner. The gas lines (two, one not shown) can be seen running in parallel to the pipe. The chamber backplate and body are fastened to the pipe with 12 copper bolts.




Figure 3.16: Exploded view of the IAXO-D1 detector chamber created using the simplified simulations geometry from figure 3.15. Showing the copper backplate (a), readout PCB (b), readout / mesh (c), chamber copper walls (d), chamber inner PTFE lining (e), ultra-thin aluminized mylar window (f), conductive copper cathode support (g) and PTFE cathode support (h). In practise the cathode parts (f, g, h) are glued together into a single piece. The whole PCB (b) is coated in copper (see figure 3.15) in order to provide a uniform electric field in the TPC (only the area inside of the TPC is connected to high voltage).




Figure 3.17: IAXO-D1 Micromegas PCB. Image obtained by combining the Gerber files of the multiple layers of the PCB. The readout (yellow square) is shown on the center-right of the PCB. The connector pads for the 240 strips are shown on the left side of the image grouped by strip direction. High voltage pads can be seen in the top-left and top-right of the image. The blue circles indicate holes used to mount the PCB. (WHAT ARE THESE CONNECTIONS EXACTLY? MESH AND WHAT ELSE?)




Figure 3.18: Detailed view of the IAXO-D1 Micromegas PCB design (rotated zoom on the bottom-left of the readout (yellow square) from figure 3.17 showing connection traces for strips in both directions). The individual pixels (yellow) can be seen arranged in strips. In order to connect the pixels in this pattern it is required to connect the pixels for one strip direction in the bottom layer of the PCB (orange). The model from figure 3.12 is based on this design.





Figure 3.19: Pictures of the IAXO-D1 Micromegas. The area outside the square chamber body (gold-colored in the left picture) is flexible. The Micromegas mesh has a 3x3 rectangular hole pattern for each pixel as shown in the bottom-right picture.



(a) Front-end card (FEC).

(b) Back-end electronics.

Figure 3.20: IAXO-D1 electronics. The front-end electronics are flexible PCB cards using the STAGE chips and are placed next to the readout in the copper pipe. The back-end electronics are based on the ARC design and are placed outside of the shielding.

### 3.3 Detector Ancillary Systems

#### 3.3.1 Gas System

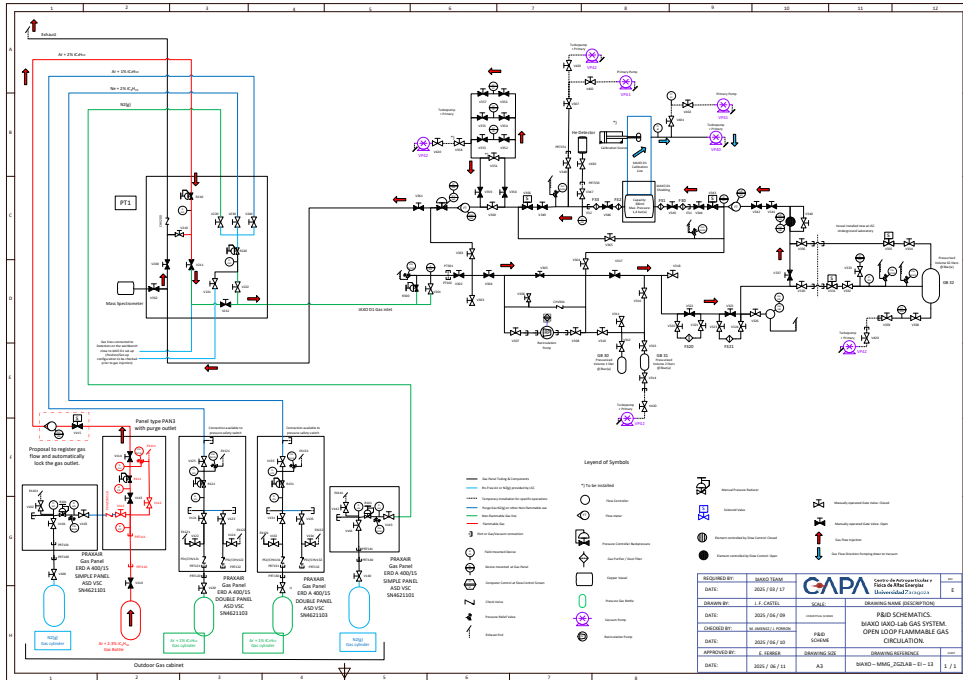


Figure 3.21: Diagram of the IAXO-D0/D1 gas system. A larger rotated version is provided as Appendix Figure 6.65, p. 282.

The gas system (see figure 3.21 for the detailed diagram) is responsible for the circulation of the gas medium through the detector. The gas flows from a high pressure gas bottle through a pressure reducer into the detector gas line. The gas flow into the detector is controlled by a set of computer-controlled pressure and flow controllers. The desired gas flow and pressure can be set to the desired values by the operator. A vacuum pump can be used to evacuate the gas from the detector and gas system in order to replace the gas medium.

A number of emergency relief valves are placed in the gas system to prevent overpressure in the detector. The window which separates the gas medium in the TPC from the outside (either atmospheric air or vacuum) is a critical detector component. The window is made of ultra-thin aluminized mylar and is placed in a very thin copper support frame. The window is designed to support up to 1.5 bar of pressure difference between the inside and outside of the detector.



Figure 3.22: Photo of the IAXO-D0 gas system. (TESIS ELISA). TODO: Actualizar con foto de IAXOLAB.

The IAXO Micromegas detectors are designed to operate with different gas mixtures. The gas mixtures considered at the time of writing are argon and a xenon-neon mixture, both using isobutane as the quencher gas. A mixture of argon with 2% isobutane at 1.4 bar was successfully used in CAST during most of its operation [42]. A mixture of xenon and neon (equal parts in partial pressure) with 2.3% isobutane was also used in CAST during the last years of operation [67].

The xenon-neon mixture presents some theoretical advantages with respect to the argon mixture. The argon mixture results in a significant escape peak in the energy region of interest of the axion due to the argon  $K_{\alpha}$  line as can be seen in figure 3.10. Natural argon also has a long-lived  $^{39}\text{Ar}$  isotope present in trace amounts which could potentially affect the background level of the detector (TODO: link to background model argon section).

The gas system is designed to operate in one of two modes: open loop or closed loop. In the open-loop mode, the gas is continuously replaced by a fresh gas flow from the bottle. The old gas is evacuated from the detector into an exhaust line. Operating under open loop is simple and should provide the best gas quality, assuming

the gas bottle is of high purity. It has the disadvantage of consuming a large amount of gas, which is not an issue if the gas mixture is inexpensive, as is the case of the argon mixture, but not for the xenon-neon mixture. In the closed-loop mode, the gas is continuously purified by a filter and recirculated back into the detector. The closed-loop mode is more complex and requires a gas filter and a gas recirculation system. Closed-loop operation is more suitable for expensive gas mixtures, such as the xenon-neon mixture.

This design direction is consistent with the IAXO pathfinder detector operated at CAST, where xenon-compatible running motivated the installation of a recirculation and purification line including isolation electrovalves, flow controllers, pressure sensors, and dedicated moisture and oxygen filters [45]. The CAST experience is especially relevant here because it showed that closed-loop operation is not merely a gas-handling convenience, but a practical requirement when moving toward more expensive mixtures and longer stable surface campaigns.

TODO: photo of the filters?

TODO: Binary gas analyzer?

TODO: Spectrometer?

### 3.3.2 High-Voltage System

The Micromegas detectors require a high-voltage system to provide the necessary electric field for the drift of the electrons and the amplification of the signal.

Two high-voltage channels are used for each Micromegas detector, one for the drift region (cathode) and one for the amplification (mesh) region. A CAEN high-voltage power supply is used to provide the detector bias.

HABLAR SOBRE LOS FIELD SHAPERS / FIELD CAGES.

HABLAR DE LOS FILTROS DE ALTA TENSION.

HABLAR DEL ALTO VOLTAJE DEL SISTEMA DE VETOS.

The CAEN high-voltage power supplies used can be controlled via a serial connection. This allows the high-voltage system to be managed by the Slow-Control, enabling the monitoring and control of the high-voltage levels remotely. In particular, the automatic control of the high-voltage of the detector is useful due to the possibility trips during long data taking periods. The Slow-Control can be programmed to automatically turn back on the high-voltage in case of a trip, to send an alert to the operators, or to turn off the high-voltage in case of repeated trips.

HABLAR DE LA LIBRERIA PYTHON HVPS?

### 3.3.3 Slow-Control System

The Slow-Control System is responsible for the monitoring and control of the ancillary systems of the Micromegas detectors. The Slow-Control continuously reads the status of the systems and can be programmed to take actions based on the readings. It can also be used to send alerts to the operators in case of a problem with the systems.

A web-based interface is used to interact with the Slow-Control, allowing the operators to monitor the status of the systems and to take actions such as turning on or off the high-voltage of the detectors. The data in this dashboard are updated in near real time, allowing the operators to quickly react to any issues that may arise. This data is also periodically stored in a database for later analysis.

The Slow-Control is built using node-red [71], a flow-based development tool for visual programming.

An instructive precedent was provided by the CAST pathfinder campaign, where a compact remote-control system was used to supervise detector high voltages, calibration-source position, readout power supply, muon-veto status, chamber pressure, and gas flow from a single interface [45]. That system also incorporated an automatic trip manager to recover the detector voltage after short discharges, a feature that became particularly important during xenon recirculation studies. This experience reinforces the view that slow control is not a secondary service, but an integral part of stable low-background operation.

## 3.4 Shielding and Veto Systems

The Micromegas detector is only one part of the complete low-background detection line. For BabyIAXO, the detector must operate inside a passive shield and in coincidence with an active veto system, so the mechanical envelope, signal feedthroughs, calibration access, gas services, high-voltage routing, and data-acquisition interfaces all have to remain compatible with the surrounding shielding. The detailed passive-shielding studies, cosmic-ray-induced background simulations, and scintillator-cadmium veto design are therefore treated in the dedicated shielding and veto chapter, Chapter 5. In the present chapter, the relevant point is the interface: Micromegas operation defines the x-ray-like event selection, timing reference, calibration strategy, and veto-coincidence information that the background model uses later.

## 3.5 Data Acquisition System

The operation of the Micromegas detector requires the coordination of several subsystems to ensure proper gas flow, high voltage or data acquisition. These systems are

monitored and controlled by a piece of software called the Slow-Control System.

The Data Acquisition (DAQ) System is responsible for the readout and storage of the detector signals. It is composed of three main components: the Front-End Card (FEC), the Front-End Module (FEM), and the Data Acquisition Computer (DAQ PC).

The overall architecture follows the same detector-readout philosophy adopted for the IAXO pathfinder line at CAST, where the 240 Micromegas strips were read through four AGET chips connected to a Feminos back-end board, with the external muon-veto signal routed to an otherwise unused AGET channel so that veto and TPC information could be recorded together on an event-by-event basis [45]. In this sense, the pathfinder operation serves as a practical bridge between the late CAST Micromegas generation and the more scalable BabyIAXO-oriented readout concepts.

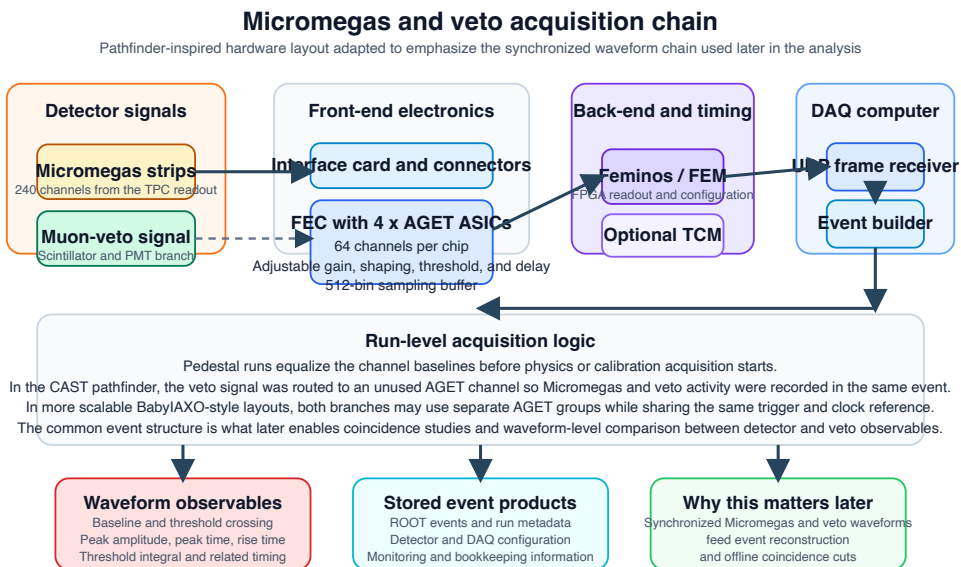


Figure 3.23: Acquisition-chain schematic for the Micromegas detector and the external veto signal, inspired by the CAST pathfinder layout described in Cristina Margalejo Blasco’s thesis [45]. The diagram emphasizes the pieces that are most relevant for this thesis: the strip and veto signals enter a common AGET/Feminos-based electronics chain, are assembled into synchronized event records by the DAQ computer, and are stored together with the waveform observables that later feed reconstruction and coincidence-based analysis.

Figure 3.23 is intentionally placed before the detailed hardware subsections because it summarizes the logic that connects the detector electronics to the later data analysis. At acquisition level, the important point is not only that the Micromegas strips are digitized, but that their waveforms are recorded in a timing-aware event structure



## Front-End Module




Figure 3.25: FEC FEM TODO: CHANGE and add labels for parts

The Front-End Card used with the IAXO Micromegas detectors consists of 4 AGET chips, each responsible of 60 channels for a total of 240 channels. The FEC is connected to the Front-End Module (FEM) as shown in Figure 3.25. The board used is called Feminos [70] and it is responsible for the communication with the FEC and the Data Acquisition Computer (DAQ PC).

The Feminos board consists mainly of an FPGA module, different connection interfaces such as RS232 and Gigabit Ethernet and a few additional components.

The Feminos is connected to the DAQ PC through an Ethernet connection. Multiple Feminos boards can be connected to the same DAQ PC, allowing for the readout of multiple detectors. In this case a Trigger Clock Module (TCM) is used to synchronize multiple Feminos boards.

The configuration of the FEC's AGET chips is done through the Feminos board via commands sent through the Ethernet connection from the DAQ PC. For instance, a command such as `aget * time 0x1` sent to the Feminos would set the shaping time of all the AGET chips in the FEC to 50 ns. (TODO: is this 50 ns or 70 ns?).

## Data Acquisition Computer

The Data Acquisition Computer (DAQ PC) is responsible for the communication with the Feminos boards and the storage of the data it receives. The DAQ PC runs a

custom software (DAQ software) that communicates over the Ethernet link with all the connected Feminos boards.

During the initial setup, the DAQ software sends the configuration commands to the Feminos boards to set the AGET chips to the desired settings. After the configuration is done, the DAQ software starts the data acquisition by sending a start command to the Feminos boards. The Feminos boards then start the readout of the AGET chips and continuously send the data to the DAQ PC in the form of UDP packets. The data is encoded in the form of data frames. The DAQ software is responsible of decoding the data frames which may come from multiple Feminos boards and constructing the final data file.

The result of an acquisition run is a file or a set of files containing some metadata related to the run and the run data which consists of a collection of events where each event is a collection of samples from the detector channels.

The DAQ software is described in more detail in Section 3.5.1.

### 3.5.1 Data Acquisition Software

The authors of the Feminos board provide a basic DAQ software written in the C programming language that can be used to communicate with the Feminos boards and start the data acquisition. The data frames are stored into a binary file that can be later decoded and analyzed. This software, with some modifications, was used during the CAST experiment data taking and during the initial tests of the IAXO Micromegas detectors.

During this work, a new DAQ software based on the original was developed in C++ to add new features and improve the data acquisition process. The sections related to the communication with the Feminos boards remain mostly the same, however in this new software the data frames are decoded and stored in a more structured way in a ROOT file.

Some of the new features added to the software are:

- Simplify the installation process by introducing a CMake build system for the software and its dependencies. Target C++17 to enhance code readability and maintainability.
- Integrate a Prometheus [74] exporter to allow for the monitoring of the acquisition process externally, for example, via the Slow-Control system.
- Write the data to a ROOT file instead of a binary file to allow for easier analysis and visualization of the data. The writing to the ROOT file is done without causing appreciable overhead to the data acquisition process.

- Develop a viewer software using `Python` and `Uproot` to visualize the data in the `ROOT` files, which can also be used to monitor the data acquisition process in real time.

## Data Storage Process

The data acquisition process is naturally I/O bound, both in the reception of the data from the Feminos boards and in the writing of the data to disk. The DAQ software needs process incoming frames faster than they are received, otherwise there is a risk of data loss.

The original DAQ software stored the data in a binary file as it was received while the new software stores the data in a `ROOT` file and performs some basic analysis of the data before writing it to the file. This means that there is a significant overhead in the data processing step compared to the original software. In order to minimize the impact of this overhead, the data processing is done in a separate thread from the data reception. Data is transferred via a shared queue between the two threads, allowing the data reception to continue while the data processing is done in parallel. This queue has a relatively big size to allow for some data processing delay without causing data loss. As long as the queue is not full, the data reception can continue. Still, the data processing thread must be able to process the data faster than it is received, otherwise the queue will eventually fill up and data will be lost. This design allows for a more costly data processing step without affecting the data reception, as long as the processing is done in a timely manner.

The output `ROOT` file is periodically flushed to disk to minimize the risk of data loss in case of a crash. In case of an unexpected stop of the DAQ software, the data stored in the `ROOT` file up to that point is still available for analysis.

## Structure of the DAQ `ROOT` File

The `ROOT` file is structured in a way that allows for easy analysis of the data. It can be opened by plain `ROOT` (without the need of dictionaries) or in `Python` via `uproot` [75].

Besides ease of use, another important aspect of the `ROOT` file is the compression of the data. Storage of the DAQ data represents a significant cost in terms of disk space, so it is important to minimize the size of the data stored.

`ROOT` automatically compresses the data stored in the file, but the compression settings can be adjusted to achieve a better compression ratio, usually at the cost of slower read/write times. A study of different compression settings was carried out in order to evaluate the `ROOT` file size and the read/write times.

There are two distinct regimes of operation: calibration runs where the incoming data rate is relatively high (up to 10 MB/s, but typically lower) and the data is mostly

noise, and background runs where the incoming data rate is much lower (around 100 kB/s).

The two main ROOT compression settings are:

- Compression algorithm: ROOT provides several compression algorithms, such as ZLIB, LZ4, and LZMA. The LZMA algorithm provides the best compression ratio, but it is also the slowest. The default algorithm is ZLIB, which provides a good balance between compression ratio and read/write times.
- Compression level: The compression level can be adjusted to achieve a better compression ratio at the cost of slower read/write times.

Type	Algorithm / Level	Write Time (s)	Read Time (s)	File Size (GB)	Branch Compression
<code>unsigned short</code>	ZLIB / Default	142.86	51.88	5.61	2.41
<code>unsigned short</code>	LZMA / Default	620.96	182.72	4.53	2.99
<code>unsigned short</code>	LZMA / 9 (Max.)	4621.88	144.56	3.76	3.61
<code>unsigned char</code>	ZLIB / Default	196.59	75.03	5.74	1.77
<code>unsigned char</code>	LZMA / Default	721.47	187.08	4.79	2.12
<code>float</code>	ZLIB / Default	302.71	86.95	8.04	3.37
<code>double</code>	ZLIB / Default	399.62	134.91	10.95	4.95

Table 3.1: Study of the effect of different ROOT compression settings on the size of the ROOT DAQ file and the read/write times. The data used is from a calibration run with 52194 events, each with 272 signals with 512 samples each. The signal values for the event are stored in a single `std::vector<Type>` per event. The signal values are 12-bit so in the case of `unsigned char` (8-bit) the data is stored by tightly packing two signal values into three bytes. ROOT v6.32.02 was used for the study. TODO: CPU / DISK INFO? (DONE IN WSL HOME DESKTOP)

The initial hypothesis was that storing the 12-bit signal tightly packed in an `std::vector<unsigned char>` would result in a smaller file size than storing the signal values as `std::vector<unsigned short>`. The study showed that this is not the case, as the file size is similar for both representations. If the data were stored without compression, the signal payload would be 25% smaller when stored as `std::vector<unsigned char>` than as `std::vector<unsigned short>`, but the compression algorithm achieves a similar compression ratio in both cases. Storing the data as `std::vector<unsigned char>` would also require additional processing to pack and unpack the data, so the approach does not offer a practical advantage.

As a result of the study, the LZMA algorithm with the default compression level was selected as the default compression setting for the DAQ ROOT files. This achieves a significantly smaller file size compared to the default algorithm with a small increase in read/write times. The DAQ software can also be configured to use a faster algorithm

in the case of calibration runs where the data rate is higher or a higher compression level in the case of background runs where the data rate is lower.

The processing speed of the DAQ ROOT files into a **REST-for-Physics** format was measured with different compression settings in the input file and no measurable difference was found between the different settings, so the choice of compression settings appears to only be relevant for the acquisition process.

### Data Acquisition Process

These are the main steps in the data Acquisition process with a FEC-Feminos setup:

An instance of the acquisition software is started on the DAQ PC. The software is configured with the IP addresses of the Feminos boards and the desired settings for the acquisition.

The start sequence is initiated by the operator. The software sends the configuration commands to the Feminos boards to set the AGET chips to the desired settings which includes starting the power-up sequence of the AGET chips.




Figure 3.26: IAXO-D0  $^{55}\text{Fe}$  calibration event sample. Showing all 240 channels with the same base level due to the pedestal subtraction.

A pedestal run is performed to measure the mean  $m_i$  and standard deviation  $\sigma_i$  of the channels when not triggered. The pedestal values are stored internally in the Feminos boards and are subtracted from the signal values during the data acquisition, this way all signals have the same base level. This base level (corresponding to zero signal) can be set in the configuration and is usually set to 250 ADC counts, as

recommended by the Feminos authors. Figure 3.26 shows a sample event with all channels at the same base level due to the pedestal subtraction.

Each channel can serve as trigger. The trigger level (the value above which a signal is considered a trigger) can be configured and it's set to three times the standard deviation of the pedestal values. The acquisition can be configured to return all signals or only the triggered signals.

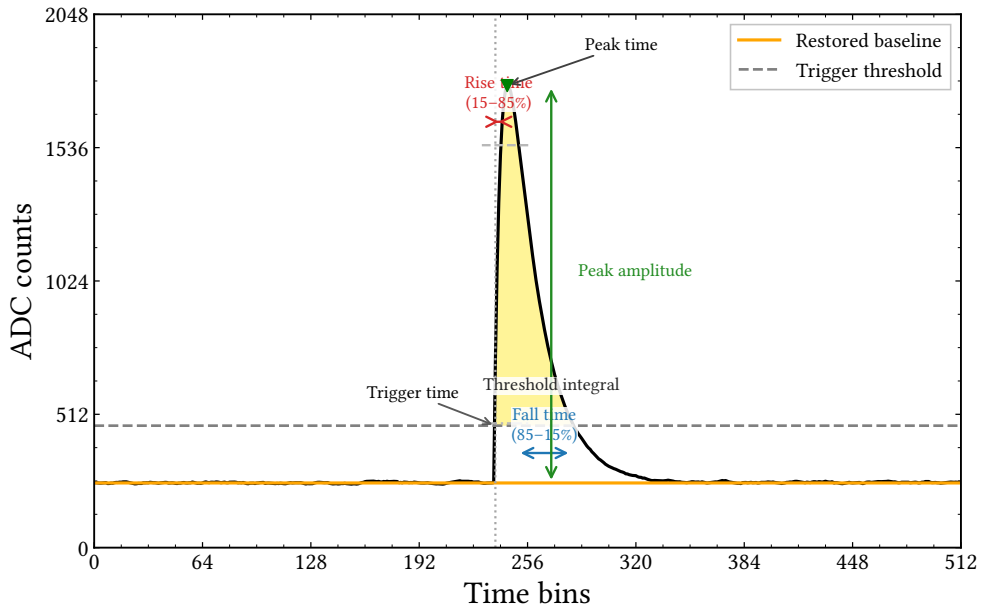


Figure 3.27: Representative single-channel Micromegas waveform after pedestal subtraction. The channel is characterized by the restored baseline, trigger threshold, peak amplitude and peak time, rise and fall times, and the threshold integral. These waveform-level quantities are the building blocks from which the later event-level observables used in the signal analysis are derived.

Figure 3.27 summarizes the morphology of a typical single-channel waveform after pedestal subtraction. At the lowest level, the relevant quantities are the baseline, the trigger-threshold crossing, the pulse maximum, and simple timing and integral descriptors such as the rise time, fall time, and threshold integral. These quantities are not merely electronics diagnostics. They define the first layer of observables from which the later x-ray selection is constructed, because they encode both the amount of collected charge and the temporal development of the signal. The DAQ settings that most directly control this waveform morphology are:

- **Gain:** Multiplicative factor applied to the pulse, directly affecting the measured peak amplitude and all integral-based observables.

- **Shaping time:** Time constant of the front-end shaper, which controls the pulse width and asymmetry and therefore modifies rise-time, fall-time, and time-over-threshold observables.
- **Sampling time:** Time between samples. It can be set up to 10 ns per bin (100 MHz). Given that the AGET chip has a 512 sample buffer, the minimum acquisition window is 5.12  $\mu$ s. This parameter fixes the time granularity with which the pulse morphology can be reconstructed.
- **Trigger delay:** Time between the trigger and the start of the acquisition window. It is typically chosen so that the trigger lies near the center of the buffer, preserving both pre-trigger baseline samples and the post-trigger tail of the pulse.
- **Trigger threshold:** Minimum excursion above the baseline required for a channel to trigger. This setting determines which signals are recorded and also influences threshold-based timing and integral observables.

These parameters can in principle be set channel by channel, but in practice they are chosen coherently for the full readout. Once the waveform has been digitized, the analysis software reduces it to per-channel observables such as peak amplitude, peak time, threshold integral, and rise time. These are then aggregated at event level into quantities stored in the `TRestAnalysisTree`, such as `rawAnaTPC_MaxPeakAmplitude`, `rawAnaTPC_ThresholdIntegral`, `rawAnaTPC_RiseTimeAvg`, `rawAnaTPC_MaxPeakTimeDelay`, and `rawAnaTPC_TripleMaxIntegral`. For x-ray-like events, one expects relatively compact and mutually synchronous strip pulses, while extended tracks, pile-up, or pathological events tend to produce broader, more asymmetric, or less synchronous responses. For that reason, the pulse description introduced here is already part of the logic of the later signal analysis rather than a purely instrumental detail.

## 3.6 Detector Calibration

The operation of a Micromegas detector involves a set of coupled parameters that must remain under control in order to guarantee stable gain, reproducible energy calibration, and a well-defined trigger threshold. In the gas system, the most relevant settings are the gas flow, chamber pressure, and gas mixture. In the detector itself, the drift and amplification voltages determine the transparency of the mesh and the gas gain. At the readout level, shaping time, trigger delay, sampling configuration, and threshold settings define how the pulse is digitized and which events are retained. In addition to these explicitly configured quantities, environmental variables such as temperature, as well as slow drifts in gas quality, can also modify the detector response. For that reason,

the relevant detector and DAQ settings are stored together with the run metadata and are monitored through regular calibration runs.

From the point of view of data taking, the acquisition is organized in runs, each one corresponding to a period with fixed detector and DAQ settings. The run duration may range from a few minutes to several hours depending on the purpose of the measurement. In practice, the most important distinction is between calibration runs and background or tracking runs. The former provide a controlled x-ray-like reference with which the detector gain, energy scale, resolution, and threshold can be monitored. The latter provide the physics data used for background characterization and, when relevant, for axion-sensitive exposure. In the CAST pathfinder campaign, the daily operating sequence was explicitly structured around this logic, with regular  $^{55}\text{Fe}$  calibrations used to track the detector response and to associate each background or tracking period with the most representative nearby calibration [45].

### 3.6.1 Standard x-ray calibration

The standard calibration procedure uses a source of soft x rays placed in front of the detector window, typically through a dedicated calibration port. The most common choice is  $^{55}\text{Fe}$ , whose dominant manganese K-shell line at 5.9 keV lies inside the energy region most relevant for the axion search and generates compact, x-ray-like events in the gas. Additional sources such as  $^{109}\text{Cd}$  may be used for complementary checks at higher energy, but  $^{55}\text{Fe}$  remains the reference source for routine operation. Since photons in this energy range interact predominantly through the photoelectric effect, the resulting signal is a localized ionization cluster that closely resembles the topology of a low-energy x-ray conversion in the Micromegas gas.

In routine operation, the calibration data are first used for a fast data-quality check. A reconstructed hit map verifies that the source illuminates the expected region of the detector and that no large-scale asymmetry or dead area has appeared. The corresponding energy spectrum is then inspected to verify the position and width of the main photopeak, which provides an immediate monitor of gain stability, energy resolution, and effective threshold. This daily quick-look procedure was an important part of the CAST pathfinder operation and is especially relevant for surface-running detectors, where small changes in gas quality, voltage settings, or noise conditions can translate into visible shifts of the calibration peak from run to run [45]. Figure 3.28 shows an example of such a quick-look spectrum using an IAXO-D0  $^{55}\text{Fe}$  calibration run. The dominant peak is associated with the manganese  $K_\alpha$  line at 5.9 keV, while the smaller contribution at higher integral is compatible with the  $K_\beta$  line at 6.5 keV. The fit to the main peak provides the run-wise calibration anchor that is later propagated to the associated physics data.




Figure 3.28: Reconstructed  $^{55}\text{Fe}$  calibration spectrum for the IAXO-D0 run R02756 using the `rawAnaTPC.TripleMaxIntegral` observable. The dominant structure corresponds to the manganese  $K_\alpha$  line, while the weaker component at higher integral is consistent with the  $K_\beta$  contribution. The fit around the main peak defines the run-wise conversion between reconstructed pulse integral and deposited energy.

The energy calibration itself is obtained from fits to the reconstructed calibration spectrum. For argon-based mixtures, the main 5.9 keV photopeak and the argon escape structure around 2.9 keV provide a convenient linear calibration of ADC units to deposited energy [45]. For xenon-based operation, where the argon escape feature is absent, the calibration can be anchored to the 5.9 keV line alone with a linear relation through the origin. Repeating this procedure run by run makes it possible to quantify the long-term evolution of the gain and to propagate the appropriate calibration constants to the associated physics runs. In a software-oriented workflow, this is important not only for detector operation but also for the reproducibility of the later reconstruction and cut-definition stages.

For the present thesis, the role of calibration extends beyond the treatment of measured source data. Dedicated calibration-like Monte Carlo samples were also produced in the collaboration repository `iaxo-simulations` [76], as discussed in Sections ?? and 6.3. These samples included  $^{55}\text{Fe}$ -like configurations and uniform low-energy gamma simulations in simplified Micromegas geometries. Their purpose was to study how x-ray-like observables evolve with energy, to validate the detector-response model, and to estimate the efficiency of the reconstruction and discrimination cuts

under controlled conditions. This point is particularly relevant in a software-focused thesis, because it shows that calibration is not only an experimental procedure used to assign an energy scale, but also a computational benchmark through which the common `REST-for-Physics` analysis chain can be tested, tuned, and compared consistently between data and simulation. Figure 3.29 illustrates this point with a direct comparison between a measured  $^{55}\text{Fe}$  calibration run and a dedicated simulation sample after both spectra are anchored to the 5.9 keV line. In order to make the simulated calibration event more representative of the measured detector response, the reconstructed Monte Carlo observable is additionally smeared to reproduce the experimental width of the 5.9 keV peak. The presence of the argon escape feature in both cases then supports the consistency of the common reconstruction chain while preserving a visually realistic comparison.




Figure 3.29: Comparison of normalized  $^{55}\text{Fe}$  calibration spectra in data and simulation after anchoring the dominant line to 5.9 keV. The same reconstructed observable, `rawAnaTPC.TripleMaxIntegral`, is used in both cases. The simulated spectrum includes an additional detector-resolution smearing chosen to reproduce the measured width of the 5.9 keV peak. With this correction, the simulation reproduces the main line structure and the argon escape feature while remaining directly comparable to the measured calibration data.

### 3.6.2 UV-light calibration R&D

A complementary calibration concept was explored during a three-month internship at CEA Saclay. This work is not part of the baseline BabyIAXO calibration strategy, which remains based on x-ray source runs, but it is worth retaining as detector R&D because it addresses a closely related problem: how to generate controlled, localized, and time-stamped primary electrons in a Micromegas gas volume. The idea is inspired by Micromegas-based photocathode detectors such as PICOSEC [77], where ultraviolet photons release photoelectrons that are subsequently drifted and amplified.

In the CEA setup, a pulsed ultraviolet source illuminated an aluminized cathode through a UV-transparent window. The lamp trigger provided a timing reference, while the anode pulse arrival time was measured after electron drift and amplification. By repeating the measurement with spacers of different thicknesses, the drift distance was changed in a controlled way and the drift velocity could be estimated from the variation of pulse arrival time with distance. The presentation study used argon–isobutane mixtures, several quencher fractions, and comparisons with `Garfield++/Magboltz` drift-velocity calculations. The measured velocities had the expected order of magnitude and qualitative field dependence, although offsets with respect to the simulation remained and were attributed to possible gas-settling, field, photocathode, or space-charge effects.

For the present thesis, the importance of this work is methodological rather than as a mature calibration proposal. A pulsed UV system could, in principle, provide single-electron or few-electron calibration, localized topological checks, timing studies, and gas-transport measurements without relying only on radioactive x-ray sources. However, the tested gas-discharge lamp had limited pulse stability, and a quantitative implementation for BabyIAXO would require a better-controlled UV source, calibrated photocathode response, stable gas conditions, and a dedicated comparison with the final detector geometry and readout. Supplementary figures from this R&D study are collected in Appendix 6.16.

---

## Chapter 4

# Computational Framework

4.1	Scope and software contributions of this thesis . . . . .	80
4.2	The ROOT Data Analysis Framework and Python interoperability . . .	81
4.2.1	Uproot and fsspec integration . . . . .	81
4.3	REST-for-Physics as the common event-processing framework . . . .	82
4.3.1	Event-based processing and the EventTree/AnalysisTree model	83
4.3.2	Reproducibility and maintainability . . . . .	85
4.4	Detector readout metadata and channel mapping . . . . .	85
4.5	Geant4 simulation infrastructure . . . . .	87
4.5.1	restG4 and Geant4Lib: the interface to Geant4 . . . . .	87
4.5.2	Geometry-generation infrastructure for IAXO simulations . . .	89
4.5.3	Python interface to REST-for-Physics . . . . .	90
4.5.4	Probability Distribution Projection cosmic-ray generator . . . .	90
4.5.5	Cosmic-ray source generators . . . . .	92
4.5.6	Radiation transport in gases: Garfield++ . . . . .	93
4.6	Visualization and online diagnostics . . . . .	94
4.6.1	Browser-based Geant4 event viewer . . . . .	94
4.6.2	DAQ event viewer: feminos-viewer . . . . .	96
4.7	Monte Carlo production at scale . . . . .	97
4.7.1	End-to-end simulation and analysis chain . . . . .	97
4.7.2	Production workflow with HTCondor . . . . .	98
4.8	geant4-python-application: Python prototyping with Geant4 . . .	102
4.9	Summary and role in the thesis . . . . .	104

## 4.1 Scope and software contributions of this thesis

This chapter describes the software tools and framework components that were developed, extended, or validated during the course of this thesis. It does not attempt to catalogue the full IAXO software stack; instead, it focuses on software that was either created as part of this work or modified in ways that directly enabled the background-model studies and veto analysis presented in later chapters.

The distinction between pre-existing infrastructure and thesis-specific contributions is important for understanding the role of each component. `REST-for-Physics` [78], `Geant4` [79–81], `ROOT` [82], `CRY` [83], and the Kotlin GDML domain-specific language were already available when this work began. The contributions of this thesis lie in integrating these tools into a reproducible end-to-end simulation and analysis workflow, adding new generators and analysis processes, developing auxiliary tools for validation and monitoring, and scaling the production to the throughput required for an IAXO background model. Tables 4.1 and 4.2 summarize the main software components and their role in the thesis.

Table 4.1: Core production and analysis infrastructure developed, integrated, or validated during this thesis.

Component	Problem	Thesis contribution	Validation or output	Role in thesis
REST-for-Physics model	Reproducible, version-tracked analysis chains for rare-event searches	Used framework as common analysis language for IAXO; configured RML chains, validated <code>EventTree/AnalysisTree</code> workflow	<code>analysis.rml</code> chain used for all productions; consistency checks vs. experimental data	Common analysis framework
<code>restG4</code> <code>Geant4Lib</code>	+ Performance, storage, multi-threading, and flexibility of <code>Geant4</code> simulations	Multi-threading, track pruning, sub-event splitting, interrupt handling, volume hash resolution	>600 + >370 commits; 1–10× storage reduction; validated with cosmic productions	Transport and background model
Micromegas/veto readout	From electronics channels to physically meaningful detector observables	Generated <code>TRestDetectorReadout</code> for IAXO-D0/D1 strips and 59-panel veto; validated against geometry	Channel-ID audit tools; cross-check of simulation readout vs. experimental readout	Detector response / veto analysis
Geometry integration	Reproducible IAXO geometries for shielding/veto scans	Integrated Kotlin GDML DSL into <code>restG4</code> , resolved hashed volume names, validated readout consistency	Production geometries linked to simulation files via Git commit	Geometry / veto scans
Track pruning	Large output files from full-track storage	Configurable volume-of-interest pruning preserving tracks leading to relevant detector deposits	1–10× output reduction; validation runs checked the veto/TPC observables used in the analysis	Production optimization
PDP cosmic generator	Inefficient cosmic secondary source generation wasting CPU on non-detector trajectories	Published Probability Distribution Projection method; samples directly from trajectories intersecting enclosing sphere	Peer-reviewed [84]; 3–37× yield improvement	Cosmic-background production
CRY-to-REST source histograms	Reusable atmospheric-secondary source terms for many geometries	Developed auxiliary generator that runs <code>CRY</code> , extracts particle-dependent $E$ - $\theta$ distributions, and stores them as <code>ROOT</code> histograms	Cross-checked against <code>EXPACS</code> , <code>HENSA</code> , and detector-level cosmic productions	Surface source term
HTCondor production	Manual job management for large campaigns	Developed <code>restG4ToCondor.py</code> with <code>DAGMan</code> , merging, output staging, and dry-run support	300-job campaigns; validation productions totaling about 40 000 final events across 600 Condor jobs	Production infrastructure

Table 4.2: Auxiliary, diagnostic, and prototyping tools developed or extended during this thesis.

Component	Problem	Thesis contribution	Validation or output	Role in thesis
Uproot + fsspec	Python ROOT-file I/O limited to local files and a few remote protocols	Delegated all <code>Uproot</code> I/O to <code>fsspec</code> , enabling SSH, cloud, and other compatible protocols without protocol-specific code	Upstreamed to <code>Uproot</code> v5.2.0; <code>feminos-viewer</code> remote support	DAQ / remote I/O enabled
Browser viewer	event No portable tool for event topology, veto patterns, geometry inspection	Developed web viewer using <code>three.js</code> ; <code>RestGeant4Event</code> to JSON; exports figures	Used to debug geometry, validate veto mapping, produce thesis figures	Visualization / validation
<code>feminos-daq-viewer</code>	/ Legacy <code>mclient</code> lacked live monitoring and ROOT output	Refactored DAQ: ROOT output, Prometheus monitoring, CMake build, live Python waveform viewer	Used during IAXO-D0/D1 operation for real-time diagnostics	DAQ / online diagnostics
<code>geant4-python-application</code>	Standard workflows are application based, limiting notebook prototyping	<code>Geant4</code> <code>Pybind11</code> wrapper with process isolation; pip-installable wheels based, on PyPI	Attenuation-length notebooks; not used for production simulations	Auxiliary prototyping

## 4.2 The ROOT Data Analysis Framework and Python interoperability

ROOT [82] is the data analysis framework developed by CERN and the foundation on which `REST-for-Physics` is built. For this thesis, the most relevant features of ROOT are its file format (`.root`), its `TTree` columnar storage, its Python bindings (`PyROOT`), and the serialization dictionary system that `REST-for-Physics` relies on for metadata persistence and versioning.

The simulations and analyses reported in this thesis used ROOT v6.34.04. The ROOT file format remains the primary data container for both Monte Carlo and experimental data within the IAXO collaboration. Every `restG4` simulation, every `restManager` processing step, and every DAQ acquisition file uses the same underlying storage format, which is essential for maintaining a single analysis language across the full chain.

While ROOT provides extensive built-in visualization and analysis tools, most of the figures in this thesis were produced with Python libraries such as `Matplotlib` and `Plotly`. This was possible thanks to two developments: ROOT’s own `PyROOT` interface, which exposes C++ objects to Python, and `Uproot` [75], a pure-Python library that reads and writes ROOT files without requiring a ROOT installation.

### 4.2.1 Uproot and fsspec integration

`Uproot` is widely used for reading ROOT files in Python and underpins many Python-based HEP analysis workflows. A three-month IRIS-HEP fellowship during this thesis provided the opportunity to contribute to `Uproot` directly. At the time, `Uproot`

supported local files and a few remote protocols (HTTP, XRootD), each requiring a hand-written implementation inside the library. The fellowship work delegated all file-I/O operations in `Uproot` to `fsspec`, a Python library that provides a uniform interface across local and remote file systems. This change, released in `Uproot` v5.2.0, simplified the codebase and made every `fsspec`-compatible protocol (including SSH, cloud object stores, and WebDAV) available to `Uproot` without additional maintenance burden.

This integration was directly useful for the present thesis in two ways. First, it enabled the `feminos-viewer` application to open remote DAQ files over SSH while they are being written, using the same code path as local files. Second, it allowed the Python-based background-analysis scripts to read simulation outputs directly from the NAF-IAXO dCache storage without copying them locally. The experience also provided a thorough understanding of the ROOT file format, which was important for designing the DAQ ROOT output and for debugging file-structure issues during analysis.

### 4.3 REST-for-Physics as the common event-processing framework

`REST-for-Physics` [78] (Rare Event Searches Toolkit for Physics) is an open-source, collaborative C++/ROOT-based framework that provides a unified environment for data acquisition, Monte Carlo simulation, detector-response emulation, and physics analysis. It was originally developed for experiments searching for rare phenomena such as neutrino interactions, dark matter, and axion signals, where precise detector modeling and reproducible analysis chains are essential.

The framework's relevance extends beyond the present thesis. Cristina Margalejo Blasco's recent thesis on the IAXO pathfinder detector at CAST independently demonstrates the same architecture being used for a Micromegas-based low-background line that is directly connected to the BabyIAXO and IAXO detector technology roadmap [45]. That work reinforces the claim that `REST-for-Physics` is most powerful when used as a common analysis language spanning transport, response, reconstruction, and event selection for Micromegas-based experiments.

`REST-for-Physics` is structured as a core Git repository with libraries and packages as Git submodules. It is strongly rooted in the ROOT ecosystem and follows ROOT conventions for class naming, I/O patterns, and dictionary-based serialization. The framework exposes three main executables: `restG4`, a user-configurable `Geant4` application initialized via RML configuration files; `restManager`, which applies a sequence of processing stages to data; and `restRoot`, a ROOT interpreter wrapper that loads the `REST-for-Physics` environment for interactive data inspection.

### 4.3.1 Event-based processing and the EventTree/AnalysisTree model

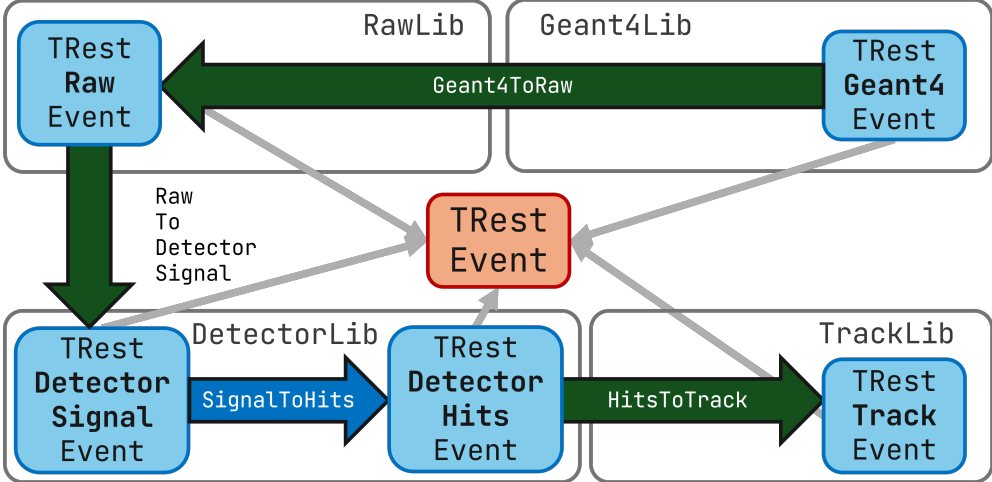


Figure 4.1: Schematic representation of the event-oriented design of REST-for-Physics. Different processing stages operate on `TRestEvent`-derived containers and may append observables to the `AnalysisTree` while preserving the corresponding event representation in the `EventTree`.

REST-for-Physics follows an event-driven architecture. The abstract `TRestEvent` class is the central data container, with libraries defining their own derived event types (e.g. `TRestGeant4Event`, `TRestRawSignalEvent`, `TRestTrackEvent`). Each library also provides processes that transform one event type into another or compute scalar observables.

Data persistence is organized around two complementary `TTree` structures. The `EventTree` stores full event representations, preserving the physics content at each stage of the chain. The `AnalysisTree`, implemented as `TRestAnalysisTree`, stores scalar observables computed by the different processes and is used for cuts, control plots, efficiency studies, and background estimates. This separation is especially valuable in rare-event searches, where thresholds and selections are often scanned many times without regenerating the complete event representation.

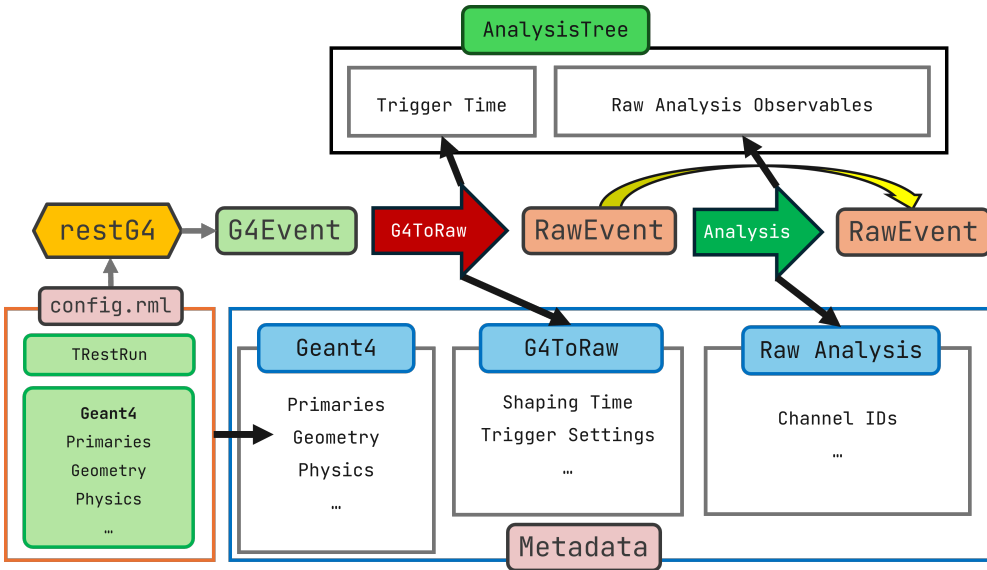


Figure 4.2: Typical REST-for-Physics Geant4 workflow. A simulation configured through RML is executed with `restG4`, after which `restManager` applies a sequence of processing stages that emulate detector response, electronics, and reconstruction while filling analysis observables.

For the IAXO-D0 and BabyIAXO studies, a common `analysis.rml` chain transformed `Geant4` truth information into Micromegas and veto observables that could be treated with the same logic used for measured data. In the TPC branch this meant deriving reconstructed hit and track observables for X-ray/background discrimination. In the veto branch it meant generating realistic waveforms and applying peak finding to obtain timing, amplitude, and multiplicity observables. The event-type evolution is one of the key strengths of the framework, because it ensures that the final analysis is performed on reconstructed observables rather than on idealized Monte Carlo truth quantities.

The `AnalysisTree` approach also provides a systematic way of turning reconstruction outputs into reusable analysis products. Observables filled by each process can be inspected interactively, processed through `ROOT` macros or `Python` scripts, or consumed by dedicated REST-for-Physics plotting utilities. In this thesis many final figures were produced with external tools, but the underlying quantities were still derived from the common REST-for-Physics reconstruction chain, ensuring that simulation and experimental analyses remained aligned at the observable level.

### 4.3.2 Reproducibility and maintainability

`REST-for-Physics` stores the full analysis configuration as metadata objects inside the `ROOT` file, including framework version, dependency versions, and RML parameters. This allows any `ROOT` file to be re-analyzed with the original configuration or to be inspected for provenance. The source code is version-controlled and openly available on GitHub; periodic releases specify recommended versions of `ROOT`, `Geant4`, and `Garfield++` to simplify reproducibility. Container images with pre-built dependencies are provided for users who want a consistent environment, and a comprehensive test suite with Google Test runs in GitHub Actions CI for every pull request. These practices were adopted throughout the production campaigns described in this thesis, and the CI infrastructure was migrated from GitLab to GitHub Actions during this work. Table 4.3 summarizes the practical provenance information that was preserved or recorded for the production campaigns used in this thesis.

Table 4.3: Reproducibility manifest for the simulation and analysis workflows used in this thesis.

Item	Stored or recorded in	Purpose
<code>REST-for-Physics</code> version / commit	<code>ROOT</code> metadata and production environment	Identifies the event classes, process implementations, and metadata schema.
<code>ROOT</code> , <code>Geant4</code> , and <code>Garfield++</code> versions	Environment metadata, release notes, and container configuration	Fixes the external software stack used for transport, I/O, and gas-parameter generation.
Geometry repository commit	Simulation metadata and <code>iaxo-geometry</code> repository	Links each production to a specific GDML geometry and veto/shielding configuration.
RML configuration files	Input repository and serialized <code>ROOT</code> metadata	Records the source, physics-list, processing, reconstruction, and observable definitions.
Random seeds and job identifiers	<code>HTCondor</code> logs and job outputs	Allows failed jobs to be diagnosed and statistically independent campaigns to be checked.
Source histograms	Input <code>ROOT</code> files and source-generation scripts	Records the atmospheric secondary distributions used for muons, neutrons, photons, electrons, and protons.
<code>HTCondor</code> logs and DAG identifiers	Batch-system output directories and merge logs	Documents job splitting, runtime, failures, restarts, output staging, and merge history.
Final output metadata and <code>AnalysisTree</code>	Merged <code>ROOT</code> files	Preserves the observables used for cuts, efficiency studies, and background-rate calculations.

## 4.4 Detector readout metadata and channel mapping

An essential intermediate layer between transport simulation and physics analysis is the detector readout description encoded in the `TRestDetectorReadout` metadata classes. In plain terms, this layer translates electronics channel numbers into physical detector positions and types, making it possible to know whether a signal came from a Micromegas strip, a veto scintillator panel, or a specific region of the detector. Without

it, a waveform would remain only a list of ADC samples attached to an integer channel number. With it, the same processing chain can move consistently from raw signals to detector signals, and from detector signals to reconstructed hits and tracks, while preserving the physical interpretation of each channel.

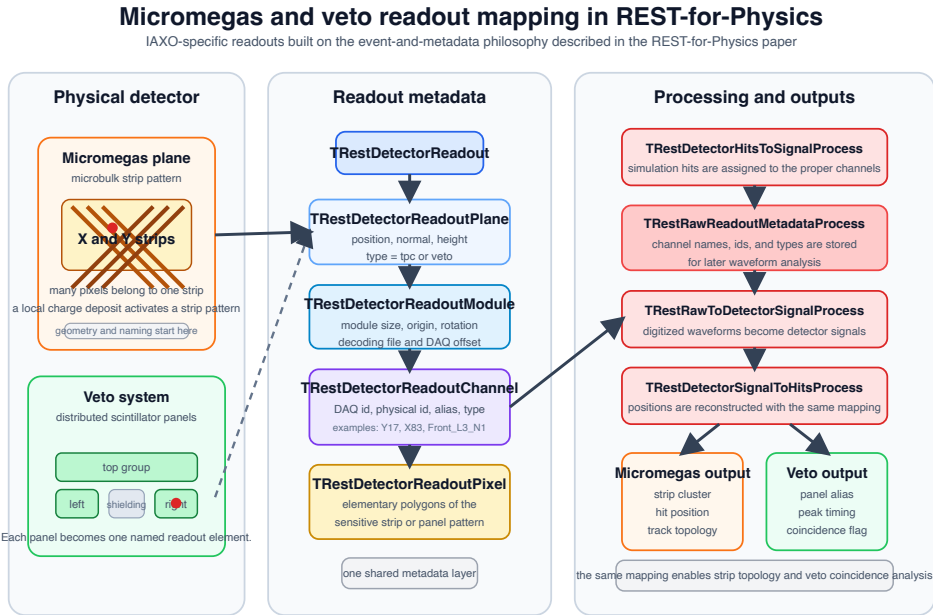


Figure 4.3: Schematic view of the readout-metadata layer used in this thesis, combining the event-and-metadata philosophy of REST-for-Physics [78] with the IAXO-specific Micromegas and veto readout definitions. The left side shows the physical detector elements, the center shows the hierarchical `TRestDetectorReadout` description that maps them to channels and aliases, and the right side shows the processes that use that mapping to move between detector hits, raw signals, detector signals, and reconstructed observables.

The readout description is organized hierarchically. `TRestDetectorReadout` stores the full detector readout as a collection of `TRestDetectorReadoutPlane` objects. Each plane defines a position and orientation in world coordinates, the normal vector that identifies the drift side, the effective height of the active volume, and a semantic type such as `tpc` or `veto`. Each plane contains `TRestDetectorReadoutModule` objects, which define the local geometry of a readout module. Inside each module, `TRestDetectorReadoutChannel` stores the correspondence between the DAQ identifier and the physical channel identifier, while `TRestDetectorReadoutPixel` provides the elementary polygons representing the actual sensitive pattern. A physical channel may

be composed of many pixels combined into a complex strip or pad geometry, which is important for the microbulk Micromegas pattern, where each strip is constructed from multiple pixels including the special edge pieces required by the real detector layout.

For the veto system, the effort was even more detector-specific. The scintillator panels do not form a single regular plane but a distributed system of individually oriented detector elements surrounding the shielding. Dedicated readout-generation code was developed to build one `TRestDetectorReadoutPlane` per veto panel, using the geometry information to determine the panel position, its outward normal, and the corresponding effective sensitive depth. Each panel was assigned a unique channel identifier and an alias matching the experimental naming convention. This mapping step was essential because the later analysis is formulated in terms of physical veto groups, layers, and aliases rather than arbitrary DAQ integers.

This metadata layer bridges the simulation-side and data-side representations of the detector response. On the simulation side, processes such as `TRestDetectorHitsToSignalProcess` use the readout geometry to decide which channels collect charge or light from a given interaction point. On the experimental side, `TRestRawToDetectorSignalProcess` and `TRestDetectorSignalToHitsProcess` use the same readout definition to transform digitized waveforms back into detector signals and then into reconstructed spatial hits. The auxiliary `TRestRawReadoutMetadataProcess` serializes channel-level information so that later waveform-analysis stages can distinguish TPC channels from veto channels without relying on external spreadsheets.

This was one of the most consequential software contributions of this thesis because it made the Micromegas and veto branches speak the same analysis language. That capability becomes especially important in the signal-analysis chapters, where waveform observables, strip topology, and veto coincidences must be combined in a single event-level selection strategy.

## 4.5 Geant4 simulation infrastructure

### 4.5.1 restG4 and Geant4Lib: the interface to Geant4

`Geant4` [79–81] is a Monte Carlo toolkit for simulating the passage of particles through matter, widely used in high-energy physics, nuclear physics, and medical applications. `Geant4` does not provide a command-line interface; users write a `C++` application that instantiates the toolkit with a specific geometry, physics list, and primary generator. These applications share a large amount of boilerplate code, and `Geant4` offers no built-in solution for serializing event data.

`REST-for-Physics` addresses this through `restG4` (the executable) and `Geant4Lib`

(the library): a modular, user-configurable **Geant4** application that is driven by RML configuration files rather than recompilation. During this work, **restG4** received over 600 commits and **Geant4Lib** over 370 commits. The most significant improvements are summarized below.

### **Geant4 event model and REST-for-Physics representation**

In the **Geant4** framework, a **run** is a collection of events sharing a fixed detector geometry and physics configuration. A run is initialized once (an expensive operation that loads geometry and cross sections) and then executes many events. A **primary generation** occurs at the beginning of each event, when one or more primary particles are created with specified energies, directions, and positions. Each primary particle is tracked through the geometry in a process called radiation transport, moving in discrete **steps** whose lengths are determined by the cross sections of active physics processes. Interactions can produce secondary particles, which are themselves tracked. The full chain of secondaries must be considered when calculating the energy deposited in a given detector region.

The principal **Geant4** data structures (run, event, track, and step) are represented in **Geant4Lib**. **Geant4** steps are stored as **REST-for-Physics** hits, which contain the position, momentum, time, energy, interaction process, and target information at a single point. **Geant4** also provides user-action hooks at various levels (event, tracking, stepping); a stacking-action hook was added during this work to tag new tracks before processing, enabling long-lived secondaries to be isolated into separate sub-events with distinct sub-event identifiers while remaining traceable to the same **Geant4** event.

### **Multi-threading and track pruning**

Radiation transport is naturally parallelizable at the event level, since events are independent. During this work, support for **Geant4** multi-threading was added to **restG4**, requiring a thread-safe event container and synchronization primitives with minimal performance impact. Support for interrupt-signal handling was also added, allowing a simulation to be stopped cleanly while preserving all output produced up to that point.

Track pruning was implemented to reduce output size and analysis time. The user specifies which detector volumes are of interest (e.g. the gas volume and the scintillator panels). Tracks and hits not associated with these volumes are removed, while the full track leading to a hit in a volume of interest is preserved, including intermediate steps through passive material, so that visualization tools retain all necessary information. This reduces output file size by a factor of 1–10, depending on the simulation configuration, without changing the observables used in the veto and TPC analyses, as checked in validation runs.

## 4.5.2 Geometry-generation infrastructure for IAXO simulations

The IAXO detector geometries used in this thesis—including the IAXO-D0 prototype, passive-shielding scans, and successive veto-layer designs—were defined through a dedicated geometry-generation workflow rather than hand-edited `Geant4 C++` code. The Geometry Description Markup Language (GDML) [85] is an XML-based format for describing detector geometries, supported natively by both `Geant4` and `ROOT`.

The complexity of the IAXO geometry (over 400 individual components in some veto iterations) made manually writing GDML impractical. A Kotlin-based domain-specific language (commonly referred to as `gdml.kt`) was developed within the collaboration to build GDML files from a higher-level, parameterized description. The author collaborated in its continued development, validation, and application to the IAXO-D0 and BabyIAXO simulation campaigns. The thesis contribution is not the invention of the Kotlin DSL itself, but the integration of that geometry-generation approach into the full simulation, analysis, veto-mapping, and visualization workflow.

The generated geometries are tracked in a dedicated Git repository, so each simulation can be associated with the commit hash of the geometry used to produce it. A hierarchical geometry description allows groups of related volumes to be enabled, disabled, or highlighted in the event viewer, connecting `gdml.kt` directly to the visualization package.

A technical challenge arose from `Geant4`'s handling of GDML assemblies: physical-volume names were converted into hashed identifiers during parsing, which prevented `restG4` from mapping simulation properties (step-size limits, production cuts) to volumes by name. Significant effort was devoted to resolving these hashed names back to human-readable representations and to implementing support for logical-volume-name references, enabling properties to be assigned consistently to all instances of a repeated component. These developments made the complex IAXO geometries usable for production simulations in which detector materials, sensitive regions, and veto-channel definitions must remain traceable across many geometry versions.

The same semantic geometry information was used to validate the veto readout mapping. When the browser event viewer highlights a panel that received an energy deposit, the displayed panel, the `Geant4` sensitive volume, the readout channel, and the analysis observable must all refer to the same detector element. Several iterations of the veto geometry and readout description were checked in this way, and the visualization package closed the loop: the Kotlin DSL produced the GDML, `restG4` transported particles through it, the detector-response chain reconstructed signals using the readout metadata, and the browser viewer made it possible to inspect whether all these representations agreed event by event.

### 4.5.3 Python interface to REST-for-Physics

REST-for-Physics inherits ROOT's PyROOT interface, which automatically generates Python bindings from the C++ class dictionaries. This allows REST-for-Physics objects to be used from Python with nearly identical syntax to the C++ API, as illustrated in Figure 4.4.

```

#include <TRestRun.h>
#include <TRestGeant4Event.h>

TRestGeant4Event* event = new TRestGeant4Event();
TRestRun* run = new TRestRun(filename);

run->SetInputEvent(event);
run->GetEntry(0);

double energy = 0;

for (const auto& track: event->GetTracks()){
    const auto hits = track.GetHits();
    for (int i = 0; i < hits.GetNumberOfHits(); ++i){
        energy += hits.GetEnergy(i);
    }
}

cout << "Total deposited energy: " << energy << " keV" << endl;

Total deposited energy: 197.705 keV

```

```

import ROOT
import REST

run = ROOT.TRestRun(filename)
event = ROOT.TRestGeant4Event()

run.SetInputEvent(event)
run.GetEntry(0)

energy = 0

for track in event.GetTracks():
    hits = track.GetHits()
    for i in range(hits.GetNumberOfHits()):
        energy += hits.GetEnergy(i)

print(f"Total deposited energy: {energy} keV")

Total deposited energy: 197.7047878262165 keV

```

Figure 4.4: Comparison of C++ and Python code computing the total energy deposited across all hits in a `restG4` simulation. Both examples produce the same result; the Python version uses PyROOT bindings automatically generated from the C++ class dictionaries.

The Python interface makes it possible to prototype analysis code quickly, integrate with the Python scientific ecosystem, and use Jupyter notebooks for interactive data exploration. The performance penalty of interpreted Python loops can be mitigated by using optimized array libraries such as Awkward Array [86] or by relying on compiled C++ backends. In the present thesis, the Python interface was used extensively for generating plots, scanning cut thresholds, and performing statistical analyses that would have been more cumbersome to develop in C++ alone.

### 4.5.4 Probability Distribution Projection cosmic-ray generator

The cosmic-ray background simulations required for this thesis launch primary particles from the atmosphere toward the detector. The conventional approach—generating particles uniformly over a large plane above the detector and discarding those that miss—is physically intuitive but highly inefficient for compact geometries like the IAXO-D0 detector surrounded by shielding.

To address this, a new cosmic-ray generator was developed and published during this work [84]. The Probability Distribution Projection (PDP) method samples directly from the subset of trajectories that intersect a sphere enclosing the geometry of interest.

For a fixed zenith angle, the allowed starting points are restricted to the ellipse obtained by projecting the enclosing sphere onto a plane tangent to it. The generator does not sample the original zenith distribution  $f(\theta)$  directly, but rather the conditional distribution  $f(\theta \mid \text{intersect sphere}) \propto f(\theta) \sec(\theta)$ , which concentrates computational effort on useful trajectories only.

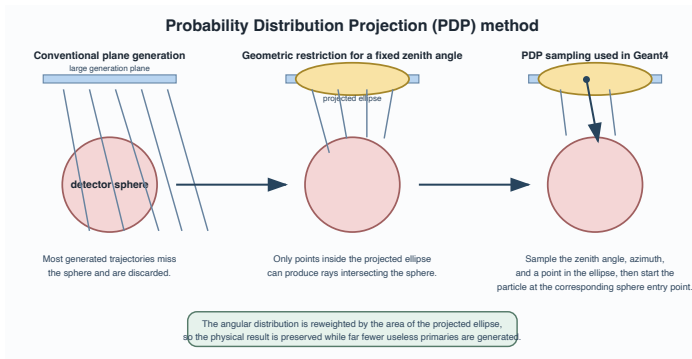


Figure 4.5: Schematic representation of the Probability Distribution Projection (PDP) cosmic-ray generator developed during this work. For a fixed zenith angle, only the region of the tangent plane whose trajectories intersect the enclosing sphere is sampled.




Figure 4.6: Summary of the IAXO-D0 benchmark reported in [84]. The left panel shows the computation yield (saved events per second) for several conventional generation-disk radii and for the PDP method. The right panel shows the corresponding recovered physical background rate. The PDP method reaches the converged rate while preserving a much higher computational efficiency.

The method was validated against the standard Monte Carlo approach using two complementary observables: the zenith-angle distribution of intersecting trajectories and the distance distribution between the intersection point and the detector axis, which probes the spatial phase-space sampled by the generator. The comparison showed

that the PDP method reproduces both observables within statistical uncertainties once the conventional Monte Carlo is run in its converged regime.

For the IAXO-D0 geometry, the PDP method improved the computation yield by a factor of approximately three compared with the fastest (and least accurate) conventional configuration. When compared at equal physical accuracy against the large generation disks needed for the conventional method to converge, the advantage reached up to a factor of about 37. This optimization was particularly valuable for the large cosmic-ray production campaigns discussed in the shielding and veto system chapter and the background model chapter, where the cumulative cost of low-efficiency simulations would otherwise have been prohibitive.

### 4.5.5 Cosmic-ray source generators

The simulation of cosmic-ray backgrounds involves two conceptually distinct tasks: generating the atmospheric secondary flux (performed with packages such as CRY [83] or CORSIKA [87]) and efficiently injecting those secondaries into the detailed IAXO **Geant4** geometry. The PDP method addresses the second task; it does not replace shower generators but provides a more efficient way of sampling the already-modeled flux around the detector.

A dedicated auxiliary program was developed to precompute the cosmic-ray secondary distributions used as input for the **Geant4** simulations [88]. The program runs CRY for a chosen site configuration, including latitude, date, altitude, and lateral generation box, and records the secondary particles crossing the generation surface. For each relevant particle species, the output is reduced to two-dimensional histograms in kinetic energy and zenith angle,  $H_i(E, \theta)$ , stored in ROOT files for muons, electrons, positrons, photons, protons, and neutrons. The histograms carry the particle-dependent source term, while the later **restG4** generator handles the geometry-dependent part of the problem: sampling a direction, choosing a valid entry point around the detector, and launching the particle into the detailed IAXO geometry.

This design deliberately decouples the atmospheric calculation from the detector transport. Instead of calling CRY inside every detector simulation, the atmospheric source term is computed once and reused across passive-shielding scans, veto-layer geometries, and production campaigns. It also makes the source term inspectable: the same ROOT histograms can be plotted, compared with EXPACS or HENSA measurements, or replaced by a different measured spectrum without changing the detector-response chain. In the final neutron studies, for example, the measured HENSA energy spectrum is passed through the same histogram-based interface, while the missing incident-direction information is supplied by the  $\sin \theta \cos^2 \theta$  angular model discussed in the veto-system chapter.

## CRY-to-REST atmospheric-source workflow

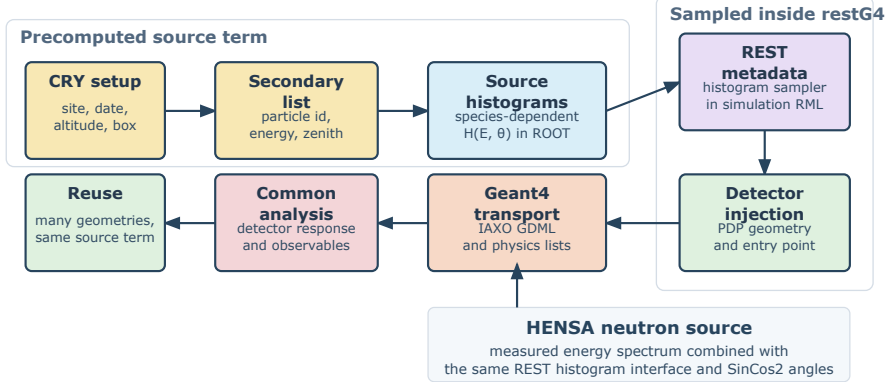


Figure 4.7: Workflow used to turn atmospheric-secondary calculations into reusable REST-for-Physics source terms. CRY is run once for a site/source configuration and the resulting secondaries are summarized as particle-specific  $H(E, \theta)$  histograms in ROOT. The `restG4` source metadata samples these histograms and combines them with the geometry-aware injection method before `Geant4` transport and the common detector-response analysis chain.

In the end, the production campaigns reported in this thesis relied predominantly on this CRY-based workflow for muons, photons, protons, electrons, and positrons. For neutrons, CRY remained essential as a generator-level cross-check and as the reference used to validate the angular prescription, while the final outdoor neutron energy source term was anchored to the HENSA measurement.

### 4.5.6 Radiation transport in gases: Garfield++

`Garfield++` is a widely used software package for simulating the behavior of gaseous detectors. Within `REST-for-Physics`, it is an optional dependency used to precompute gas-transport parameters such as electron drift velocity, diffusion coefficients, and Townsend and attachment coefficients as a function of the applied electric field. These parameters (shown in the Micromegas detector chapter, Figures 3.4 and 3.5) are then used by `REST-for-Physics` to model the drift and diffusion of ionization electrons in the TPC gas volume without performing a full `Garfield++` particle-by-particle transport, which would be computationally prohibitive.

The full `Garfield++` microscopic transport was not integrated into the `restG4` simulation chain; it was used only for generating the precomputed gas-parameter tables that the `REST-for-Physics` detector-response processes consume. Explicit integration of `Garfield++` with `Geant4`, as described in [89], remains a potential future

improvement, but was not required for the background-model studies presented in this thesis.

## 4.6 Visualization and online diagnostics

### 4.6.1 Browser-based Geant4 event viewer

REST-for-Physics provides a ROOT-based 3D event viewer using TEve, but this interface has significant limitations: CPU-based rendering leads to performance issues with complex geometries, the backend technology has not seen major updates in recent years, and the interface crashes frequently with large events.

During this thesis, a dedicated browser-based event viewer was developed for REST-for-Physics Geant4 output files. The package converts selected windows of REST-for-Physics ROOT files into a compact JSON scene representation and renders the detector geometry and event history in a web browser using `three.js`. The central design choice was to keep the heavy ROOT/REST-for-Physics dependency on the server or conversion side, while keeping the browser client simple and portable.

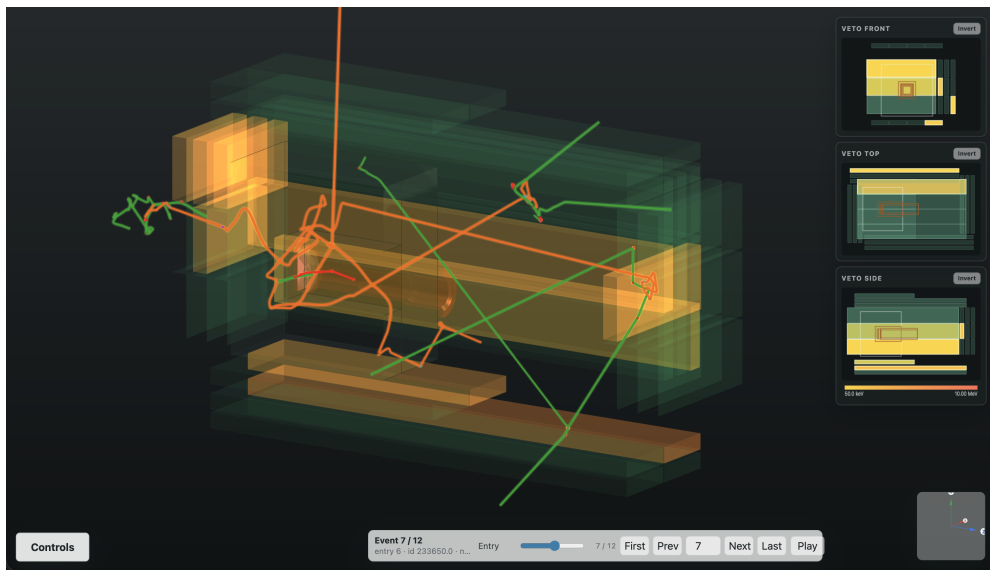


Figure 4.8: Browser-based REST-for-Physics Geant4 event viewer developed during this thesis. The main panel shows the IAXO detector geometry with simulated particle tracks from a cosmic-neutron event and highlighted scintillator panels receiving energy deposits. The four side panels provide compact orthographic veto projections of the front, left, top, and right scintillator groups, with the active detector elements visible at the center.

The viewer was developed with three primary use cases in mind:

1. **Geometry and readout validation.** The viewer renders the detector with a role-aware material model: gas volume, copper vessel, shielding, scintillator panels, cadmium layers, light guides, photomultipliers, and auxiliary elements are assigned distinct colors, transparency values, and visibility controls. This allowed systematic checks that the Kotlin-generated GDML, the `restG4` geometry import, the sensitive-volume definitions, the readout metadata, and the reconstructed veto observables were mutually consistent. If the viewer showed a veto panel highlighted by the analysis but no plausible track or capture near that panel, the problem could be traced to the mapping, timing, or reconstruction rather than to the transport simulation.
2. **Veto and delayed-capture topology debugging.** The event history is displayed as a time-dependent object: particle tracks are colored by particle type, and the user can play the event in time with a configurable hit lifetime. This was particularly useful for cosmic-neutron studies, where the relevant signature can involve a prompt cascade, a delayed neutron capture, and energy deposits in different veto layers separated by microseconds. Several projection modes were added because the IAXO geometry is elongated and surrounded by a layered veto system; the four-panel orthographic view with compact veto maps immediately shows whether an event activates opposite sides, adjacent layers, a local cluster, or a delayed capture-like pattern.
3. **Shareable event inspection and figure generation.** The deployed viewer supports browsing a curated set of ROOT files, converting only the requested event window, and caching the resulting scene. The cache makes it possible to share stable URLs pointing to a particular file, event entry, and view direction, which is useful when discussing a suspicious topology with collaborators or preparing a thesis figure. Publication-oriented export tools include a single-frame export with veto projections attached and a sequence export that samples the event timeline and writes a reproducible set of frames. The neutron event sequence shown in the background model chapter was produced with this workflow.

The development of this package was closely connected to the geometry and readout work described in Sections 4.4 and 4.5.2. In this sense, the visualization package was not only a presentation tool but also a development instrument used to find geometry mistakes, volume-orientation problems, wrong particle color mappings, missing process labels, and inconsistencies in veto-channel interpretation.

The `Phoenix` event display [90], an open-source project supported by the High Energy Software Foundation (HSF), provided a useful reference during development.

The IAXO viewer was not developed as a fork of `Phoenix`, but the `Phoenix` architecture demonstrated how modern browser-based event displays can replace traditional desktop-only visualization tools in high-energy physics.

### 4.6.2 DAQ event viewer: `feminos-viewer`

`REST-for-Physics` provides a `ROOT`-based event viewer for visualizing events throughout the processing chain, from raw signals to processed events. However, it lacks the ability to visualize events in real time as they are being acquired by the data acquisition system. During this thesis, the DAQ software was significantly refactored and collected in the `feminos-daq` repository [91], which modernized the original `mclient` program used with the FEMINOS electronics. Key improvements included storage of raw data in a regular `ROOT` file, a Prometheus exporter for external monitoring, a CMake-based build system, and a live viewer for the acquired data.




Figure 4.9: The `feminos-viewer` application developed as part of the `feminos-daq` acquisition software. The left panel shows the raw FEMINOS waveforms for a single event, with individual readout channels colored according to their position in the detector. The right panel displays the corresponding hit map, showing the physical X/Y positions of the active readout channels for the selected event. The interface includes controls for navigating events, selecting readout mappings, and switching between waveform, activity, and observables views.

A Python application, `feminos-viewer`, was developed to visualize raw data in

real time as it is being acquired. It uses `Uproot` to read the `ROOT` file and `Matplotlib` to render waveforms. It can display events from closed files as well as from files that are still being written, either by opening a local or remote file or by attaching to the output of a running acquisition process. The viewer uses readout mappings to translate signal identifiers into physical detector coordinates and includes online observables such as channel activity and energy-like summaries.

`feminos-viewer` is complementary to the browser-based `Geant4` event viewer: the former is an acquisition and waveform-quality monitoring tool, while the latter is a simulation and event-history tool. Both share the broader goal of making detector information inspectable without writing ad hoc analysis code for every diagnostic question. During data taking, `feminos-viewer` was used to identify detector and acquisition issues quickly, including malformed events, inactive or noisy channels, and inconsistencies between waveform activity and detector-coordinate mapping.

## 4.7 Monte Carlo production at scale

The background-model studies required large-scale Monte Carlo campaigns spanning multiple particle species, shielding configurations, and veto designs. This section describes the production infrastructure that made these campaigns feasible.

### 4.7.1 End-to-end simulation and analysis chain

The complete workflow used for production simulations follows a fixed sequence of stages:

1. **Source generation.** Atmospheric secondary distributions (muons, neutrons, photons, electrons, protons) are computed with `CRY` at the relevant site coordinates and stored as `ROOT` histograms.
2. **restG4 transport.** The source histograms are sampled by the primary generator inside `restG4`, and particles are transported through the IAXO detector geometry using `Geant4`. The output is a `TRestGeant4Event` stored in the `EventTree`.
3. **Detector response.** `Geant4` energy deposits are converted into `TRestDetectorHitsEvent` and `TRestDetectorSignalEvent` representations, where quenching, light attenuation, diffusion, and smearing are applied.
4. **Raw-signal emulation.** Signals are converted to `TRestRawSignalEvent` through waveform shaping, sampling, trigger positioning, and dynamic-range effects.

5. **Reconstruction.** Waveforms are reconstructed back into detector hits and tracks, yielding event representations directly comparable to experimental data.
6. **AnalysisTree observables.** Each processing stage appends scalar observables (hit multiplicities, energies, positions, veto peak times and amplitudes) to the `AnalysisTree`.
7. **Cuts and background-rate extraction.** The observables are used to apply fiducial, topological, and energy cuts; the surviving events are normalized by the simulated exposure to produce background rates.

This chain is configured through two RML files stored in the collaboration repository `iaxo-simulations` [76]: a source-specific `simulation.rml` and a common `analysis.rml`. This separation made it possible to treat source generation, transport, response emulation, and final selection as distinct but reproducible stages.

## 4.7.2 Production workflow with HTCondor

Simulations were orchestrated using `HTCondor`, a workload management system designed for High-Throughput Computing (HTC) [92]. Unlike traditional HPC environments that optimize for instantaneous floating-point performance, `HTCondor` maximizes total computational work over long periods through dynamic matchmaking between job requirements and available resources.

A dedicated `Python` script, `restG4ToCondor.py`, was developed to facilitate the submission of `restG4` jobs to the `HTCondor` system. The script automates the creation of job description files and manages output data including log files, error reports, and the resulting `ROOT` files. Its command-line interface extends the `restG4` interface, making it easy for users familiar with `restG4` to adapt to batch submission.

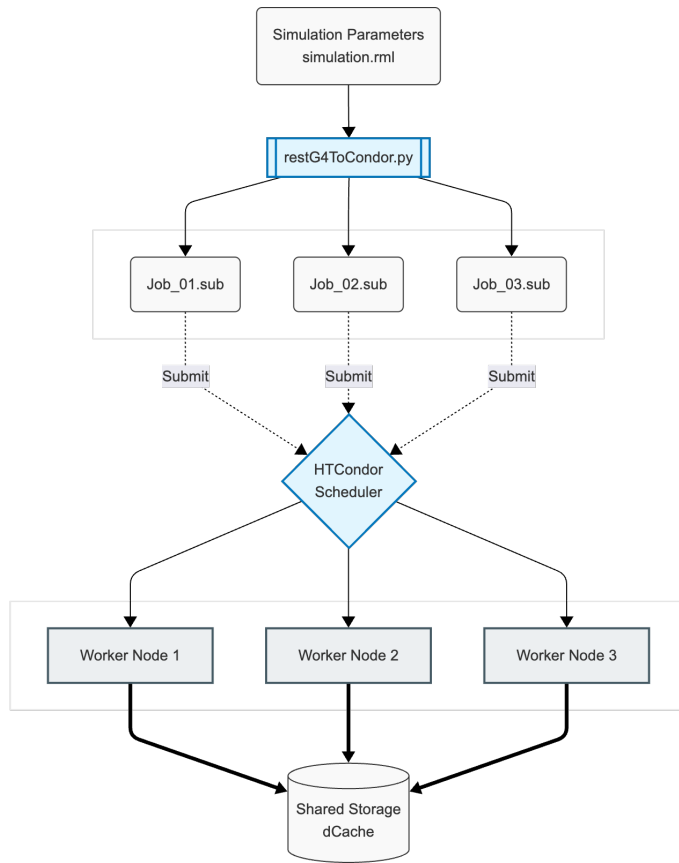


Figure 4.10: Schematic representation of the HTCondor-based Monte Carlo data production workflow using the `restG4ToCondor.py` script.

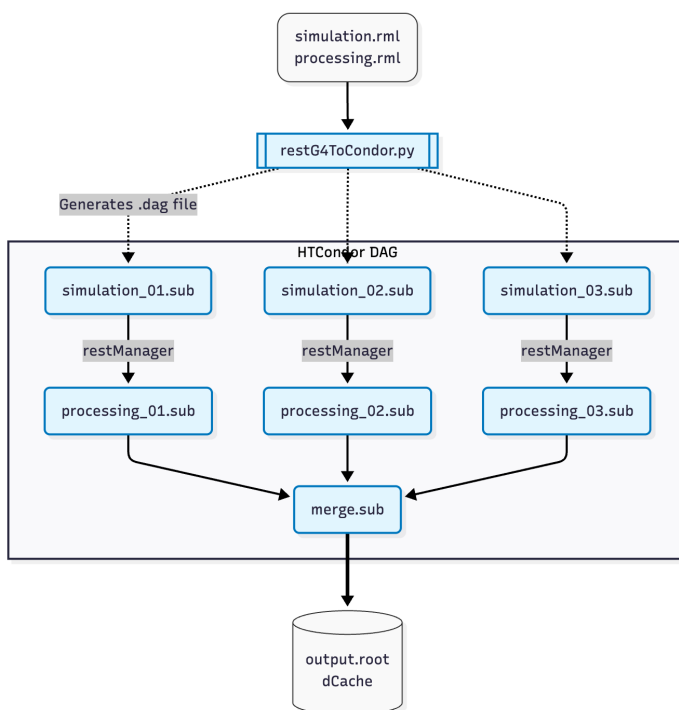


Figure 4.11: Schematic representation of the HTCondor-based Monte Carlo workflow with processing and merging. Each simulation job is followed by a processing job, and a final merging job combines all processed files into a single output file.

Table 4.4 summarizes the typical workflow steps for a large Monte Carlo production campaign, from job splitting to the final extraction of analysis observables.

Table 4.4: Typical HTCondor production workflow steps.

Step	Action	Key configuration
Split / submit	<code>restG4ToCondor.py</code> splits the total number of events or time budget across $N$ independent jobs and submits them to HTCondor (optionally via DAGMan)	<code>--n-jobs</code> , <code>--entries</code> , <code>--time</code>
Simulate	Each job runs <code>restG4</code> with the specified RML, geometry, source configuration, and environment variables	<code>--rml</code> , <code>--geometry</code> , <code>--env</code>
Process	<code>restManager</code> applies the analysis RML to the output of each simulation job	<code>--rml-processing</code>
Merge (optional)	<code>restManager</code> combines the processed outputs into a single file	<code>--merge</code>
Analyze	Observables from the <code>AnalysisTree</code> are extracted and used for cuts, efficiency studies, and background-rate calculations	Python / ROOT macros

The multi-job productions used in this thesis were run on the National Analysis Facility (NAF) at DESY, referred to here as NAF-IAXO, which provides a centralized computing environment with dCache storage [93] and a HTCondor batch system. Typical campaigns involved 300 parallel jobs of 2–8 hours each, producing tens of thousands of saved events per particle species. The production infrastructure was validated with cosmic-neutron and cosmic-muon campaigns totaling approximately 40 000 final events across 600 Condor jobs.

## 4.8 `geant4-python-application`: Python prototyping with Geant4

`geant4-python-application` is a project developed during this thesis to make small Geant4 studies accessible from Python notebooks and scripts. It is a teaching and prototyping tool, not a replacement for the `restG4` production pipeline.

The standard Geant4 workflow is based on user-written C++ applications. Although Python bindings and third-party Python interfaces exist, they do not provide the specific `restG4`-like, pip-installable, process-isolated application wrapper targeted here for notebooks, teaching, and rapid prototyping. The project wraps a generic user-configurable Geant4 application using Pybind11 to generate Python bindings. Because a Geant4 application cannot be reinitialized once started, process isolation was implemented with Python's `multiprocessing` module, allowing each worker process to run its own independent Geant4 instance.

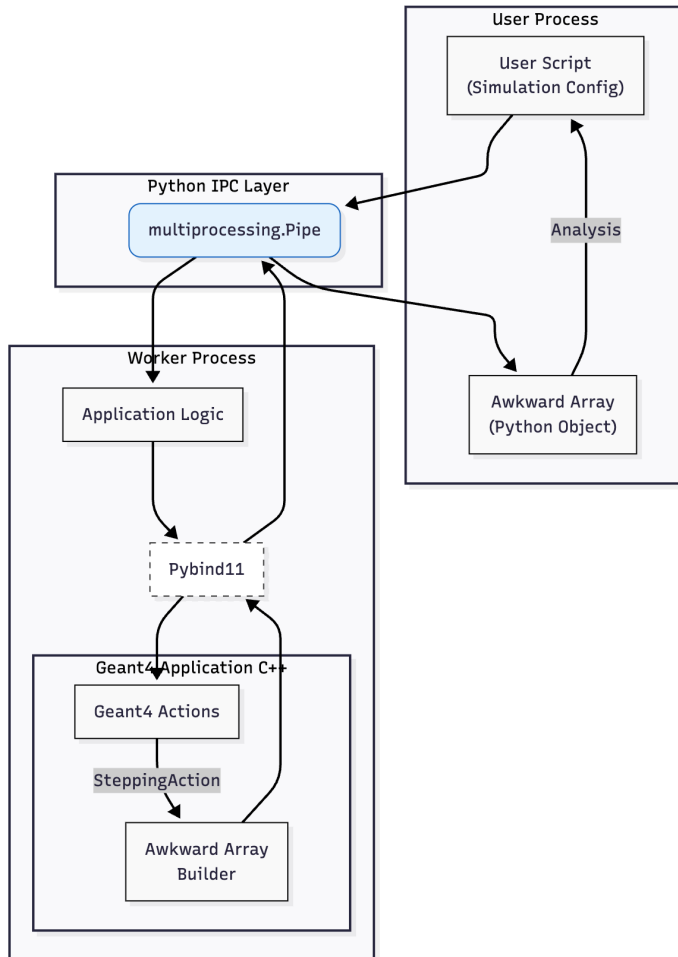


Figure 4.12: Architecture of the `geant4-python-application` project, showing how a `Geant4` application is wrapped via `Pybind11` and how independent worker processes provide isolation for repeated simulations.

The following concrete outputs were demonstrated with this tool:

- A Jupyter notebook computing the X-ray attenuation length in a given material by running 5 000 `Geant4` events and analyzing the results with Awkward Array [86] (Figure 4.13).
- Parallelization of independent simulations through Python’s `concurrent.futures` interface, with each task launching an isolated `Geant4` worker process rather than attempting to reinitialize `Geant4` inside a single process.
- Distribution as a pip-installable package on PyPI (Python Package Index), with

pre-compiled binary wheels for Linux, Windows, and macOS across multiple Python versions, built via a GitHub Actions CI pipeline.

```
def compute(i: int):
    thickness = thicknesses[i]
    geometry = construct_geometry_gdml(material="G4_Pb", thickness_mm=thickness)

    with Application(gdml=geometry) as app:
        app.command("/gun/particle gamma")
        app.command("/gun/energy 477.6 keV")
        app.command("/gun/direction 0 0 1")
        app.command("/gun/position 0 0 -1 m")

        max_total_count = 100000

        total_counts = 0
        valid_counts = 0

        while total_counts < max_total_count:
            # batch size can be adjusted to optimize performance
            events = app.run(5000)
            total_counts += len(events)
            valid_counts += np.sum(events.energy_in_volume("target") > 0)
            # check if relative error is below 1%
            if valid_counts > 0 and np.sqrt(valid_counts) / valid_counts < 0.01:
                break

        ratios[i] = valid_counts / total_counts

    if valid_counts > 0:
        # add some upper limit to the error (68% confidence level)
        ratios_error[i] = np.sqrt(valid_counts) / total_counts
    else:
        ratios_error[i] = 1.14 / total_counts
```

Figure 4.13: A code sample from a Jupyter notebook using the `geant4-python-application` package to compute the X-ray attenuation length. The `compute` method runs 5000 Geant4 events, and the result is obtained through Awkward Array operations.

While `geant4-python-application` was not used for the production simulations of this thesis, it demonstrated the feasibility of embedding Geant4 inside a Python workflow and contributed to the broader goal of improving the accessibility and maintainability of HEP simulation software.

## 4.9 Summary and role in the thesis

This chapter has described the software framework and tools that enabled the background-model and veto-system studies presented in later chapters. The contributions can be summarized as follows:

- `REST-for-Physics` provided a reproducible event-processing architecture in which the same analysis chain could be applied to Monte Carlo and experimental data, ensuring that observables remained aligned across both branches.

- **Readout metadata** linked detector geometry to electronics channels and analysis observables, making it possible to move from raw waveforms to physically meaningful detector hits and veto peaks within a single framework.
- **restG4 and Geant4Lib** provided the **Geant4** transport layer, extended with multi-threading, track pruning, sub-event splitting, and geometry-name resolution to support the demanding IAXO detector geometries and the throughput required for background-model campaigns.
- **The PDP cosmic-ray generator** accelerated the simulation of atmospheric secondaries by concentrating computation on trajectories that actually intersect the detector, enabling large-scale cosmic-background productions that would otherwise have been computationally prohibitive.
- **The browser-based Geant4 event viewer** served as a validation tool for geometry, readout mapping, and veto topology, and as a means of producing publication-quality event figures and sharing event topologies with collaborators.
- **feminos-daq and feminos-viewer** provided live waveform monitoring and detector-coordinate mapping for the acquisition system, supporting real-time data-quality assessment during IAXO-D0 and IAXO-D1 operation.
- **The HTCondor production infrastructure** enabled large Monte Carlo campaigns on the NAF-IAXO batch system, with automated job submission, processing, merging, and output staging.
- **Python tools** (Uproot + fsspec integration, PyROOT, and `geant4-python-application`) supported interoperability between the C++ simulation framework and the Python analysis ecosystem, enabling rapid prototyping, monitoring, and interactive data exploration.

Collectively, these components formed the computational backbone of the thesis. They allowed the background model to be constructed from reproducible simulations with validated geometry, calibrated detector response, and statistically meaningful event samples, and they ensured that the veto tagging strategies could be developed and tested within the same analysis environment used for the experimental data.



---

## Chapter 5

# Shielding and Veto System

5.1	Motivation: surface backgrounds and veto requirements . . . . .	108
5.1.1	Design requirements and constraints . . . . .	108
5.1.2	Surface cosmic-ray source term . . . . .	110
5.1.3	Relevant interaction processes in the shielding and veto materials	116
5.1.4	Interaction of high-energy neutrons and photons with the detector	124
5.1.5	Effect of detector inclination during solar tracking . . . . .	129
5.2	Limits of passive shielding . . . . .	131
5.2.1	Lead-thickness scan: electromagnetic suppression versus neutron multiplication . . . . .	131
5.2.2	Borated-HDPE moderation studies . . . . .	133
5.3	Neutron-sensitive active veto concept . . . . .	136
5.3.1	Plastic scintillators . . . . .	136
5.3.2	From prompt muon veto to neutron-sensitive stages . . . . .	139
5.4	Multilayer veto optimization . . . . .	145
5.5	Final veto geometry and readout . . . . .	148
5.5.1	Prototype construction and commissioning . . . . .	151
5.5.2	Electronics and readout . . . . .	152
5.6	Waveform-level veto observables . . . . .	154
5.6.1	Boosted veto BSD selection . . . . .	157
5.7	Simulation performance . . . . .	161
5.8	Experimental validation with IAXO-D0 data . . . . .	167
5.9	Systematic limitations and remaining work . . . . .	183
5.9.1	Design-level systematics . . . . .	183
5.9.2	Limitations of the neutron-enriched comparison . . . . .	184
5.10	Summary and outlook . . . . .	185

## Introduction

This chapter describes the design, optimization, construction, and validation of the shielding and active veto system developed for the BabyIAXO/IAXO-D0 Micromegas detector line. The central background problem addressed here is the surface-level cosmic-ray-induced background, in particular the high-energy neutron component that cannot be sufficiently suppressed by passive lead shielding alone. The chapter follows the evolution of the veto concept from the initial source-term studies and passive-shielding simulations to the final multilayer scintillator–cadmium design, waveform-level veto observables, prototype construction, commissioning, and comparison with experimental data.

The design strategy is driven by the fact that high-energy cosmic neutrons are difficult to detect directly. Instead, the lead shielding converts part of the primary neutron flux into showers of lower-energy secondary neutrons and photons. The veto system is therefore optimized not only as a muon veto, but also as a neutron-sensitive active shield exploiting prompt scintillation signals, delayed neutron-capture signatures, and event multiplicity.

The general `REST-for-Physics/restG4` simulation infrastructure and the full cosmic-source-generation workflow are described in the software and background-model chapters. Here, only the aspects directly relevant to the shielding and veto design are summarized. The radiation transport studies discussed below were performed with `Geant4` through `restG4` [78–81].

The chapter is organized around four linked claims. First, the surface source term makes high-energy neutrons a design-limiting background for a low-threshold Micromegas detector. Second, passive lead shielding remains essential for photons but cannot by itself remove the neutron component, and can even transform it into a more distributed secondary shower. Third, repeated scintillator–cadmium stages provide a practical way to tag that shower through prompt, delayed, and multiplicity-rich veto observables. Finally, the commissioned IAXO-D0 prototype validates the same waveform-level logic in data, although the absolute neutron fraction remains limited by source-normalization and detector-response systematics.

## 5.1 Motivation: surface backgrounds and veto requirements

### 5.1.1 Design requirements and constraints

Before presenting the simulations and prototype results, it is useful to state explicitly the requirements that shaped the veto concept. The active shielding was not designed

as an independent detector placed around the Micromegas chamber. It is part of the same background-rejection chain as the passive shield, the X-ray entrance geometry, the Micromegas reconstruction, and the waveform-level veto analysis. The design problem is therefore constrained by three coupled goals: preserve the low-energy X-ray acceptance, suppress the surface cosmic-ray background, and remain compatible with the available scintillator modules, mechanical envelope, and readout electronics.

<b>Requirement / constraint</b>	<b>Consequence for the veto concept</b>	<b>Practical anchor</b>
Surface-level operation	Cosmic muons and neutrons must be rejected without relying on underground overburden. The active system must therefore tag both prompt charged-particle activity and neutron-induced secondary showers.	Sea-level source terms at Zaragoza/DESY-like latitudes; no overburden assumed in the baseline design studies.
Low-energy X-ray acceptance	The veto must reject background while preserving signal-like Micromegas candidates in the keV region. Veto decisions are therefore applied as analysis cuts, not as a requirement that an event have veto activity.	Prototype validation uses the 2–7 keV analysis window within the broader IAXO X-ray region.
Lead shield required for gamma suppression	A thick lead castle suppresses external photons but also converts high-energy neutrons into showers of lower-energy neutrons, photons, and charged fragments. The veto must tag this shower, not only the primary neutron.	Baseline passive shield of about 20 cm lead around the detector.
Open detector geometry	The X-ray beam path, services, and mechanical clearances prevent a perfectly hermetic passive shield. The active layers must compensate for these unavoidable openings and for particles entering through nonideal paths.	Dedicated checks include finite apertures, service-side exposure, and detector inclination during tracking.
Neutron time structure	Useful neutron tags are not purely prompt. The readout and analysis must retain delayed capture-like activity while controlling random coincidences.	Capture-time scale $\tau \approx 45 \mu\text{s}$ ; prototype-like veto window of $100 \mu\text{s}$ .
Available scintillator and readout hardware	Panel dimensions, segmentation, light collection, channel mapping, and timing settings are constrained by re-used scintillators, PMTs, and the AGET-based front end.	Final simulated design uses 59 panels; the commissioned prototype records up to 57 veto signals.
Accidental-veto and live-time control	Thresholds, multiplicity requirements, and delayed windows must improve rejection without turning accidental activity into excessive signal loss.	Advanced veto cuts are benchmarked against calibration data and preserve about 97% of events relative to the prompt-veto stage.

Table 5.1: Main design requirements and constraints that guided the BabyIAXO/IAXO-D0 veto concept. The veto is optimized as part of a coupled passive-shield, Micromegas, and waveform-analysis system rather than as a standalone counter.

Several of these requirements pull in opposite directions. Increasing the lead thickness improves gamma attenuation but increases the importance of neutron multiplication in the shield. Extending the delayed window recovers more capture-

like activity but raises the accidental-coincidence probability. Increasing segmentation improves topology and channel-level diagnostics but is limited by the number of available readout channels and by the inherited scintillator geometry. The rest of this chapter follows these trade-offs in order: first the surface source term and material interactions, then the limits of passive shielding, and finally the active scintillator–cadmium veto response at waveform level.

### 5.1.2 Surface cosmic-ray source term

The veto design studies rely on a common description of the atmospheric-secondary background at surface level. Only a short summary is given here, since the `CRY-to-REST` source-generation workflow is described in Section 4.5.5 and the component normalization is part of the broader background-model effort. For the veto studies, the relevant secondary populations are muons, neutrons, photons, electrons, and protons. The photon, electron, proton, and mixed-secondary source-term comparisons were generated primarily with `CRY` [83], stored as two-dimensional energy–zenith histograms, and then reused as external sources in the `restG4` simulations. For muons, the production simulations use the correlated sea-level `CosmicMuons` parameterization implemented in `REST`, while `CRY` provides an independent generator-level cross-check. For neutrons, the final surface simulations use the measured HENSA outdoor spectrum discussed below as the nominal energy source term, while `CRY` and `EXPACS` provide spectral and angular guidance. This separation allowed the same atmospheric source term to be propagated through many alternative shielding and veto geometries.

The `CRY` source terms correspond to sea level, with latitude settings appropriate to Zaragoza (41.65° N) and DESY (53.34° N), a reference date of 2024-01-01 for the solar-modulation setting, and the maximum available lateral generation box of 300 m. For the veto studies, the most relevant output of this stage is the spectral shape of the atmospheric secondaries reaching the laboratory level, rather than the exact event-by-event ancestry of the primary cosmic ray.

The histogram and formula source terms generated in this way were reused as inputs to the detector simulations, which made it possible to compare alternative shielding layouts and veto geometries under a common source definition. The resulting `CRY`-based distributions were checked against `EXPACS`, the sea-level neutron measurements of Gordon *et al.*, and the HENSA measurement campaign in Zaragoza [94–97]. This comparison is used in two different ways: `CRY` and `EXPACS` establish a stable generator-level reference, while HENSA anchors the measured neutron field used for the final neutron simulations.




Figure 5.1: Atmospheric-secondary source terms from EXPACS and CRY, shown as flux per logarithmic energy interval in units of  $\text{cm}^{-2} \text{s}^{-1}$ . The same particle legend is used in both panels, with charge-separated lepton components when available. The shaded interval marks the 1 MeV–100 GeV range used for the neutron-focused veto design studies.

The CRY and EXPACS distributions are in good agreement for the broad energy range relevant to the veto studies, which provides confidence that the simulated source term captures the main features of the atmospheric-secondaries flux. This agreement is especially important for cosmic neutrons, whose absolute rate is difficult to normalize but whose spectral shape largely determines the energy region that dominates the detector response.

**Muon source term.** The production muon simulations use the *CosmicMuons* source in *restG4*, which implements the sea-level parameterization of Guan *et al.* [98]. This model is a Gaisser-like atmospheric-muon flux with an effective zenith-angle correction, sampled as a two-dimensional distribution in kinetic energy and zenith angle rather

than as independent one-dimensional marginals [99]. This is useful for a surface veto because inclined and high-energy muons are not equivalent to vertical low-energy muons: both the detector path length and the shielding traversal change with angle. `CRY` is therefore kept as a consistency check and as the source model for several historical scans, while the final muon production uses the correlated `CosmicMuons` input shown in Fig. 5.2.

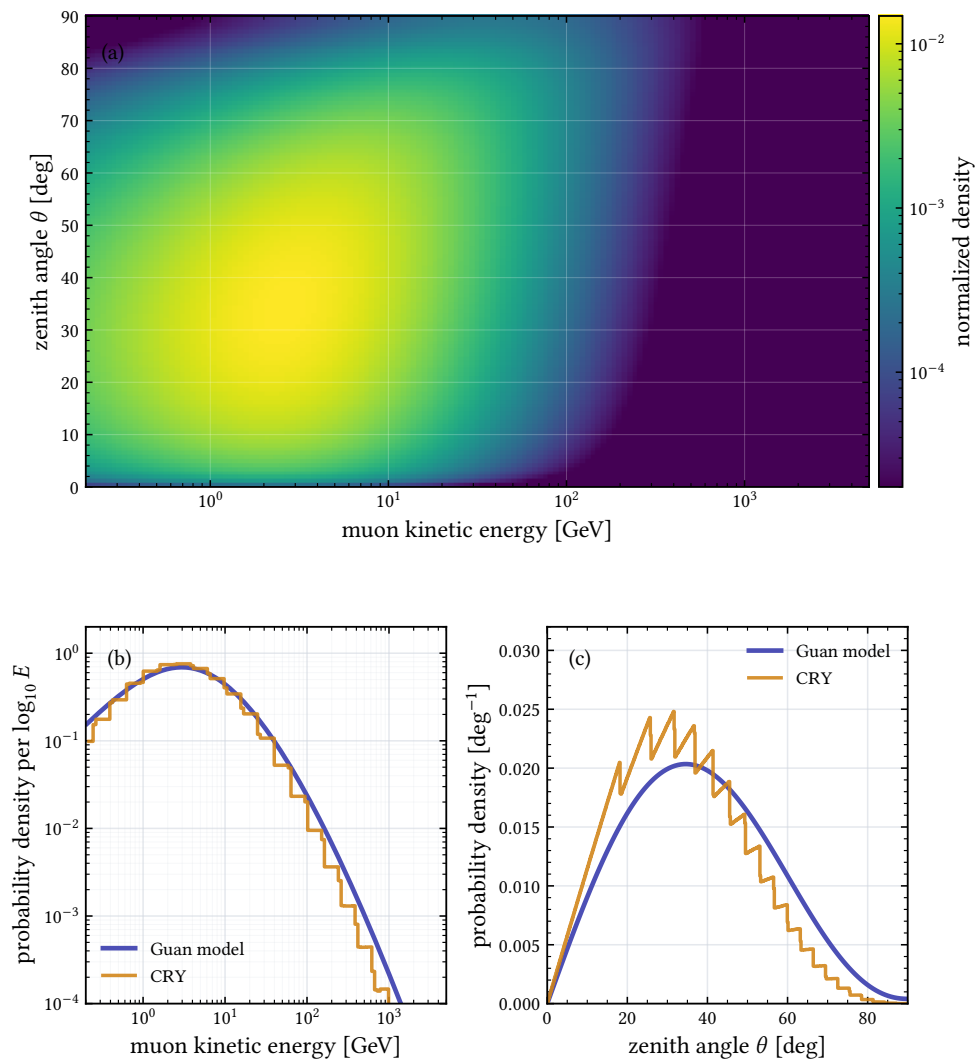


Figure 5.2: Muon source model used for the production `restG4` simulations. Panel (a) shows the normalized two-dimensional Guan sea-level parameterization sampled by the `CosmicMuons` generator, as a function of kinetic energy and zenith angle. Panels (b) and (c) compare its one-dimensional marginals with the `CRY` muon source used as a generator-level cross-check.

**HENSA neutron source term.** The nominal neutron source used in the final surface simulations is the outdoor HENSA spectrum unfolded up to 10 GeV. HENSA is an extended-energy neutron spectrometer based on moderated  $^3\text{He}$  counters and the Bonner-sphere principle; it measures the inclusive neutron field and obtains the energy spectrum through an unfolding of the detector count rates [100]. This point is important

for BabyIAXO: at ground level, the neutron field is not a pure atmospheric-cascade component, but also contains neutrons produced, moderated, and scattered in the surrounding soil, concrete, and laboratory structures. The HENSA outdoor spectrum therefore includes the environmental, or ambient, neutron component to which the measurement technique is sensitive, and is closer to the field that a surface experiment sees than a generator-only cosmic-neutron spectrum.




Figure 5.3: Outdoor HENSA neutron spectrum unfolded up to 10 GeV and used as the nominal neutron-energy source term in the surface simulations [97]. The ordinate is the bin-integrated flux converted to  $d\Phi/d \log_{10} E$ . The low-energy part is an inclusive measured ambient field, while the MeV and high-energy regions indicate the spectral components most relevant for moderation, lead-induced secondary production, and veto response.

Since HENSA provides the energy spectrum but not the incident direction of each neutron, the spectrum is paired in `restG4` with the analytic downward angular model

$$p(\theta) = 3 \sin \theta \cos^2 \theta, \quad 0 \leq \theta \leq \pi/2, \quad (5.1)$$

where  $\theta$  is measured from the downward vertical axis. This form corresponds to a  $\cos^2 \theta$  zenith intensity multiplied by the  $\sin \theta$  solid-angle factor. It is implemented through the `SinCos2` angular generator and was validated with a geantino sample. For CRY-based source terms, the simulations use the sampled two-dimensional energy–zenith histograms directly, while for HENSA the measured energy spectrum is combined with Eq. 5.1. The comparison in Fig. 5.4 shows that this analytic model reproduces the

marginal neutron angular distribution obtained from CRY to the accuracy needed for the veto-design source term.




Figure 5.4: Zenith-angle model used for the HENSA-driven neutron simulations. The analytic  $p(\theta) = 3 \sin \theta \cos^2 \theta$  law is compared with a `restG4 SinCos2` validation sample and with the marginal neutron zenith distribution sampled from CRY. The agreement justifies using the analytic angular law for HENSA, which constrains the neutron energy spectrum but not the event-by-event incidence direction.

Label used below	Source and geometry role	Purpose in the chapter
Historical CRY scans	CRY atmospheric secondaries propagated through parameterized or early shielding geometries.	Establish the original passive-shielding and geometry trends under a common generator-level source model.
Final HENSA neutron scans	Outdoor HENSA neutron spectrum up to 10 GeV, combined with the validated $\sin \theta \cos^2 \theta$ angular law.	Provide the nominal neutron source term for the layer optimization and neutron-sensitive veto studies.
Final muon production	Correlated sea-level CosmicMuons parameterization in <code>restG4</code> , cross-checked with CRY.	Quantify prompt muon rejection and validate the timing alignment of the waveform analysis.
Prototype data	Commissioned 57-signal IAXO-D0 surface veto data with calibration-trigger and background-trigger samples.	Test the waveform-level prompt, delayed, and multiplicity observables under real accidental-veto conditions.

Table 5.2: Simulation and validation taxonomy used in this chapter. The labels separate historical design scans from the final HENSA/CosmicMuons productions and from the experimental prototype validation.

### 5.1.3 Relevant interaction processes in the shielding and veto materials

The veto concept is motivated by the different roles played by the materials surrounding the Micromegas detector. The relevant question is not only whether a material has a large cross-section for a given particle, but whether the resulting interaction helps to reduce, transform, tag, or accidentally produce a signal-like event. For this reason, the material discussion is kept focused on the components that directly determine the shielding and veto response.

Component	Where and dominant interactions	Design relevance
<b>Lead</b>	<i>Location:</i> passive castle around the detector. <i>Interactions:</i> photon attenuation; fast-neutron elastic, inelastic, and neutron-emission channels.	Suppresses external gamma radiation, but can generate secondary neutron showers.
<b>Plastic scintillator</b>	<i>Location:</i> segmented veto panels outside the lead. <i>Interactions:</i> neutron scattering on H and C; charged-particle stopping power; scintillation yield and quenching.	Moderates secondary neutrons and provides prompt veto signals from charged particles and recoil protons.
<b>Cadmium</b>	<i>Location:</i> thin sheets between scintillator stages. <i>Interactions:</i> thermal and epithermal neutron capture, dominated by $^{113}\text{Cd}$ in natural Cd.	Converts moderated neutrons into delayed gamma-cascade signatures near the scintillators.
<b>Copper</b>	<i>Location:</i> chamber, pipe, supports, and readout structures. <i>Interactions:</i> photon attenuation and fluorescence; neutron scattering and activation.	Radiopure structural material, but also a possible source of X-ray lines near the analysis region.
<b>Detector gas</b>	<i>Location:</i> Micromegas conversion and drift volume. <i>Interactions:</i> X-ray photoabsorption; neutron-induced nuclear recoils; ionization yield and quenching.	Signal medium and the place where neutron-induced recoils can mimic compact X-ray events.
<b>Mylar/Al window</b>	<i>Location:</i> X-ray entrance window and cathode. <i>Interactions:</i> transmission in the 1–10 keV region.	Controls the low-energy signal efficiency and threshold behavior.
<b>Borated HDPE</b>	<i>Location:</i> passive neutron-shielding option. <i>Interactions:</i> H moderation and $^{10}\text{B}$ neutron capture.	Tested as an auxiliary passive moderator/absorber, but insufficient by itself for the high-energy neutron component.

Table 5.3: Detector and shielding materials most relevant to the veto design. The table emphasizes the interaction quantities that motivate the simulations, rather than attempting to provide a complete inventory of all mechanical materials in the setup.

Lead illustrates the central compromise. Its high density and atomic number make it an efficient attenuator for environmental photons in the keV–MeV range, as quantified by photon attenuation data such as NIST XCOM [55]. For fast cosmic neutrons, however, lead is not an efficient moderator. In an elastic collision with a nucleus of mass number  $A$ , the maximum fraction of neutron kinetic energy that can be transferred is  $4A/(1+A)^2$ . This fraction is unity for hydrogen, about 0.28 for carbon, and only about 0.019 for lead. Consequently, elastic scattering in lead removes little energy per collision, while inelastic and neutron-emission reactions become important at MeV energies. The nuclear data underlying this interpretation were checked against

the data actually used by the **Geant4** high-precision neutron transport. For natural lead, the reference curves were computed from the ENDF/B-VIII.0 incident-neutron evaluations by weighting the  $^{204}\text{Pb}$ ,  $^{206}\text{Pb}$ ,  $^{207}\text{Pb}$ , and  $^{208}\text{Pb}$  isotope files by their natural abundances [101]. The **Geant4** comparison points were extracted from the G4NDL4.6 high-precision neutron tables distributed with **Geant4** 11.0.3, using the same natural-abundance weighting [102, 103]. Elastic, inelastic, and capture channels were read from the corresponding HP cross-section tables, while the  $(n, 2n)$  contribution was read from the HP inelastic final-state channel because **Geant4** treats it as a channel inside the inelastic process rather than as an independent hadronic process. This procedure isolates the physics-list nuclear data from detector-geometry effects; it is therefore a validation of the transport inputs, not a replacement for the full detector simulations. The same ENDF/B-VIII.0 versus G4NDL4.6 procedure was also applied to the two active-veto materials that control neutron tagging after the lead shower has formed: BC408, approximated as  $\text{C}_9\text{H}_{10}$ , and natural cadmium. Those material-comparison studies are kept as supplementary validation plots in appendix 6.7, figures 6.44 and 6.45, because they support the scintillator-cadmium design choice without changing the main lead-multiplication argument.




Figure 5.5: Natural-lead neutron cross sections from ENDF/B-VIII.0 evaluated data compared with the G4NDL4.6 high-precision neutron tables used by Geant4 11.0.3. Solid curves show the evaluated reference; open markers show the corresponding Geant4 HP tables. The lower panel gives the pointwise Geant4/ENDF ratio. Total, elastic, inelastic, and  $(n, 2n)$  channels agree at the level needed for the veto-design interpretation, confirming that the simulated lead-induced neutron multiplication is driven by the same fast-neutron structure seen in evaluated nuclear data.

The event display in figure 5.22a shows the transport-level origin of the veto signature, but the veto does not identify the incoming primary neutron directly. The operational signature is instead the correlated secondary shower produced after the lead shielding has converted a hard neutron into softer neutrons, photons, and charged secondaries. This motivates a detector response that preserves two complementary components: prompt scintillation from charged secondaries and recoil protons, and delayed activity from neutrons that moderate and capture in cadmium. The concept is summarized in figure 5.22b.




Figure 5.6: Elastic-scattering moderation metrics for hydrogen, carbon, cadmium, and lead. The comparison shows why hydrogen-rich plastic scintillator is much better suited than lead for slowing secondary neutrons: a neutron can transfer all its kinetic energy to hydrogen in a single head-on collision, whereas the maximum transfer to lead is only about 2%.

Cadmium enters at the opposite end of the neutron-energy scale. It is not a shield for the primary MeV–GeV cosmic neutrons. Its purpose is to absorb neutrons after they have been slowed in the scintillator and surrounding low- $Z$  materials. Natural cadmium is especially effective in this regime because the  $^{113}\text{Cd}$  isotope has a very large thermal neutron-capture cross-section. A 1 mm natural-cadmium sheet is already optically thick for thermal neutrons: using  $\rho_{\text{Cd}} = 8.65 \text{ g cm}^{-3}$ ,  $A_{\text{Cd}} = 112.4 \text{ g mol}^{-1}$ , and a thermal natural-cadmium capture cross-section of about  $2.5 \times 10^3$  barn gives an uncollided transmission of order  $10^{-5}$  through 1 mm at 0.025 eV. This is consistent with the practical use of millimeter-scale cadmium as a thermal-neutron cut-off absorber [104]. The design requirement is therefore not to make the cadmium thicker, but to place it where enough neutrons have first been moderated and where the resulting capture gamma cascade can be detected by nearby scintillators. The corresponding Geant4 nuclear-data cross-check for natural Cd is shown in appendix figure 6.45; the same appendix also includes the BC408 scintillator comparison in figure 6.44, documenting the elastic, inelastic, and capture inputs used for moderation and prompt scintillator response.




Figure 5.7: Thermal-region cadmium capture and estimated uncollided transmission through a 1 mm natural-cadmium sheet. The curves use the  $1/v$  behavior anchored to the ENDF/B-VIII.0 thermal capture scale for natural Cd and  $^{113}\text{Cd}$  [101]; the transmission estimate shows that 1 mm is already sufficient for thermalized neutrons. The figure should not be interpreted as a fast-neutron shielding curve.

The same capture process also explains why cadmium can produce an observable veto signature rather than merely remove moderated neutrons from the transport. Evaluated prompt-gamma data for  $^{113}\text{Cd}$  show a cascade of characteristic photons after neutron capture, with strong lines in the few-hundred-keV to MeV range [105, 106]. Figure 5.8 compares those evaluated prompt-gamma lines with the photon energies emitted by **Geant4** in the HENSA-driven three-layer veto simulation. The simulation records 99,021 cadmium captures and 372,611 emitted prompt-gamma tracks, corresponding to an average of 3.76 photons and 8.73 MeV emitted as gamma energy per capture. The relevant design consequence is that the cadmium sheets should be placed close to active scintillator, because the useful signal is the transport and partial absorption of this gamma cascade in nearby panels. At the same time, this comparison is treated as a validation diagnostic rather than as a claim that the default **Geant4** capture model reproduces all cadmium line intensities. The high-energy  $\sim 9.04$  MeV  $^{113}\text{Cd}$  line is present in evaluated data, but with a much smaller listed production yield than in the default simulated cascade, by roughly two orders of magnitude in this line-window comparison. This is consistent with the more general limitation discussed in recent work on data-driven neutron-capture cascade models: the standard **Geant4** capture/de-excitation modules are useful transport tools, but they do not necessarily reproduce all measured prompt-gamma lines and correlations

in the MeV range [107]. For this reason, the cadmium study includes a dedicated `Geant4` physics-model comparison rather than a single production configuration. The comparison separates three choices that affect different parts of the calculation. First, the neutron transport and capture probabilities are governed by the high-precision neutron data used by the `QGSP_BIC_HP` hadronic list and the associated elastic HP model [102, 103]. Second, the exact isotope-data policy is controlled by requiring `ParticleHP` to skip missing isotope data instead of silently falling back to a nearby element or natural composition. Third, the prompt-gamma cascade after capture can be generated either from the default `ParticleHP` final-state data or through the `PhotonEvaporation` de-excitation model [102]. The last choice is the most relevant one for cadmium, because it changes how the capture  $Q$ -value is partitioned among emitted photons while leaving the preceding neutron moderation and capture geometry largely unchanged.

The campaign is enabled through an extension of the `restG4` physics-list interface that forwards the `usePhotoEvaporation` and `skipMissingIsotopes` options to `G4ParticleHPManager` before the hadronic neutron processes are constructed. This distinction is important for the veto optimization because the design does not rely on a single gamma line being simulated with spectroscopic accuracy. It relies on three more robust ingredients: moderation of secondary neutrons in hydrogen-rich material, high capture probability in cadmium once those neutrons are thermalized, and emission of an MeV-scale photon cascade close enough to nearby scintillators to deposit detectable energy. The model comparison therefore tests whether the veto conclusion survives a different de-excitation prescription, while also identifying which line-level predictions should be treated as a physics-model systematic. In the partial validation campaign, the default and strict-`ParticleHP` variants give nearly the same cascade: 3.88 and 3.87 photons per cadmium capture, respectively, and 8.80 and 8.70 MeV of emitted gamma energy per capture. The `PhotonEvaporation` variant gives 6.66 photons per capture and 9.00 MeV per capture, while reducing the 9.043 MeV line-window yield from about 0.33 photons per cadmium capture to about 0.001. This result motivates the default physics convention used for the veto study: HP neutron transport is kept as the nominal treatment for neutron propagation and capture rates, strict isotope handling is enabled for production runs, and `PhotonEvaporation` is used as the nominal cadmium-capture cascade model when prompt-gamma energy partition is part of the observable. The default `ParticleHP` cascade is retained as a systematic comparison and for continuity with earlier productions.




Figure 5.8: Prompt-gamma emission from neutron capture in cadmium. The left panel overlays the evaluated  $^{113}\text{Cd}$  prompt-gamma lines from the IAEA PGAA database [105] with the gamma energies emitted by **Geant4** for cadmium captures in the HENSA three-layer veto simulation. The line yields are normalized using the natural-cadmium thermal capture cross-section from the same database, while the histogram is normalized per simulated cadmium capture. The right panel shows the simulated cascade multiplicity, with an average of 3.76 emitted photons per cadmium capture.

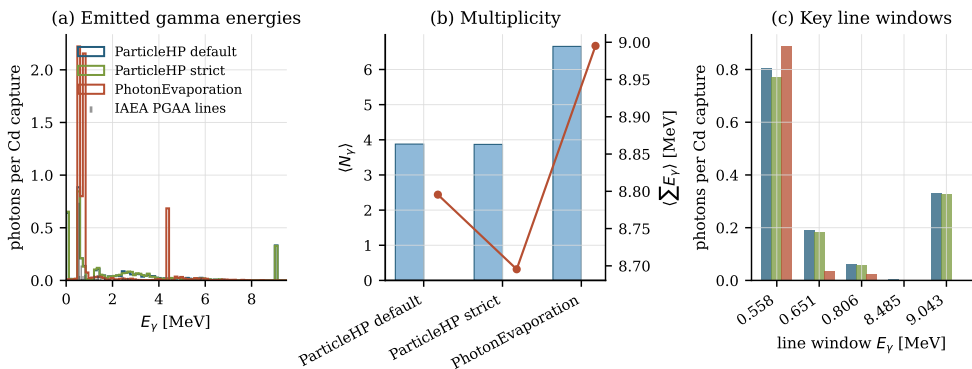


Figure 5.9: Model comparison for prompt-gamma emission from cadmium captures in the HENSA three-layer veto geometry. The default **ParticleHP** model and the strict-isotope variant have nearly identical multiplicities, energy sums, and line-window yields. The **PhotonEvaporation** variant preserves the total capture-energy scale but shifts the cascade toward more numerous, lower-energy photons and suppresses the high-energy 9.043 MeV line that is overproduced by the default model. The comparison uses the partial *r2* validation campaign; the conclusion is therefore used as a model-systematics diagnostic rather than as a high-precision spectroscopy result.

These cross-section arguments are used here only as physical motivation. The actual veto efficiency depends on geometry, shielding gaps, secondary-particle multiplicity,

capture position, capture time, gamma transport, scintillator quenching, light attenuation, PMT thresholds, and Micromegas topology. Those coupled effects are the reason why the design is evaluated with the full `Geant4/restG4` simulation chain rather than with cross-section estimates alone.

### 5.1.4 Interaction of high-energy neutrons and photons with the detector

Atmospheric neutrons span many orders of magnitude in energy, from thermal and environmental energies to the GeV scale. The source spectra discussed above show that the integrated flux is large at low energy, but this is not the same as saying that low-energy neutrons dominate the detector background. The relevant question for the veto design is which primary-neutron energies survive the shielding and produce activity in the Micromegas sensitive volume.




Figure 5.10: Geometry of the IAXO detector and shielding as implemented in the `Geant4` simulation. The copper chamber is shown in the center, surrounded by the 20 cm thick lead shield.

This was studied in two complementary samples. The first is the original CRY-based baseline lead-shielding simulation, used to establish the high-energy response of the detector-plus-shield system. The second uses the recent outdoor-HENSA neutron

source propagated through the three-layer cadmium veto configuration. The baseline configuration consists of a 20 cm thick lead shield surrounding the copper chamber, as shown in figure 5.10; a thin layer of copper is used as inner shielding.

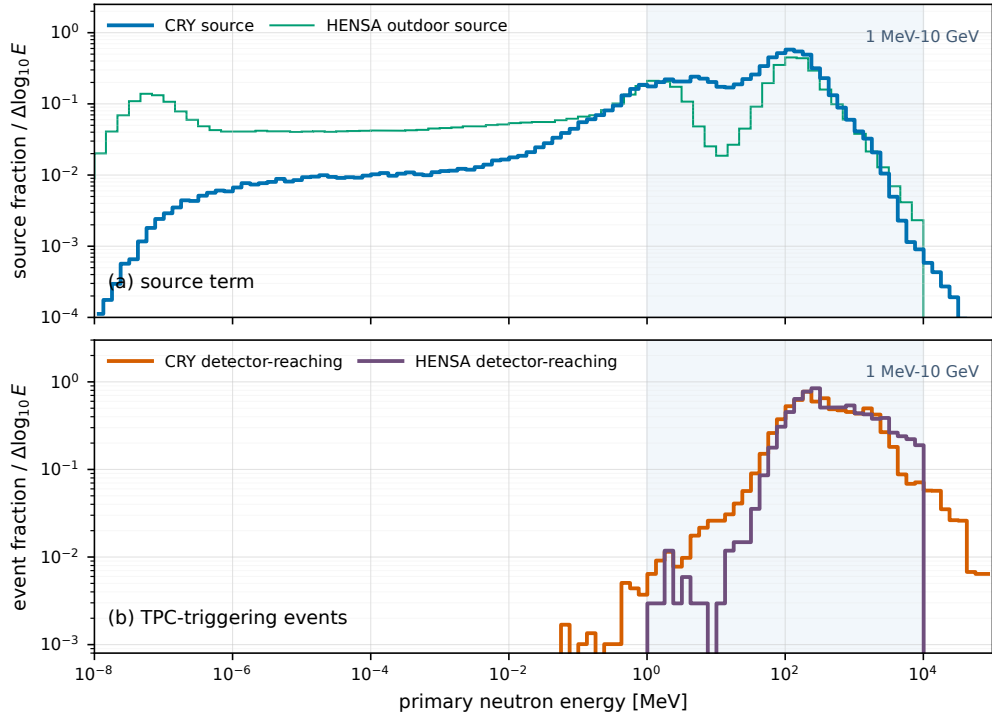


Figure 5.11: Primary-neutron energy distributions before and after detector selection. The upper panel compares the CRY neutron source with the outdoor-HENSA source term. The lower panel shows the primary energies of events that produce Micromegas-sensitive activity in the old CRY baseline study and in the recent HENSA-driven three-layer cadmium simulation. All curves are normalized to unit area to emphasize the change in spectral shape.

Energy range	CRY source	CRY detector	HENSA source	HENSA detector
10 meV–1 eV	0.56%	0.00%	14.83%	0.00%
1 eV–1 keV	2.77%	0.05%	12.53%	0.00%
1 keV–1 MeV	14.67%	0.29%	21.10%	0.00%
1 MeV–10 MeV	21.02%	1.37%	12.87%	0.41%
10 MeV–100 MeV	28.99%	12.87%	12.20%	8.13%
100 MeV–1 GeV	30.53%	57.20%	25.05%	59.68%
1 GeV–10 GeV	1.43%	25.37%	1.43%	31.78%
10 GeV–100 GeV	0.02%	2.77%	0.00%	0.00%
100 GeV–1 TeV	0.00%	0.09%	0.00%	0.00%
<b>1 MeV–10 GeV total</b>	<b>81.97%</b>	<b>96.80%</b>	<b>51.55%</b>	<b>100.00%</b>

Table 5.4: Energy fractions of the source and detector-reaching neutron distributions shown in Fig. 5.11. The detector columns are normalized to the events that produce Micromegas-sensitive activity, not to the total source flux. The HENSA detector column uses the three-layer cadmium simulation, whose generated source is restricted to the 1 MeV–10 GeV production interval.

Figure 5.11 and Table 5.4 show the same qualitative behavior for both source models. Sub-MeV neutrons are an important part of the measured or generated source term, especially for the HENSA outdoor spectrum, but they do not dominate the detector-triggering population. The events that reach the Micromegas sensitive volume are concentrated in the 100 MeV–10 GeV region, where neutrons are energetic enough to create secondary activity in the lead and copper surrounding the detector. This justifies keeping the high-energy part of the neutron source in the veto simulations even though its integral flux is small, and it also motivates the active veto: the problematic neutron background is not simply a low-energy moderation problem, but a secondary-shower problem produced by fast primaries interacting in the shielding.

For the original CRY baseline study, the ratio between the detector-reaching and input primary-energy distributions can be used as a relative detector-response diagnostic. This ratio should not be interpreted as an absolute detection efficiency: each bin compares the fraction of detector-reaching events with the fraction of generated primaries in that same energy interval, and the vertical normalization is arbitrary. It is nevertheless useful for identifying which primary energies are preferentially converted into Micromegas activity.




Figure 5.12: Relative detector-response diagnostics for neutron and photon primaries in the baseline shielding configuration. The response ratio compares the detector-reaching and generated primary-energy distributions and is normalized only for shape comparison. The shaded regions indicate the ranges used for the power-law guides, giving  $R(E) \propto E^{0.97}$  for neutrons and  $R(E) \propto E^{0.98}$  for photons.

Figure 5.12 shows that, above the MeV scale, the relative response grows approximately as

$$R(E) = CE^\alpha, \quad (5.2)$$

with  $\alpha$  close to unity for both neutron and photon primaries in the simulated range. This trend is physically reasonable: more energetic primaries create more secondary particles in the shielding and therefore have a larger probability of transferring energy into the Micromegas sensitive volume. The photon case was studied in the same baseline detector-plus-shielding configuration [97]; its similar response shape does not imply similar mitigation behavior, because the electromagnetic component is still much more efficiently attenuated by additional lead, as discussed in Section 5.2. For neutrons, the same lead that suppresses photons becomes the material where energetic primaries can multiply and generate the correlated secondary shower seen by the TPC and the veto system.

### Secondary-particle production in the lead shielding

The lead shield should therefore not be regarded as a passive absorber only. For high-energy neutrons it is also the main site of secondary-particle production. Dedicated simulations with the HENSA outdoor neutron spectrum and the nominal three-layer veto geometry confirm that the population emerging from the shield is dominated by neutrons with significantly lower energies than the primaries, with a spectrum that peaks around the MeV scale. In the current sample, 1,745 detector-reaching events contain  $9.19 \times 10^4$  neutrons produced in the lead, of which  $6.58 \times 10^4$  leave the lead volume and  $1.83 \times 10^4$  do so with a momentum direction pointing toward the TPC. This shift toward lower energies is important because it makes the secondary shower much better matched to moderation and detection in the hydrogen-rich scintillator panels of the veto system.




Figure 5.13: Energy distributions of secondary neutrons produced in the lead shielding and of the subset that leave the shield in the HENSA outdoor three-layer-veto simulation. The spectra are shown as  $E dN/dE$  per detector-reaching event. The lower panel gives the fraction of exiting neutrons whose momentum points toward the TPC. The shielded configuration softens the outgoing neutron population toward the MeV range, while the directionality study shows that the secondary shower is distributed over the full detector volume rather than being confined to the TPC side alone.

This behavior clarifies why the veto can be effective even though the primary cosmic neutrons themselves are difficult to detect directly. The relevant observable is not the

high-energy primary neutron in isolation, but the correlated shower of secondaries generated when it interacts in the lead.




Figure 5.14: Neutron-history classification of the Micromegas signal origin for the HENSA outdoor three-layer-veto production sample. The darker bars show the fraction of neutron-induced TPC events in each category, while the lighter bars show the corresponding fraction of the total gas energy deposited in the same sample. Direct neutron deposits in the gas are only a small contribution; most events are produced through neutron-induced photons, electrons, recoils, fragments, or other secondaries generated in the surrounding detector and shielding materials.

Figure 5.14 shows that the final interaction responsible for the TPC signal is very often produced by a secondary particle created in the surrounding shielding and detector materials. The event count is dominated by gamma/electron routes, whereas the deposited gas energy is dominated by recoil and fragment routes. This result provides the physical justification for designing the veto system to tag the shower as a whole, rather than attempting to identify only the incoming primary neutron. It also motivates the later use of veto observables based on prompt energy, delayed activity, multiplicity, and spatial correlations, since these quantities are sensitive to the topology of the full secondary cascade.

### 5.1.5 Effect of detector inclination during solar tracking

As a solar axion helioscope, BabyIAXO must track the Sun across the sky during its daily observation window. This tracking is achieved by physically rotating and tilting the entire magnet and detector assembly, so that the optical axis remains pointed toward the Sun. Consequently, the detector does not operate in a fixed orientation:

the inclination angle with respect to the horizontal plane changes continuously during each tracking period.

This raises a natural question for the background model: does the cosmic ray-induced background depend significantly on the detector orientation? For the muon component, some dependence might be expected on physical grounds, since the muon flux has a strong zenith-angle dependence and the effective cross-sectional area of the TPC presented to vertically incident muons varies with inclination. For the neutron component, the situation is less clear, as the secondary neutron shower produced inside the lead shielding is largely isotropic.

To address this question, dedicated `Geant4` simulations were performed for the horizontal detector orientation and for tracking inclinations spanning the  $-25^\circ$  to  $+25^\circ$  range. The full three-layer veto geometry was retained, while the incident cosmic field was rotated with respect to the detector axes. This is equivalent to tilting the detector with respect to the downward cosmic-ray field, but avoids modifying the GDML geometry for each angular point. The scan was performed for both cosmic muons and HENSA-driven neutrons, and the convention and response are summarized together in Fig. 5.15.

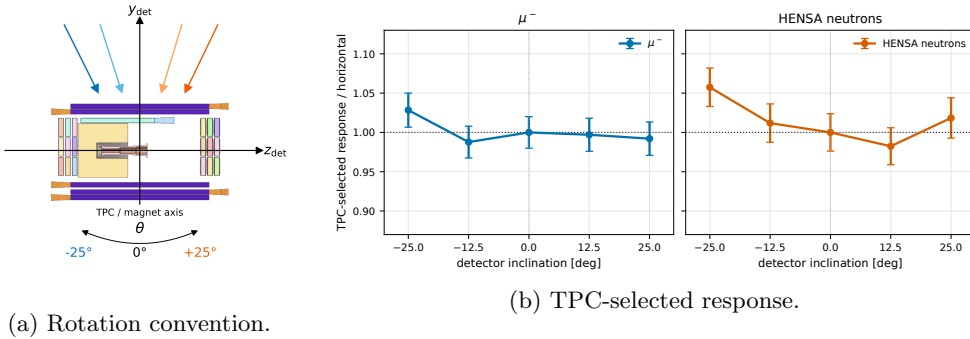


Figure 5.15: Detector-inclination study for the three-layer veto geometry. The detector, shielding, and veto volumes are kept fixed, while the incident cosmic source direction is rotated in the detector  $y_{det} - z_{det}$  plane. The response is normalized to the horizontal configuration for each source.

The TPC-selected rates in Fig. 5.15b show no monotonic or operationally significant dependence across the tracking-angle range. For HENSA neutrons, the relative response spans from  $0.982 \pm 0.024$  at  $+12.5^\circ$  to  $1.057 \pm 0.024$  at  $-25^\circ$ , compared with the horizontal configuration. For muons, the corresponding range is  $0.988 \pm 0.020$  to  $1.028 \pm 0.022$ . The excursions remain at the few-percent level, so detector inclination is not a leading effect for the cosmic-background studies presented in this chapter. The result justifies using a fixed horizontal geometry for the larger production campaigns, while treating

residual orientation effects as a subdominant systematic.

## 5.2 Limits of passive shielding

The passive shield defines the boundary conditions for the veto design. Lead is retained because it is highly effective against environmental photons in the keV–MeV range, while copper provides a radiopure inner layer close to the detector. At the same time, the detector line cannot be enclosed by a perfect passive castle: the X-ray entrance path, gas services, readout cabling, and mechanical supports introduce openings and penetrations. For a surface helioscope, passive shielding must therefore be evaluated not only as an absorber, but also as part of the geometry in which cosmic-ray-induced secondary particles are produced.

Neutron shielding is qualitatively different from gamma shielding. Hydrogen-rich materials such as polyethylene or water are useful moderators because neutron elastic scattering on hydrogen can transfer a large fraction of the neutron energy, as discussed in Section 5.1.3. Boron- or lithium-loaded materials can then absorb moderated neutrons through neutron-capture or absorption reactions. Such captures may also generate gammas or charged particles, so the absorber must be placed in a geometry where those secondaries do not become a new detector background. These constraints make passive neutron shielding a design trade-off rather than a simple matter of adding more material.

In the IAXO-D0/BabyIAXO geometry, the passive shield must also remain compatible with the X-ray optical axis. A copper pipe connects the detector to the vacuum beam line and produces the most important intentional weakening of the shielding. The lead-thickness and moderator studies presented below therefore ask a narrow design question: can practical passive shielding suppress the remaining cosmic-ray-induced background, or must the detector rely on an active veto?

### 5.2.1 Lead-thickness scan: electromagnetic suppression versus neutron multiplication

A parameterized lead-shield geometry was used to scan the passive-shielding thickness with two variants: an idealized  $4\pi$  coverage and a more realistic configuration including the beam-pipe opening. The atmospheric-secondary distributions generated with CRY (Figure 5.1) were used as input. The scan was designed to motivate the original shielding choice and, just as importantly, to test whether lead alone could be a complete solution for a surface detector. The full per-particle scans are preserved in Appendix 6.13; the main text keeps only the photon and neutron trends because they drive the design decision.




Figure 5.16: Design summary of the lead-thickness scan. Increasing the lead thickness strongly suppresses the photon-induced background, especially for ideal  $4\pi$  coverage, which motivates the original lead-shielding concept. The neutron-induced background does not follow the same attenuation trend: it is weakly affected by additional lead and shows a broad non-monotonic structure associated with secondary production in the high- $Z$  shield. The scan therefore gives the central passive-shielding message of this chapter: lead is necessary for the electromagnetic component, but it is not sufficient against cosmic neutrons.

The photon response confirms the expected role of the lead castle. In the idealized  $4\pi$  configuration the photon-induced background drops by more than four orders of magnitude over the scanned range; with the beam-pipe opening the suppression is weaker, about a factor  $3 \times 10^2$ , because leakage through the opening becomes comparatively important once the rest of the solid angle is shielded. Lead therefore remains part of the baseline design, but the geometry of penetrations matters.

The neutron response leads to the opposite conclusion. High-energy neutrons interact in the lead and can produce secondary neutron showers, so the residual detector background is not removed simply by increasing the lead thickness. In this scan the neutron level remains within the same order of magnitude across the full range

and even rises at intermediate thickness before decreasing mildly at large thickness. The beam-pipe configuration changes the neutron rate less dramatically than it changes the photon rate and does not alter the design conclusion. The shield is therefore both a necessary gamma attenuator and a source of correlated neutron-induced activity that must be tagged by the active system.

Component	Passive-shielding behavior	Design consequence
<b>Muons</b>	<i>Pb trend:</i> approximately flat within the scan. <i>Pipe opening:</i> minor compared with the intrinsic penetration of muons.	Lead does not solve this component; prompt active vetoing is required.
<b>Photons</b>	<i>Pb trend:</i> strong suppression as thickness increases. <i>Pipe opening:</i> important at large Pb thickness, when leakage through openings becomes comparatively visible.	Retain thick Pb and minimize penetrations where compatible with the X-ray line.
<b>Electrons and protons</b>	<i>Pb trend:</i> reduced by material and topology, but not the design-limiting passive-shielding case. <i>Pipe opening:</i> geometry-dependent and secondary-production dependent.	Keep as simulation checks; treat residual proton-like cascades with the same active-veto logic as neutrons.
<b>Neutrons</b>	<i>Pb trend:</i> weak and non-monotonic improvement because Pb can multiply high-energy neutrons. <i>Pipe opening:</i> less decisive than for photons and does not change the conclusion.	Passive Pb shielding is insufficient; a neutron-sensitive veto is required.
<b>Borated HDPE</b>	<i>Pb trend:</i> modest improvement for practical layered configurations. <i>Placement:</i> depends on where moderation and capture occur relative to Pb.	Useful as a design study, but insufficient as a stand-alone mitigation.

Table 5.5: Design interpretation of the passive-shielding studies. The table summarizes the trends used in the chapter; the detailed per-particle scan figures are collected in Appendix 6.13.

### 5.2.2 Borated-HDPE moderation studies

The lead-thickness scan motivates the consideration of passive materials specifically aimed at neutron moderation and absorption. High-density polyethylene (HDPE) is a standard hydrogen-rich moderator, and borated HDPE adds neutron absorption through the  $^{10}\text{B}$  reaction



after the neutron has been slowed in the polymer matrix. This makes borated HDPE an attractive passive candidate in principle: hydrogen can reduce the neutron energy, while boron can absorb part of the thermalized population. The question for BabyIAXO is whether this mechanism remains effective for the relevant surface-neutron problem, where many of the problematic events are produced by MeV–GeV primaries interacting in and around the lead shield.

A dedicated simulation study tested whether 5% borated HDPE could be used as an additional passive layer without abandoning the 20 cm lead shield required for gamma suppression [97]. A single external borated-HDPE layer produced only a small reduction in the neutron-induced background. A layered Pb/borated-HDPE/Pb configuration was therefore tested, keeping the total Pb thickness fixed at 20 cm while scanning both the borated-HDPE thickness and the partition of lead between the inner and outer sides of the moderator. The geometry is shown schematically in Figure 5.17. In the scan notation used below,  $L_1$  is the inner lead thickness,  $L_2$  is the borated-HDPE thickness, and the remaining outer lead thickness is  $L_3 = 200 \text{ mm} - L_1$ .




Figure 5.17: Layered passive-shielding configuration investigated for neutron moderation [97]. The total lead thickness is kept fixed while an intermediate borated-HDPE layer is scanned, so the study tests whether moderation can be added without abandoning the baseline gamma-shielding role of lead.




Figure 5.18: Neutron-induced background level predicted for layered Pb/borated-HDPE/Pb configurations with fixed 20 cm total lead thickness [97]. Panel (a) shows the absolute background level, while panel (b) shows the residual background relative to the corresponding no-moderator configuration,  $B_0 = B(L_2 = 0)$ , for the same lead partition. The scan demonstrates that borated HDPE can reduce the neutron-induced rate, but only by a modest factor and with a strong dependence on the relative placement of the lead and moderator.

The best sampled configuration reduces the neutron-induced background by about 40% with respect to the plain 20 cm Pb reference, corresponding to a suppression factor of only about 1.6; the most favorable lead partitions give reductions of order 30–40%. The reduction is largest when a substantial fraction of the lead remains outside the moderator, so that the moderator acts on neutrons produced upstream in the shield. For configurations with most of the lead inside the moderator, the gain is much smaller, and thin moderator layers can even increase the residual rate within the statistical precision of the scan. This behavior is consistent with the energy scale of the source term: a practical moderator thickness does not fully thermalize the MeV–GeV neutron component before it has already produced secondaries in the lead or surrounding materials.

Borated scintillator layers, acting simultaneously as active and passive shields, were also explored and gave only modest additional improvement for the highest-energy neutron component. Borated HDPE was therefore not adopted as a baseline passive layer. The study nevertheless defines the boundary condition for the active-veto design: lead is kept because it efficiently suppresses photons, but neither more lead nor practical borated-HDPE moderation reduces the high-energy neutron background to the required level. The active shield must instead detect the correlated secondary shower produced in and around the passive shield.

## 5.3 Neutron-sensitive active veto concept

This section follows the evolution of the veto concept from a conventional surface muon shield to a neutron-sensitive active system coupled to the passive lead shielding. The emphasis is not only on the final geometry, but on the sequence of questions that connected the source-term studies, the passive-shielding limits, the neutron-tagging mechanism, the choice of three scintillator layers with cadmium, and the subsequent validation with prototype data.

Step	Question and study	Design consequence
<b>Initial surface-background studies</b>	Which cosmic components matter? CRY plus literature-based atmospheric-secondary simulations.	Muons dominate the raw rate; neutrons remain problematic after shielding and Micromegas cuts.
<b>Passive lead shielding</b>	Can thicker lead solve the problem? Lead-thickness scan with and without the beam-pipe opening.	Gamma backgrounds improve strongly; neutron-induced background does not.
<b>Passive neutron shielding</b>	Can HDPE or borated HDPE solve it? Layered Pb/borated-HDPE/Pb simulations.	Passive moderation helps, but remains insufficient on its own.
<b>Active-material ordering</b>	Where should the active scintillator sit? Simplified neutron sandwich scans with Pb, B-HDPE, and capture/moderator sheets.	Couple the scintillator directly to the lead-generated secondary shower; thin sheets alone do not replace a multi-stage veto.
<b>Neutron-interaction studies</b>	How can high-energy neutrons be tagged? Geant4 studies of secondary production in lead and transport to the veto.	Tag the secondary shower rather than the primary neutron itself.
<b>Cadmium-capture concept</b>	How can delayed neutron activity be recovered? Cadmium vs. no-cadmium simulations and capture-time studies.	Add cadmium sheets between scintillator layers to exploit delayed capture gammas.
<b>Layer optimization</b>	How many layers are needed? 1-, 2-, 3-, and 4-layer comparisons.	Three layers provide the best efficiency/complexity compromise.
<b>Readout simulation</b>	Are deposited-energy cuts realistic? Waveform-level simulation with attenuation, shaping, and peak finding.	Use timing, multiplicity, and peak observables rather than energy-in-volume alone.
<b>Prototype validation</b>	Does the system work in data? IAXO-D0 veto commissioning and background-discrimination analysis.	Timing- and multiplicity-based observables improve rejection beyond a simple muon veto.

Table 5.6: Design history of the veto system, summarizing how each study informed the final concept.

Detailed campaign metadata, including the response level and timing convention used in each scan, are collected in Appendix 6.8. The main text keeps only the design consequences, so that the reader can follow how the veto evolved from a prompt muon counter into a neutron-sensitive, waveform-level system.

### 5.3.1 Plastic scintillators

Plastic scintillators were a natural choice for the active shielding because they combine large-area coverage, good timing performance, and a high hydrogen content, which is advantageous for neutron moderation. In the BabyIAXO/IAXO-D0 case, the collaboration also had access to a substantial stock of NE-110 panels and

photomultipliers from a former time-of-flight spectrometer [108]. The veto design was therefore developed around this existing hardware, with segmentation and panel lengths constrained by the available modules. Several material names appear in the simulation studies: NE-110 for the reused prototype bars, EJ-208 as the modern datasheet reference used for optical properties, and BC408 as the PVT-like scintillator used in the simplified material scans. For the neutron-transport argument these are treated as closely related polyvinyltoluene-like organic scintillators; the exact optical constants matter for the light-collection model, while the moderation and recoil-proton response are driven mainly by the hydrogen and carbon content [109, 110].

<b>Property</b>	<b>Value</b>	<b>Unit</b>
Density	1.023	g/cm <sup>3</sup>
Scintillation Efficiency	9200	photons / MeV
H Atoms	$5.17 \times 10^{22}$	atoms/cm <sup>3</sup>
C Atoms	$4.69 \times 10^{22}$	atoms/cm <sup>3</sup>
H:C Ratio	$\sim 1.1 : 1$	–
Wavelength of Max. Emission	435	nm
Light Attenuation Length	400	cm
Pulse Width (FWHM)	4.2	ns

Table 5.7: Physical properties of EJ-208 plastic scintillator, used here as the modern datasheet reference for the NE-110 material employed in the veto panels [109]. The simplified neutron-material scans use a BC408-equivalent PVT composition, whose H:C ratio and density are sufficiently close for the moderation argument.




Figure 5.19: Cutaway view of the IAXO-D0/BabyIAXO prototype setup highlighting the segmented veto system surrounding the lead shield and the photomultiplier-coupled scintillator panels [111]. Although the operational system is multi-layered, the figure illustrates the basic module concept used throughout the veto design: each scintillator panel is read out by a dedicated PMT through a light guide.

A scintillator converts the kinetic energy of ionizing radiation into detectable light, which is subsequently collected and amplified by a photomultiplier tube. The fundamental figure of merit is the scintillation efficiency, defined as the ratio of photons produced to the energy deposited.

For standard calibration purposes, plastic scintillators are exposed to gamma sources (e.g.,  $^{60}\text{Co}$  or  $^{137}\text{Cs}$ ), which induce energetic electron recoils via Compton scattering. Since the light yield for electrons is approximately linear with energy, this establishes an "electron-equivalent" energy scale ( $E_{ee}$ ), typically measured in units of MeVee (MeV electron equivalent).

However, the scintillation efficiency is strictly dependent on the particle type and its specific energy loss ( $dE/dx$ ). This is particularly relevant for the detection of neutrons, which is a primary focus of this work. Neutrons interact within the plastic scintillator primarily by elastic scattering off hydrogen nuclei, producing recoil protons. Unlike the electrons from gamma interactions, these recoil protons are heavily ionizing. The resulting high density of ionization centers along the proton track leads to a saturation of the fluorescent molecules, causing a significant reduction in light output known as quenching.

In organic scintillators such as NE-110, this non-linear response is semi-empirically

described by Birks' law [112, 113]. Consequently, a neutron-induced recoil proton depositing 10 MeV of kinetic energy will produce significantly less light than an electron depositing the same amount. This necessitates the application of a quenching correction factor in simulation studies to accurately predict the veto efficiency for neutron backgrounds.

It is worth noting that while Birks' law governs the quenching in the solid scintillator, the visible-signal reduction for nuclear recoils in the gaseous Micromegas volume is driven by a different mechanism and is treated through Lindhard-type models [114]. During this work both effects were incorporated in the simulation chain, since a realistic veto/background comparison requires the scintillator and the gas detector to be treated consistently at the level of visible signal rather than only at the level of deposited energy.

### 5.3.2 From prompt muon veto to neutron-sensitive stages

Given the IAXO background requirements and the limits of passive shielding, the key remaining challenge is the high-energy cosmic-neutron component [48]. The early active-shielding concept inherited from previous detector configurations was essentially a single-layer muon veto: effective for prompt charged tracks, but not yet optimized for the softer, more distributed, and often delayed signatures expected from neutron-induced secondary showers.




Figure 5.20: Early veto concept based on the previous IAXO-D0 background model [48]. A single scintillator layer is adequate for prompt muon tagging, but offers too little redundancy for delayed and low-amplitude neutron-related activity.

The next design question is where the active material should sit relative to the lead and moderator volumes. Rather than retaining the older deposited-energy limiting

scans as a separate result, the comparison was repeated with the same Birks-corrected visible-energy response used in the current simulation chain. The simplified geometry reduces the veto to a 10 cm scintillator slab, treated as the only sensitive volume, and places lead, borated HDPE, and thin capture/moderator sheets before and after it along the neutron direction. The purpose is not to predict the final veto performance, but to isolate whether the active layer benefits mainly from direct neutron interactions, from secondaries produced in lead, or from local capture/moderator sheets before the segmented multi-layer waveform response is introduced.




Figure 5.21: Representative simplified sandwich comparison for high-energy neutrons incident on a lead-coupled scintillator slab. The curves compare a BC408-equivalent PVT scintillator, the same scintillator with 5 mm cadmium sheets, and a 5% boron-loaded EJ-254 option in the same Pb/scintillator/Pb ordering. Panel (a) shows the Birks-quenched visible-threshold response; panel (b) shows the primary-energy dependence at 500 keV<sub>ee</sub>. The full ordering scan and the quenching diagnostic are retained in Appendix 6.9.

Together these scans give the design lesson that replaces the older deposited-energy limiting studies. The scintillator should not be treated as an isolated neutron counter: its useful response comes from sampling the secondary shower generated by the surrounding shielding. The fixed-threshold energy projection shows the same point from the source-term perspective: the bare slab is most effective for low-energy neutrons that can interact directly in the scintillator, whereas the lead-coupled geometries are the ones that remain efficient for the harder primaries relevant to the surface cosmic-neutron spectrum. This is the essential reason for placing the veto around, and not instead of, the lead shielding.

## Neutron-tagging mechanism

Additional simulation studies presented during the IAXO collaboration meetings clarified the relevant energy and time scales of this problem [97, 115, 116]. Although the input cosmic-neutron spectrum extends from sub-MeV energies to the multi-GeV range, the detector-triggering events are dominated by primary neutrons in the MeV-to-GeV region, with the effective primary-energy distribution peaking around the GeV scale for the detector-plus-shielding system. After interacting in the lead shielding, however, the shower that emerges toward the veto is strongly softened: the outgoing secondary neutrons are concentrated at much lower energies, typically around the MeV scale. This separation of scales is central to the veto concept.

The relevant microscopic processes were introduced in Section 5.1.3. Here the same physics can be summarized operationally. Lead is a poor moderator but an efficient converter for fast neutrons: inelastic scattering,  $(n, xn)$  channels, and high-energy nuclear reactions create a softer secondary shower distributed through the shielding volume. The hydrogen-rich scintillator then samples the moderated and charged part of that shower through recoil-proton and charged-secondary light, while cadmium samples the thermalized tail through capture-gamma cascades. The veto is therefore designed to observe the history of the neutron-induced cascade, not the incoming neutron itself.

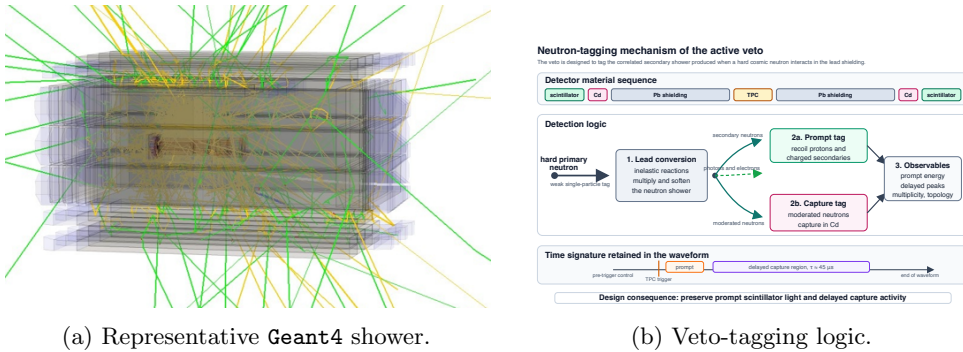


Figure 5.22: Neutron-induced shower topology and active-veto response. In the event display, a high-energy primary neutron undergoes non-elastic reactions in the lead shielding and generates secondary neutrons and photons, shown in yellow and green. The design schematic summarizes the corresponding veto logic: scintillators tag recoil-proton and charged-secondary light, while cadmium sheets tag the moderated component through delayed capture-gamma cascades.

This mechanism is the central reason for using a multilayer active veto rather than a single passive moderator or a single scintillator counter. The lead shielding increases the probability that a high-energy neutron creates detectable secondaries, while the hydrogen-rich scintillator and cadmium sheets sample different stages of the moderation

history. A neutron-induced background event need not look like a clean through-going track. It can instead appear as a moderate prompt signal, several delayed peaks, and a spatially distributed multiplicity pattern. The analysis is therefore designed to retain these correlated waveform-level observables.

The timing structure of the event is as important as its topology. The interaction that gives rise to the Micromegas trigger is essentially prompt on the timescale of the detector: in the dedicated simulations, the relevant activity in the shielding and the first interaction in the sensitive volume occur within less than about  $1\ \mu\text{s}$  from the primary neutron generation [97, 115]. By contrast, neutron capture is delayed by the stochastic moderation chain. The capture-time distribution follows approximately an exponential law with characteristic constant  $\tau \simeq 45\ \mu\text{s}$ , measured relative to the Micromegas trigger time [115, 116]. An ideal post-trigger delayed window of  $100\ \mu\text{s}$  would therefore contain approximately  $1 - \exp(-100/45) \approx 89\%$  of delayed captures before threshold and acceptance effects are considered. In the prototype-like acquisition mode used later in this chapter, however, the veto is read out in a  $100\ \mu\text{s}$  total window with the Micromegas trigger placed at  $30\ \mu\text{s}$ , corresponding to approximately  $[-30, +70]\ \mu\text{s}$  relative to the trigger. The practically accessible post-trigger delayed fraction is then  $1 - \exp(-70/45) \approx 79\%$ . This is the reason why the veto analysis is not restricted to a narrow prompt coincidence: doing so would discard a large fraction of the cadmium-capture information that motivated the multilayer design.

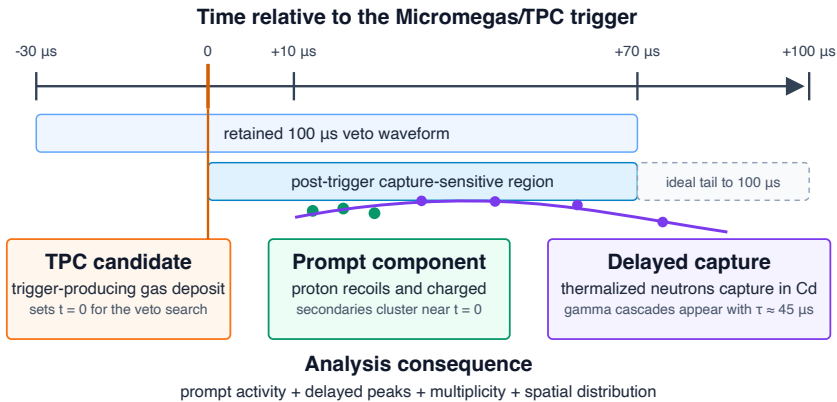


Figure 5.23: Characteristic timing structure of neutron-induced events in the veto system. The Micromegas trigger defines  $t = 0$ , prompt veto peaks are expected within a few microseconds, and capture-related activity can extend over several tens of microseconds. The prototype-like waveform window retains about  $70\ \mu\text{s}$  of post-trigger time, which contains approximately 79% of an ideal exponential capture tail with  $\tau \simeq 45\ \mu\text{s}$ .

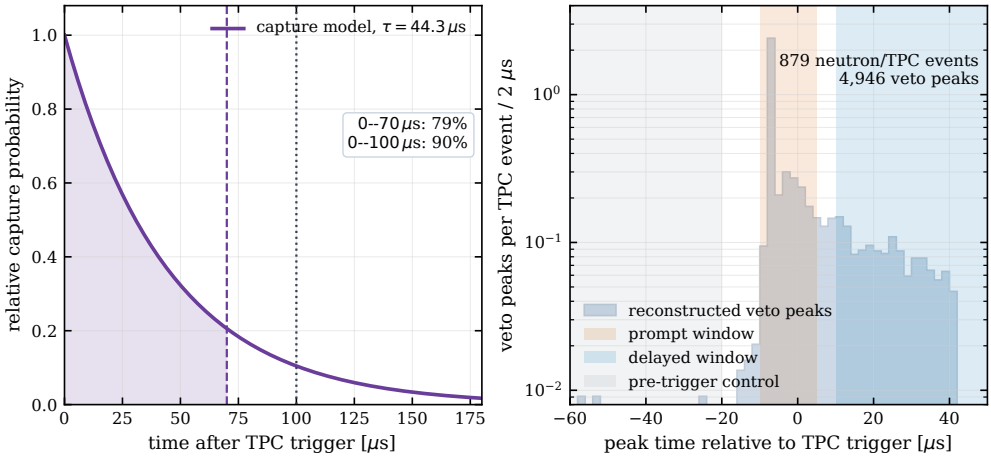


Figure 5.24: Timing diagnostics for the neutron-capture veto strategy. Left: exponential capture-time model with  $\tau = 44.28 \mu\text{s}$ , showing the fraction of an ideal delayed tail retained by  $70 \mu\text{s}$  and  $100 \mu\text{s}$  post-trigger windows [115]. Right: reconstructed veto-peak times in the waveform-level cosmic-neutron sample used for the neutron-pattern study. The reconstructed distribution contains both prompt activity and a late component, but the late peaks are detector-level observables and should not be interpreted as one-to-one truth labels for individual captures.

The same mechanism can be inspected directly in the recent HENSA-driven three-layer cadmium simulation. The diagnostic summarized in Appendix 6.11, Fig. 6.51, scans the `Geant4` history of the 2,706 neutron events that survive the common TPC selection and records explicit neutron-capture processes together with the reconstructed veto peaks. In total, 73,295 neutron-capture steps are found. At the event level, 99.3% of the selected events contain at least one neutron capture, and 98.3% contain at least one capture in a cadmium sheet. Cadmium is also the dominant capture site in the full population of capture processes: 78.9% of all explicit captures occur in the cadmium sheets, while 19.3% occur in the plastic scintillator. Within the cadmium sheets, the capture isotopes are dominated by  $^{113}\text{Cd}$ , which accounts for about 98% of the recorded cadmium captures in this sample.

For events whose first TPC activity is prompt, 79.7% of cadmium captures occur within the first  $70 \mu\text{s}$  after the TPC time, and 88.8% within  $100 \mu\text{s}$ . This reproduces at truth level the same time-window scale inferred from the exponential capture-time model. The bottom-right panel also shows why capture truth cannot be used as a direct analysis observable: about one quarter of the selected events contain a cadmium capture but no reconstructed veto peak, while another quarter produce only prompt or near-trigger peaks. The useful experimental signature is therefore not the capture process itself, but the subset of captures and secondary shower activity that survives gamma transport, scintillator response, thresholds, and waveform reconstruction.

The delayed component is therefore physically central, but it is not sufficient by itself as an analysis cut. Late veto peaks also occur in calibration and mixed-background samples through accidental activity, electronic noise, and unrelated environmental coincidences. This can be seen in a simple event-mixing test in which veto peak trains are time-scrambled with respect to the TPC trigger. A requirement of at least one late peak accepts about 52% of the simulated neutron/TPC events but also about 56% of the mixed accidental sample. Adding a prompt requirement reduces the accidental acceptance to about 14% while retaining about 46% of neutron events; requiring the late activity to span at least two veto groups gives a cleaner, more topology-aware tag with about 25% neutron acceptance and about 7% accidental acceptance. The practical capture strategy is therefore to combine prompt activity, delayed peaks, multiplicity, visible energy, and spatial distribution rather than to search for an isolated delayed pulse.




Figure 5.25: Illustrative capture-like veto selections evaluated against an event-mixed accidental model. A late peak alone is not neutron-specific because the accidental acceptance is comparable to the neutron acceptance. The useful separation appears when delayed activity is combined with prompt activity and segmentation, which is the analysis-level expression of the scintillator–cadmium design.

This is the principal design strategy followed in the remainder of the section: repeated scintillator–cadmium stages are used to increase the probability of recording at least one prompt or delayed signature from the neutron-induced shower, while the event selection uses the waveform-level context to reject accidental delayed activity.

A dedicated neutron-history check on the waveform-level cosmic-neutron sample confirms that the veto must be interpreted as a tag of the full secondary cascade, not only of the incoming neutron. The study followed the `Geant4` ancestry of the particle that deposited energy in the Micromegas gas for 879 simulated neutron events with

nonzero TPC energy. The classification is a truth-level diagnostic used to understand the detector response; it is not used as a selection variable in the final analysis. Most TPC deposits are produced by electromagnetic descendants, which explains why a neutron-initiated event can pass an x-ray-like Micromegas selection while still carrying a neutron-like prompt or delayed veto history.

Neutron-history class	Events	Fraction	Detector interpretation
Neutron-induced gamma/electron	552	62.8%	A shielding interaction produces electromagnetic secondaries; a gamma or electron later deposits energy in the gas. This is the most x-ray-like class.
Gas recoil or fragment	242	27.5%	A recoil, fragment, proton, or other heavy charged secondary deposits energy in the gas, usually with broader topology or larger energy.
Direct neutron in gas	44	5.0%	A neutron reaches the gas and scatters there directly; this class is small but can have weak veto activity.
Charged meson secondary	41	4.7%	A rare high-energy cascade produces a charged meson or related secondary that reaches the gas, typically with strong veto activity.

Table 5.8: Truth-level history categories for simulated cosmic-neutron events that deposit energy in the Micromegas gas. The table motivates the veto strategy: the scintillator–cadmium system tags the cascade history surrounding an otherwise x-ray-like TPC deposit.

## 5.4 Multilayer veto optimization

The geometry study compared 1-layer (19 panels), 2-layer (39 panels), 3-layer (59 panels), and 4-layer (79 panels) configurations. The first optimization was performed with an idealized deposited-energy observable in the scintillator volumes [115]. That scan was useful for design intuition, but the final comparison was repeated with the outdoor HENSA neutron source term and the detector-response chain used in the background-model simulations: multilayer `Geant4` transport, quenching, light attenuation, waveform generation, and reconstructed veto observables. The threshold in this final scan is therefore a calibrated visible-energy proxy, not the same quantity as the older truth-level deposited-energy threshold.

The response-chain comparison preserves the original design conclusion but gives more conservative absolute rejection values. For example, the earlier idealized scan gave approximate 3-layer rejection values of 0.85 and 0.52 at 4 and 8 MeV, whereas the final HENSA response-chain scan gives 0.58 and 0.42 at the same nominal thresholds. For the 4-layer configuration the corresponding values change from about 0.88 and 0.54

to 0.63 and 0.48. This difference is expected: quenching, attenuation, waveform-level thresholds, and reconstruction losses all move the analysis away from a perfect deposited-energy observable. The relevant stability check is the ordering, which is unchanged. The second and third scintillator–cadmium stages give the large improvement, while the fourth stage adds only a smaller residual gain.

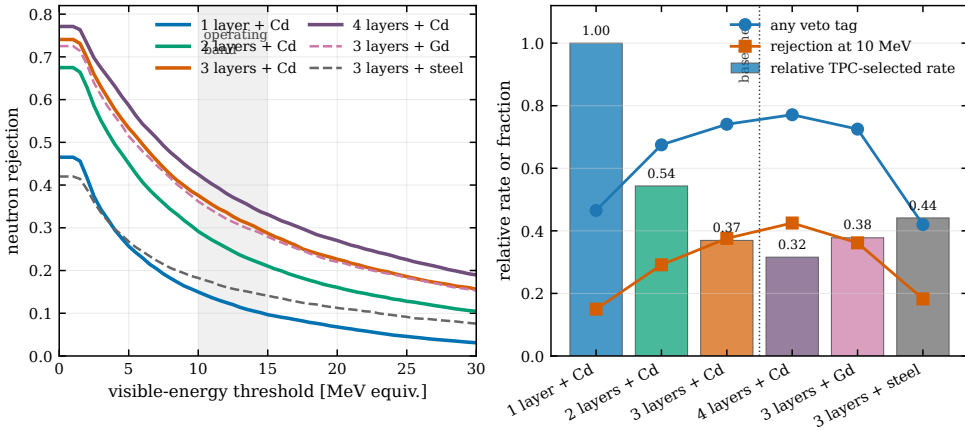


Figure 5.26: Final HENSA outdoor neutron layer-scan summary after the common TPC selection and detector-response processing. Panel (a) shows neutron rejection as a function of reconstructed veto visible-energy threshold; the shaded band marks the approximate 10–15 MeV operating region set by calibration acceptance. Panel (b) compares the relative TPC-selected event rate, any-veto-tag fraction, and 10 MeV rejection for the layer and capture-material variants. The first three cadmium-assisted stages provide the dominant gain, while the fourth layer gives a smaller residual improvement with more channels and accidental-veto exposure.

Relative to the one-layer cadmium configuration, the event rate after the common TPC selection is reduced to about 0.54, 0.37, and 0.32 for the two-, three-, and four-layer cadmium configurations, respectively. At the same time, the fraction of events with a reconstructed veto tag increases from about 47% in the one-layer geometry to 74% in the three-layer geometry and 77% in the four-layer geometry. This confirms the design interpretation: three layers do not maximize rejection in an unconstrained sense, but they capture most of the available veto information before the mechanical complexity, readout channel count, and accidental-coincidence exposure increase further. The capture-material comparison reinforces the same interpretation. A three-layer gadolinium variant gives a similar neutron-tagging response to cadmium in this observable set, whereas replacing the capture sheets with stainless steel strongly reduces the reconstructed veto-tag fraction and the 10 MeV rejection.

Configuration	Panels	Events	Rel. rate	Veto tag	Rej. 10 MeV
1 layer + Cd	19	19755	1.000	0.465	0.150
2 layers + Cd	39	10736	0.543	0.675	0.292
3 layers + Cd	59	7300	0.370	0.741	0.376
4 layers + Cd	79	6240	0.316	0.771	0.425
3 layers + Gd	59	5153	0.378	0.725	0.362
3 layers + steel	59	6019	0.441	0.420	0.183

Table 5.9: Final HENSA outdoor neutron layer-scan summary after the common TPC selection and detector-response processing. The cadmium configurations use 300 processed ROOT files each, corresponding to 840 job-hours per geometry; the gadolinium and stainless-steel capture-material variants use 230 processed files each, corresponding to 580 job-hours. The relative rate is normalized to the one-layer cadmium configuration using the simulated event rate per job-hour. The 10 MeV rejection column uses the reconstructed veto visible-energy proxy without adding accidental-noise overlays.

The detailed threshold-only and geometry-only plots from this scan are kept in Appendix 6.10, where they serve as provenance for the compact design summary in Fig. 5.26.

Layer optimization cannot be based only on the neutron rejection curve, because each additional instrumented layer also increases the number of channels that can contribute accidental veto activity. This effect was estimated by using experimental calibration-triggered data as an empirical accidental-veto template and scaling the amount of independent accidental activity with the number of active layers. For the measured veto-noise environment, a threshold of about 10 MeV in calibrated visible-energy proxy gives approximately 90% calibration acceptance for the three-layer configuration, while a more conservative threshold around 14–15 MeV gives approximately 95% acceptance. When this accidental component is included, the optimization becomes a balance between neutron rejection and calibration acceptance rather than a search for the lowest possible threshold.




Figure 5.27: Noise-aware veto-layer trade-off at a fixed 10 MeV calibrated visible-energy-proxy threshold. The accidental component is sampled from experimental calibration-triggered veto activity and scaled with the number of active layers. The plot illustrates why the fourth layer is not automatically preferred: it increases neutron rejection but also increases accidental loss of calibration-like events.

The same conclusion appears when the veto observables are treated as an event-selection problem rather than as a single threshold. As discussed in Section 5.6.1, a boosted veto BSD trained on aggregate waveform observables rejects substantially more HENSA neutron events than a single visible-energy threshold at the same calibration acceptance. Thus the veto design is not only a matter of adding active material; it is also a matter of retaining enough multiplicity and quality information in the readout to distinguish neutron-correlated veto activity from ordinary accidental peaks.

These results also clarify the role of delayed peaks. Delayed activity is physically central to the scintillator-cadmium design, but a late peak by itself is not a clean neutron tag because real calibration data already contain abundant accidental late activity. The useful signature is more specific: prompt or near-trigger activity together with delayed veto peaks, visible energy, and multiplicity distributed across the veto system. In the peak-level cosmic-neutron study discussed later in this chapter, the condition “prompt plus late” rejects nearly half of the neutron sample while preserving about 96% of calibration-like events, whereas the mere existence of late peaks has little discrimination power. This is the analysis counterpart of the geometry choice: the multilayer scintillator-cadmium structure is designed to create prompt and delayed opportunities for detection, while the event selection must combine timing, multiplicity, and energy information to keep accidental coincidences under control.

## 5.5 Final veto geometry and readout

The final simulated geometry comprises 59 scintillator volumes arranged in six groups — Top, Bottom, Left, Right, Front, and Back — with three layers per group. The nominal symmetric layout would contain 60 panels. In the implemented 59-panel

design one short inner top panel is absent, so that **Top L1** contains three panels rather than four. The asymmetry is treated consistently in the **Geant4** geometry and reflects the mechanical/service accommodation of the realized design. The group/layer logic is summarized in table 5.10; the detailed per-channel PMT serial map used during integration is not reproduced here because it is not needed for the design argument and evolved during prototype commissioning.

<b>Group</b>	<b>Layers</b>	<b>Panels per layer</b>	<b>Lengths</b>	<b>Notes</b>
Top	3	3 / 4 / 4	80 / 150 cm	Final simulated design omits one short inner top panel, giving 59 rather than 60 panels.
Bottom	3	4 / 4 / 4	150 cm	Full long-panel coverage below the detector.
Left	3	3 / 3 / 3	150 cm	Long side coverage.
Right	3	3 / 3 / 3	150 cm	Long side coverage.
Front	3	3 / 3 / 3	80 cm	Short modules on the beam-pipe side in the final simulated geometry.
Back	3	3 / 3 / 3	80 cm	Short modules on the rear side in the final simulated geometry.

Table 5.10: Summary of the final 59-panel simulated veto geometry. The commissioned IAXO-D0 prototype reused 150 cm and 65 cm panels in a 57-signal implementation [111]; this differs slightly from the later 59-panel geometry used in the optimization studies.

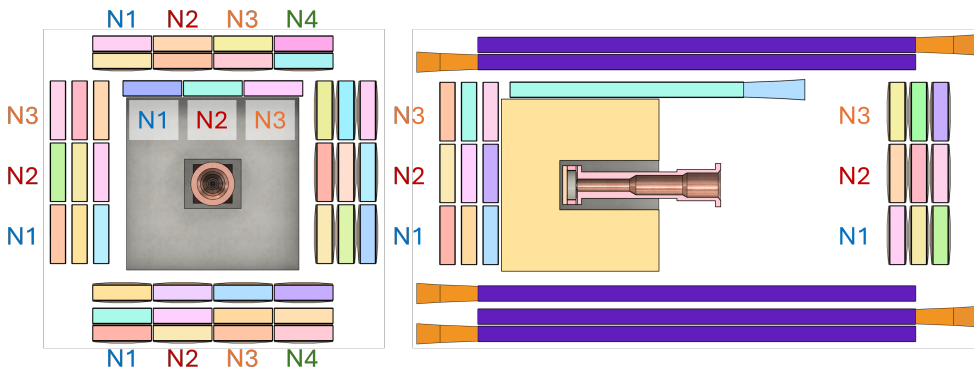



Figure 5.28: Final 59-panel veto geometry used in the *Geant4* studies. The expanded and orthographic views show the three-layer segmentation and the naming convention used to relate geometry volumes and readout channels.

### 5.5.1 Prototype construction and commissioning

The final simulated geometry discussed in the optimization sections uses long 150 cm panels on the sides and bottom together with short 80 cm modules on the front, back, and inner-top positions. The commissioned IAXO-D0 prototype, by contrast, was built from re-used NE-110 bars originally 3 m long and then cut to lengths of 150 cm or 65 cm for the first surface implementation [111]. Each panel is 20 cm wide and 5 cm thick, with a PMT coupled through a light guide at one end. This distinction between the later 59-panel simulated design and the earlier 57-signal commissioned prototype should be kept in mind whenever simulation and data are compared.

In the simulated design, the veto is described as repeated scintillator–cadmium stages, with thin cadmium capture sheets placed between successive scintillator layers so that moderated neutrons emerging from one active volume can be captured close to the next one. In the commissioned prototype, 1 mm cadmium sheets were physically installed between neighboring panels and layers around the detector [111], thereby realizing the same scintillator–cadmium stage concept in the mechanically accessible layout. The overall arrangement follows the multi-layer shielding logic described above, with overlapping panels around the lead shield but unavoidable openings associated with the beam pipe, support structure, and service routing.

The commissioned IAXO-D0 veto did not yet realize the full 59-panel BabyIAXO geometry. In the prototype data-taking configuration up to 57 veto signals could be recorded per event, corresponding to the installed panel set [111]. This distinction is important when comparing simulation and data: the design studies in this chapter refer to the final 59-panel simulated geometry, whereas the first experimental validation was carried out with the slightly reduced commissioned prototype implementation.

Commissioning relied strongly on atmospheric-muon calibration. For a flat 5 cm scintillator panel with density close to  $1 \text{ g cm}^{-3}$ , the most probable muon signal corresponds to about 10 MeV visible energy [111]. This Landau-like reference was used for gain equalization across the installed channels. At the same time, dedicated measurements on the long panels confirmed that the light collected at the far end of a 150 cm bar can be attenuated by roughly a factor of two, which sets a practical limit on how directly calibrated veto amplitudes can be interpreted as deposited energies.




Figure 5.29: Photograph of the commissioned IAXO-D0 surface prototype with the triple-layer veto installed around the lead-shielded detector [111]. This is the 57-signal implementation used for the first experimental validation, not the later 59-panel simulated design adopted in the optimization studies.

The construction and commissioning were carried out within the BabyIAXO detector team. The work presented in this thesis is centered on the simulation framework, the analysis strategy, and the interpretation of the veto response, together with contributions to the detector integration and commissioning tasks required to connect the simulation program with the first surface data.

### 5.5.2 Electronics and readout

Each scintillator panel is read out by a dedicated photomultiplier tube coupled through a light guide. In the BabyIAXO prototype both the Micromegas detector and the veto system were instrumented with four AGET ASICs each, with 64 channels per chip [72, 111]. The readout architecture is therefore a design constraint, not a mere implementation detail: segmentation, timing windows, and channel mapping must all fit within a synchronized dual-branch electronics system.

A relevant feature of the scintillator response is the attenuation of the collected light along the panel length. Measurements reported for the prototype show that, at the far

end of a 150 cm scintillator, the collected light can be reduced by about a factor of two with respect to interactions close to the photomultiplier [111]. Consequently, the pulse height measured in a veto channel cannot be interpreted as a direct deposited-energy estimator without accounting for the interaction position, and the calibrated veto energy should be regarded as a lower limit on the true deposited energy. This effect is incorporated in the detector simulations through light-attenuation models applied to the scintillator and light-guide system, allowing a more realistic description of the waveform amplitudes expected in each channel.

The shaped analog signals are digitized and stored in an adjustable acquisition window. Internally, each channel is sampled continuously over 512 bins by a circular buffer, which makes it possible to choose freely the delay between the start of the time window and the trigger [111]. This flexibility is particularly important for the present application because the Micromegas and veto signals evolve on very different timescales.

In order to operate both subsystems with independent readout settings, the Micromegas and veto signals are handled as separate AGET-based branches synchronized by a common trigger and clock reference [111]. During prototype operation the Micromegas time window was set between 10 and 30  $\mu\text{s}$  depending on the configuration, whereas the veto system was read out in a 100  $\mu\text{s}$  total window with the Micromegas trigger located at 30  $\mu\text{s}$  inside that window [111]. In trigger-centered coordinates this corresponds approximately to  $[-30, +70] \mu\text{s}$ . This common timing architecture allows veto observables and Micromegas candidate events to be merged offline into a single event record with a shared time reference.

For neutron-tagging studies, the veto acquisition window must extend well beyond the prompt response, since delayed signals produced after moderation and capture in the scintillator–cadmium structure are expected on timescales of several tens of microseconds [111, 115, 116]. The precise per-channel high-voltage and AGET mapping evolved during commissioning and are therefore not tabulated here, but the stable group/layer geometry is summarized in table 5.10. The electronics design is thus a central ingredient of the discrimination strategy itself.

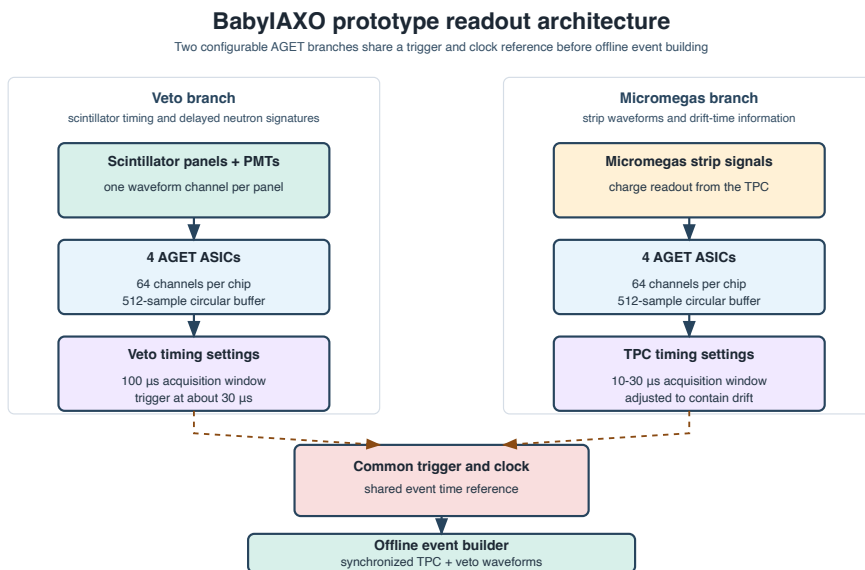


Figure 5.30: Schematic view of the BabyIAXO prototype readout architecture. Both detector branches use four AGET ASICs, but they operate with different timing settings while remaining synchronized by a common trigger and clock reference before offline event building.

## 5.6 Waveform-level veto observables

At analysis level, the waveform information recorded in the veto channels is reduced to a set of observables characterizing the presence of prompt or delayed activity. This distinction is important for the logic of the chapter. The early geometry-optimization studies used idealized deposited-energy thresholds to compare alternative layouts, while the HENSA layer validation and the performance figures use waveform-level observables extracted from shaped traces.

For this reason, the simulation framework is designed to propagate the detector response up to waveform level. In the *REST-for-Physics* chain developed for IAXO-related studies, each scintillator is treated as a dedicated readout channel coupled to one photomultiplier, so the position of the *Geant4* energy deposit can be translated into a propagation distance along the panel and through the light guide. The optical attenuation is then applied using the measured attenuation length of about 400 cm, after which the signal is mapped to the corresponding veto channel and converted into a shaped waveform [115]. In the configuration adopted for the veto studies, the simulated waveforms were sampled in 500 ns bins with a shaping time of 3000 ns, while dedicated TPC timing studies used finer values of 20 ns and 1200 ns for the Micromegas

branch. This approach allows simulated veto data to be analyzed with essentially the same logic as experimental data, which is particularly valuable for validating cuts and estimating the rejection power of different veto configurations.

In the production analysis chain, the shaped traces are processed with `TRestRawPeaksFinderProcess`, which operates on the selected channel type and identifies local maxima using configurable baseline, threshold, distance, and window parameters defined in the common `analysis.rml` file. The process returns the time, amplitude, and multiplicity of the relevant peaks that are later used for veto decisions. Working at waveform level is essential because neutron-tagging events are not characterized by a single large deposited energy. Instead, they often produce a sequence of prompt and delayed peaks with weak amplitude–time correlation, and the most useful observables are therefore the peak times, multiplicity, and channel distribution rather than the total integrated energy alone.

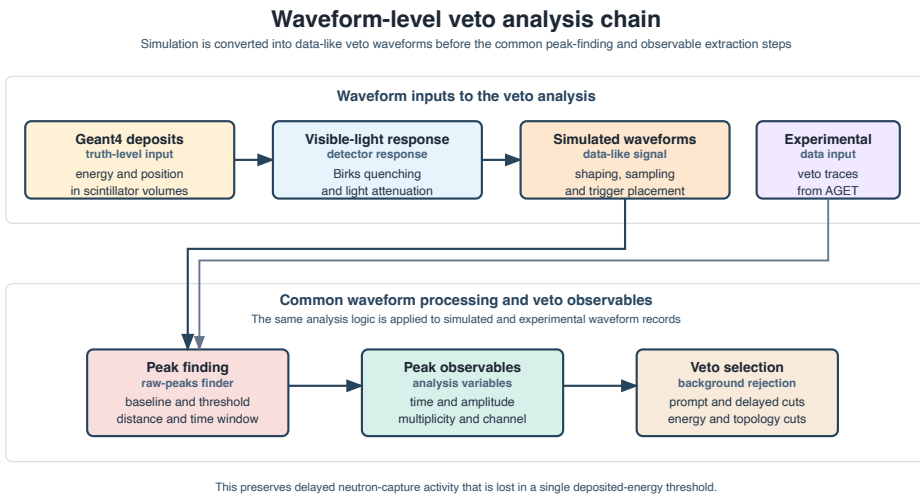


Figure 5.31: Simplified signal-processing chain used to transform `Geant4` energy deposits in the veto scintillators into analysis-level observables. The simulation includes quenching, light attenuation, waveform formation, and peak extraction through the REST raw-peaks finder, so that prompt and delayed neutron-related activity can be characterized through the same observables used in the experimental analysis.

In practice, the use of waveform-based observables provides a more robust description of veto activity than a simple energy-in-volume treatment. This is especially relevant for neutron-induced backgrounds, where the topology of the event in the veto system may contain delayed signals associated with moderation and capture rather than a single prompt pulse. The discrimination capability of the veto system therefore depends not only on its geometrical coverage, but also on the achievable threshold, timing

performance, and segmentation of the readout.

Observable	Typical behavior	Interpretation
<b>Peak time</b>	Prompt component clustered near $t = 0$ ; delayed component extends over several tens of $\mu\text{s}$ with $\tau \simeq 45 \mu\text{s}$ .	Separates immediate moderation/recoil activity from delayed neutron capture in cadmium or scintillator.
<b>Peak multiplicity</b>	About 70% of events yield one peak, about 20% two peaks, and about 5% three or more.	Encodes the stochastic sequence of moderation and capture processes and motivates multiplicity-based veto cuts.
<b>Peak amplitude</b>	Soft distribution with no strong correlation with peak time.	Reflects quenched proton recoils, attenuation along the scintillator, and the weakly localized nature of neutron-induced activity.
<b>Channel correlation</b>	Much weaker than for muons, with especially low correlation between the front group and the rest of the veto.	Confirms the absence of a track-like topology for neutron-induced events and limits the usefulness of geometric coincidence cuts.

Table 5.11: Main waveform-level observables used to characterize neutron-related activity in the veto system [115, 116].

The same reconstructed observables were used to classify the 879 TPC-depositing neutron events by their veto pattern. Explicit neutron capture was present in nearly all of these events at truth level, but a reconstructed veto tag is not guaranteed because the visible response depends on capture-product transport, quenching, light attenuation, waveform shaping, and peak finding. The resulting classes show why a single geometry-only coincidence is insufficient: the useful information is distributed across prompt activity, delayed peaks, multiplicity, and the occasional absence of a reconstructed veto signal.

Timing class	Events	Interpretation
Prompt plus late	412	Prompt secondary activity is followed by late peaks, giving the clearest neutron-sensitive signature.
No veto	153	No reconstructed peak is found; this is the irreducible weakness of an active-veto strategy.
Pre-TPC only	148	Veto activity precedes the TPC signal, consistent with a cascade reaching the veto before the gas deposit.
Near-TPC prompt	73	Activity is concentrated close to the Micromegas trigger time.
Late only	49	Delayed activity appears without a prompt tag, characteristic of capture-like signatures.
Prompt plus delayed	37	Prompt activity is accompanied by delayed peaks in the post-trigger region.
Delayed only	7	A small but distinctive class with only delayed activity.

Table 5.12: Observable veto-timing classes in the TPC-depositing neutron sample. The labels are reconstructed-pattern categories and should not be read as mutually exclusive truth processes.

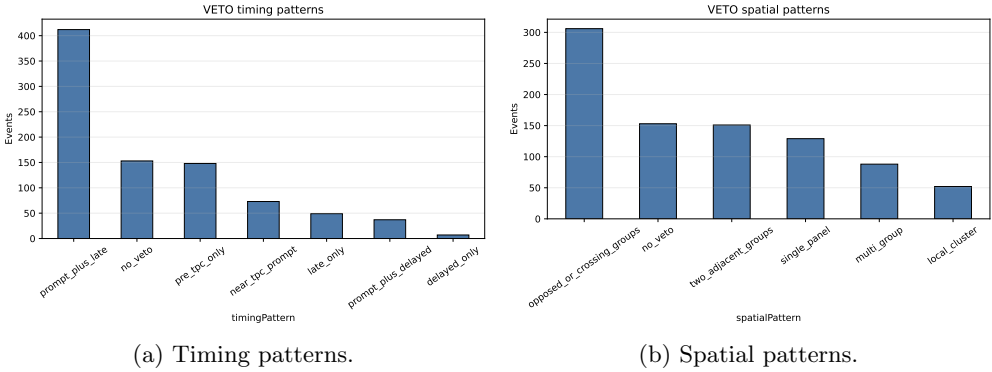


Figure 5.32: Reconstructed veto-pattern categories for neutron-induced TPC events. Spatial patterns are diverse, while the timing distribution provides the most useful handle for defining neutron-enriched control populations.

### 5.6.1 Boosted veto BSD selection

The waveform observables described above can be used either as explicit cuts or as inputs to a multivariate selection. In this thesis the latter approach is treated as the veto analogue of the background-signal discrimination (BSD) selector used for the Micromegas analysis. The purpose is not to assign a unique particle identity to each event, but to build a calibrated ranking variable that separates neutron-correlated veto activity from signal-like accidental veto activity. Only reconstructed analysis observables are used as inputs; truth labels, generator information, event number, and run number are not used by the classifier.

Two related implementations were studied. The design-facing version is a binary boosted-tree selector trained on the three-layer HENSA outdoor-neutron simulation and on experimental calibration-triggered veto activity. For the neutron class, each simulated neutron event is overlaid with one calibration event sampled from the measured accidental-veto template, so that the classifier sees neutron-correlated signals on top of a realistic random-veto environment. For the signal-like control class, the input is the calibration-triggered veto activity itself. The five aggregate inputs are the total visible veto-energy proxy, the reconstructed peak multiplicity, the number of unique veto channels, the number of high-quality veto signals, and the integrated above-threshold waveform response. This is intentionally a compact feature set: it tests how much discrimination can be obtained from observables that are robust against details of the final channel mapping and do not require truth-level information.

The discriminator output is treated as an ordering variable, not as a calibrated neutron probability. Operating points are defined by the calibration acceptance, i.e. by the fraction of calibration-like events that are not vetoed by a given score threshold. At 90% calibration acceptance, a single visible-energy threshold rejects about 51% of the HENSA three-layer neutron sample. The boosted veto BSD increases the rejection to about 76%, and still rejects about 68% of the neutron sample at 95% calibration acceptance. The maximum of the relative  $S/\sqrt{B}$  proxy occurs at a less conservative calibration acceptance of about 83%, where the neutron rejection reaches about 82%. The feature-ranking study shows that the dominant information is carried by the total visible veto-energy proxy and by the number of good veto signals, with peak multiplicity providing a smaller but non-negligible contribution. The unique-channel count is much less important in this aggregate study, consistent with the weak track-like topology of neutron-induced veto activity.




Figure 5.33: Noise-aware boosted veto BSD performance for the three-layer HENSA neutron sample. The upper panels compare neutron rejection and relative  $S/\sqrt{B}$  as a function of calibration acceptance for a single visible-energy threshold, a structured rectangular rule, and the boosted-tree discriminator. The lower panel shows the permutation importance of the aggregate inputs.

The distinction between muon and neutron inputs is also important. Muon-induced events are dominated by prompt, high-amplitude scintillator coincidences associated with a through-going track. In the aligned 20,000-event muon simulation, the median event has five veto peaks, five unique channels, and essentially no delayed-window activity. Neutron-induced events are more heterogeneous: the median event has three veto peaks and three unique channels, but the delayed window contains a median of two high peaks and a visible-energy proxy of about 986 in the analysis units. The neutron response also has a broader time distribution because the veto is sensitive not only to prompt secondaries but also to moderation and capture-related activity. This is why delayed peaks are physically important, but also why they cannot be used alone: the calibration data already contain accidental late activity, so the useful signature is the combination of timing, energy, multiplicity, and signal quality.




Figure 5.34: Comparison of waveform-level veto observables in the aligned 20,000-event muon and neutron simulation samples. Muons are concentrated in the prompt window, whereas neutron-induced events populate a broader delayed tail and show a much larger delayed-peak multiplicity.

An expanded machine-learning study was also performed for the experimental IAXO-D0 data. That study uses three training classes: calibration data, muon simulation with calibration-noise overlay, and neutron simulation with calibration-noise overlay. The input vector extends the aggregate veto variables with prompt-window, neutron-window, delayed-window, and pre-prompt veto observables, together with a small set of Micromegas topology quantities. The Micromegas variables are included only for the experimental candidate-ranking study, where the goal is to identify events whose full detector record resembles the neutron+noise template; the design-facing veto discriminator described above uses only veto observables. The held-out validation gives recalls of 99.2%, 95.8%, and 91.7% for calibration-like, muon-like, and neutron-like classes, respectively. The largest confusion is between neutron-like and muon-like events, which is expected because energetic neutron showers can contain prompt secondaries, while atypical muons can leave weak or delayed residual activity after the prompt-veto selection.




Figure 5.35: Held-out validation of the experimental neutron-like event classifier. The confusion matrix is row-normalized by true class. The feature-importance panel shows that prompt veto energy, delayed-to-total veto energy, total veto energy, and Micromegas topology all contribute to the candidate-ranking classifier.

This multivariate analysis should therefore be interpreted as a structured extension of the explicit veto cuts, not as an independent truth label. For the design study it quantifies the gain available when the full waveform-level veto information is retained. For the experimental comparison it provides a way to rank neutron-enriched event populations while preserving a direct calibration-data estimate of accidental acceptance. The final veto interpretation remains anchored to control samples and to explicit efficiency points, rather than to an unqualified classifier probability.

## 5.7 Simulation performance

Unless explicitly stated otherwise, the performance results in this subsection correspond to the final 59-panel simulated geometry, atmospheric secondaries generated with CRY at Zaragoza latitude, and the full waveform-level analysis chain including quenching, light attenuation, shaping, and peak finding. The veto readout is treated with the prototype-like  $100 \mu\text{s}$  total acquisition window and trigger placement at  $30 \mu\text{s}$ , so that the simulated observables can be compared directly with the data-taking mode of the commissioned IAXO-D0 prototype.

**Definitions.** In table 5.13, *raw* denotes the particle-specific background rate before Micromegas event selection. *Energy + topo + window* denotes the residual rate after the Micromegas energy-region, topological, and readout-window cuts. *+ veto* denotes the same analysis-selected sample after adding a simplified waveform-level veto requirement. The *veto efficiency* is therefore the fraction removed by the veto from the

analysis-selected sample, not from the raw particle flux. The *residual background* is the surviving rate after all applied cuts. Unless stated otherwise, times in this subsection are quoted relative to the Micromegas trigger at  $t = 0$ . The *threshold* refers to the veto-peak threshold used to define a triggerable veto signal; the *prompt window* denotes activity around the Micromegas trigger, and the *delayed window* the later part of the 100  $\mu\text{s}$  total acquisition used to retain capture-related peaks. When the experimental cut windows are quoted later in the original publication convention, they are given explicitly from the start of the veto acquisition window.

### Muon-dominated component

Cosmic muons dominate the raw background rate by a wide margin because of their large flux and high penetration power. Their veto signature is correspondingly robust: they traverse one or more scintillator layers and deposit large signals, typically well above any practical threshold. Figure 5.36 is a deposited-energy diagnostic illustrating this separation between muon signals and the threshold region. Figure 5.38 shows the same behavior at channel level in the segmented detector. The TPC-triggered muon population is compared with the full sample of veto-intersecting muons, showing how the TPC coincidence selects a more geometrically constrained subset while preserving the track-like correlation bands across opposite veto groups.




Figure 5.36: Total energy deposited across all scintillator volumes for muon events, shown as a normalized event-fraction diagnostic. The TPC-triggered sample is concentrated around tens of MeV in the scintillators, far above the few-MeV threshold region relevant for waveform-level veto decisions.

The same conclusion can be inspected at the level of the main veto groups.

Figure 5.37 separates the deposited energy by detector side for the muon events that also produce a TPC signal. These distributions are not used as independent cuts in the final analysis, but they are useful diagnostics of the veto geometry: through-going muons usually illuminate more than one scintillator group, and the relative response of the groups reflects the direction and acceptance of the segmented enclosure. This group-resolved view provides the bridge between the integrated energy spectrum and the channel-level correlation matrix shown next.

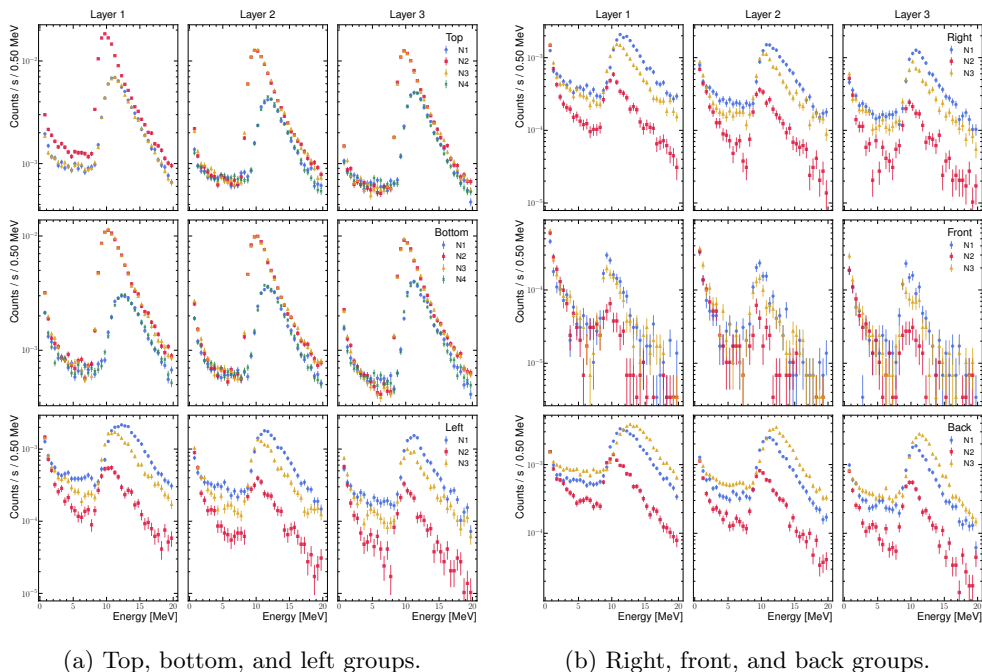


Figure 5.37: Group-resolved scintillator-energy distributions for simulated cosmic muon events that also produce a Micromegas/TPC signal. The figure is a detector-topology diagnostic: it shows that the muon veto response is not confined to a single panel family, but is distributed across the segmented enclosure according to the track direction and detector geometry.

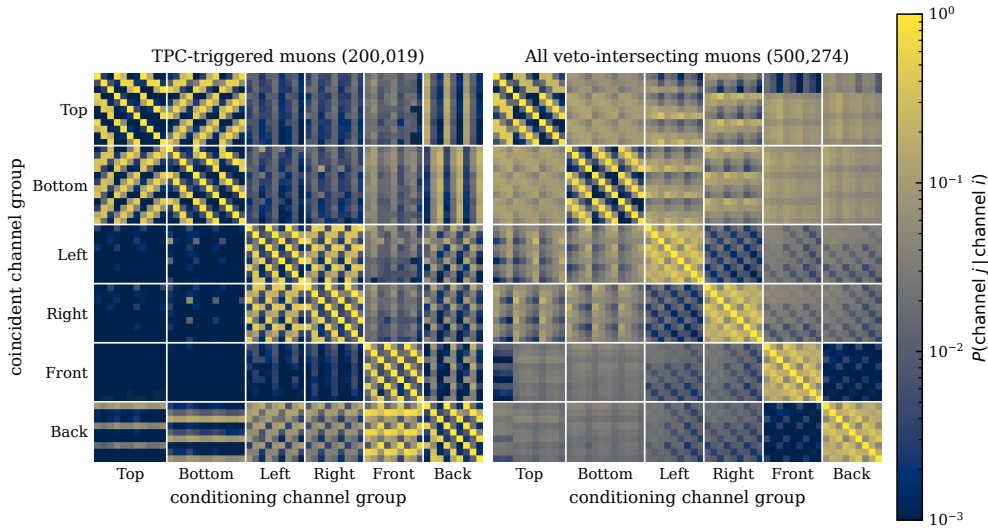


Figure 5.38: Channel-level conditional trigger probability matrices for cosmic muon events in the veto system. The left panel requires a TPC coincidence, while the right panel includes all veto-intersecting muons regardless of whether they ionize the TPC. The strong diagonal and off-diagonal bands reflect the track-like topology of muon events across the segmented veto system.

With a simplified waveform-level veto condition requiring at least one veto peak above threshold, more than 99% of the muon-induced candidates remaining after the Micromegas analysis cuts are rejected [117]. The much larger reduction from the raw muon rate to the final post-veto rate in table 5.13 is therefore not a contradiction: it reflects the combined action of the Micromegas energy, topology, readout-window, and veto cuts.

### Neutron- and proton-dominated component

The neutron case is more difficult because the veto signals are softer, more distributed, and often delayed. High-energy primary neutrons interact in the lead shielding and produce showers of secondaries concentrated around the MeV scale. Their recoil-proton signals in the scintillator are strongly quenched, and their spatial pattern is weak compared with the muon case. Figure 5.39 shows the corresponding waveform-level response. Muons and proton-induced cascades usually produce large veto signals once they reach the scintillators, whereas neutron events have a broad low-energy component and gamma events frequently produce no triggerable veto peak. This is why neutron rejection is much more sensitive to the practical veto threshold and to delayed peak recovery than muon rejection.




Figure 5.39: Waveform-level veto response separated by primary-particle type. Panel (a) shows the normalized distribution of total reconstructed veto peak energy, while panel (b) shows the fraction of events that would remain above a simple total-energy threshold. The curves are normalized independently and should therefore be read as response-shape diagnostics, not as relative cosmic-particle rates.

The same loss of clean spatial structure is visible in figure 5.40, where the strong track-like bands seen for muons are replaced by a more diffuse channel-correlation pattern.




Figure 5.40: Channel-level conditional trigger probability matrices for HENSA outdoor-neutron simulations in the three-layer veto geometry. The left panel shows the TPC-coincident subset, with veto peaks reconstructed in 1597 of 2131 analyzed events. The right panel uses the veto-only processing chain without requiring a TPC signal, with veto peaks reconstructed in 132065 of 274790 analyzed events. Compared with the muon matrices in figure 5.38, the neutron-induced activity is less track-like and less cleanly described by opposite-panel coincidences, motivating the use of time, multiplicity, and delayed-activity observables.

The delayed component is therefore essential. In the simplified waveform-level performance studies, the prototype-like  $100 \mu\text{s}$  total acquisition window captures both prompt fast-neutron activity and a large fraction of the delayed capture signal. Because only about  $70 \mu\text{s}$  of this window lie after the trigger, it retains roughly 79% of an exponential delayed component with  $\tau \approx 45 \mu\text{s}$  before threshold and acceptance effects. Under this definition, the veto reduces the residual neutron-induced background by about 30% after the Micromegas analysis cuts [117]. A further important finding was that cosmic-ray protons are not negligible: because they can also generate secondary cascades in the shielding, the same veto logic suppresses the residual proton component by about 70% [117].

Particle	Raw	Energy + topo + window [cts keV <sup>-1</sup> cm <sup>-2</sup> s <sup>-1</sup> ]	+ Veto
Muon	$3.33 \times 10^{-4}$	$1.85 \times 10^{-7}$	$\lesssim 10^{-9}$
Neutron	$3.33 \times 10^{-6}$	$3.33 \times 10^{-7}$	$\sim 2.3 \times 10^{-7}$
Proton	$7.84 \times 10^{-7}$	$7.40 \times 10^{-8}$	$\sim 2.2 \times 10^{-8}$
Gamma	$1.67 \times 10^{-7}$	0 obs.	—
Electron	$5.08 \times 10^{-8}$	0 obs.	—

Table 5.13: Cosmic-ray-induced background rates from the **Geant4+CRY** simulation at Zaragoza latitude [117]. The veto column corresponds to a simplified waveform-level veto cut applied after the Micromegas analysis cuts. Entries listed as “0 obs.” mean that no simulated events survived the analysis chain in the available sample and should therefore be interpreted as finite-statistics upper limits rather than exact zero rates.

## 5.8 Experimental validation with IAXO-D0 data

The prototype system was operated together with the IAXO-D0 Micromegas detector at surface level in Zaragoza, providing the first direct validation of the veto concept under realistic background conditions. This comparison is especially relevant because the experimental analysis uses waveform-level veto observables, including peak time, peak amplitude, multiplicity, and channel pattern. It therefore tests the same class of quantities emphasized in Sections 5.6 and 5.7, rather than only idealized deposited-energy thresholds.

The scope of the comparison must be defined carefully. The simulations presented above correspond to the final 59-panel geometry, while the first experimental implementation was a commissioned 57-signal prototype based on the same scintillator-cadmium concept. The experimental setup also includes detector effects that are only approximately represented in the simulation, such as channel-dependent gain calibration, light attenuation along the scintillator bars, evolving threshold settings during commissioning, mechanical gaps, and the exact readout configuration. For this reason, the data should not be treated as a one-to-one measurement of the simulated neutron-veto efficiency. Instead, they provide a validation of the physical discrimination strategy: prompt tagging of muon-like events, together with delayed and multiplicity-based rejection of events compatible with neutron-induced activity.

## Published surface cut-flow validation

Stage	Selection	Events	Bkg. level	Calib. eff.
Raw acquisition	Full detector, full energy range	1,305,996	–	–
Micromegas cuts	ROI 2–7 keV, focal spot, Micromegas topological selection	257	–	81.9%
+ Muon veto	Prompt veto discrimination	56	$9.78 \times 10^{-7}$	–
+ Advanced veto cuts	Overall multiplicity < 13; sum of energy < 7034 keV in [23,53] $\mu$ s; multiplicity < 12 in [35,68] $\mu$ s	49	$(8.56 \pm 1.22) \times 10^{-7}$	79.4%

Table 5.14: Experimental cut flow reported for the surface IAXO-D0 + veto prototype data set, corresponding to an effective duration of 52.1 days [111]. Background levels are quoted in counts  $\text{keV}^{-1} \text{cm}^{-2} \text{s}^{-1}$ . The advanced veto cuts preserve about 97% of the calibration efficiency with respect to the Micromegas-cut reference value.

Starting from the full raw acquisition, the Micromegas analysis selects events in the 2–7 keV region of interest, inside the focal-spot region, and passing the standard topological selection. This reduces the data set to 257 X-ray-like candidates with a calibration efficiency of 81.9%. The application of the prompt muon-veto selection then reduces this sample to 56 events. This is the dominant rejection step in the experimental data, corresponding to a reduction of

$$1 - \frac{56}{257} = 78.2\% \quad (5.4)$$

with respect to the Micromegas-selected sample. This behavior is consistent with the simulation picture: muons dominate the raw surface background and, when they produce events that survive the Micromegas cuts, they are usually accompanied by large and prompt scintillator signals.

The advanced veto selection removes 7 additional events from the 56 events left after the prompt veto, leaving 49 events in the final sample. This corresponds to an additional rejection of

$$1 - \frac{49}{56} = 12.5\% \quad (5.5)$$

after the prompt-veto stage. The associated binomial uncertainty is approximately 4.4 percentage points, so the effect is statistically visible but not yet precise enough to extract a detailed neutron-rejection efficiency. The final measured background level is

$$(8.56 \pm 1.22) \times 10^{-7} \text{ counts keV}^{-1} \text{cm}^{-2} \text{s}^{-1}, \quad (5.6)$$

where the quoted uncertainty is dominated by counting statistics. The calibration efficiency decreases from 81.9% after the Micromegas cuts to 79.4% after the full veto selection, corresponding to a relative efficiency of  $79.4/81.9 \approx 97\%$ . This is an important

result because it shows that the delayed and multiplicity-based selection can provide additional rejection without introducing a large penalty for signal-like X-ray events.

The advanced veto cuts reject events with unusually high veto multiplicity or with substantial veto activity in delayed sub-windows, both of which are difficult to reconcile with an isolated X-ray-like Micromegas signal. In the original publication convention, the windows  $[23,53] \mu\text{s}$  and  $[35,68] \mu\text{s}$  are quoted from the start of the  $100 \mu\text{s}$  veto-acquisition window. Since the Micromegas trigger is placed at  $30 \mu\text{s}$ , these correspond approximately to  $[-7, +23] \mu\text{s}$  and  $[+5, +38] \mu\text{s}$  relative to the trigger. The first window therefore straddles the prompt region and early delayed activity, while the second isolates the delayed component more cleanly. This timing convention is important when comparing the experimental cuts with simulation results, where times are generally expressed relative to the Micromegas trigger. The veto-energy variable entering these cuts is a calibrated scintillator-energy proxy rather than a direct local deposited energy, because light attenuation and channel-dependent gain corrections are folded into the observable.




Figure 5.41: Peak time versus maximum veto-energy proxy for the events rejected by the advanced veto selection after the prompt muon cut [111]. The activity is concentrated in moderate-energy and delayed signals rather than in the large prompt pulses characteristic of through-going muons, which supports the interpretation of the advanced veto as a delayed/multiplicity-rich selection.




Figure 5.42: Experimental Micromegas spectrum before veto cuts (red), after the prompt muon veto (green), and after the advanced delayed/multiplicity-rich veto selection (blue) [46, 111]. In the original slide and paper this last curve is labeled as a “neutron” cut; here it is interpreted more cautiously as a selection enriched in delayed or multiplicity-rich veto activity rather than as event-by-event neutron identification.

### Interpretation limits of the published selection

The interpretation of the advanced veto selection should be made with some caution. Although the original analysis refers to this stage as a neutron cut, the experimental observable is not a direct neutron identifier. The veto does not detect the primary neutron itself in an event-by-event sense. Instead, it tags activity expected from neutron-induced cascades: secondary particles produced in the lead shielding, moderation in the plastic scintillator, and delayed capture-related signals in the scintillator-cadmium structure. The rejected events occupy a region of moderate veto-energy proxy and delayed activity, rather than the high-amplitude prompt signals characteristic of through-going muons. It is therefore more precise to describe the advanced selection as a delayed/multiplicity-rich veto selection enriched in neutron-like activity.

This distinction is also consistent with the Monte Carlo studies. In the full waveform-level simulations, the muon component is strongly suppressed by the prompt veto, while the residual neutron-induced background is reduced only partially. The simulation predicts that, after Micromegas cuts, the neutron component is reduced by about 30% by the simplified waveform-level veto, while proton-induced residuals are reduced more strongly, by about 70%. These simulated values are not directly comparable to the 12.5% additional rejection observed in data, because the experimental sample contains a mixture of background components and because the prototype geometry, thresholds,

and calibration are not identical to the final simulation configuration. Nevertheless, the data show the expected qualitative structure: the prompt veto provides the dominant rejection of muon-like events, and a second class of delayed or high-multiplicity veto observables provides additional background discrimination.

The published cut flow is therefore kept as the experimental reference point for the veto-analysis discussion. It provides a baseline against which the refined selections developed in this thesis can be tested: at comparable calibration acceptance, a new delayed, multiplicity, likelihood, or machine-learning selection should reject more of the post-prompt background sample, or provide a clearer physical separation between prompt muon-like, accidental, and neutron-like veto activity. This comparison is more informative than replacing the published figures, because it keeps the first experimental validation and the later simulation-driven optimization conceptually separate.

Observable	Simulation / design expectation	Prototype-data check
<b>Geometry and readout</b>	Final design studies use 59 scintillator volumes in three layers, with one waveform channel per veto volume.	Commissioned IAXO-D0 setup records up to 57 veto signals with the same scintillator-cadmium concept.
<b>Timing scale</b>	Prompt activity occurs within about $1\ \mu\text{s}$ ; delayed captures follow $\tau \simeq 45\ \mu\text{s}$ , motivating a long veto window.	Data use a $100\ \mu\text{s}$ veto window with the Micromegas trigger at $30\ \mu\text{s}$ and delayed sub-windows in the analysis.
<b>Muon rejection</b>	Waveform-level simulations suppress the residual muon component to below the finite-statistics sensitivity after Micromegas cuts.	The prompt muon veto is the dominant experimental reduction after Micromegas cuts, reducing the selected sample from 257 to 56 events.
<b>Delayed and multiplicity rejection</b>	Simulations predict partial rejection of neutron- and proton-dominated residuals, about 30% and 70%, respectively, after Micromegas cuts.	Advanced delayed and multiplicity cuts remove 7 of the 56 post-muon-veto events, a 12.5% additional reduction with about 97% relative calibration efficiency.
<b>Current limitation</b>	Absolute neutron normalization remains tied to the final source-term normalization and detector-response thresholds.	The validation data set covers 52.1 days, so the delayed-component estimate is statistically limited.

Table 5.15: Compact comparison between the veto-simulation expectations and the first IAXO-D0 surface data. The entries are not one-to-one efficiency measurements because the simulated final geometry and the commissioned prototype differ slightly, but they test the same prompt, delayed, and multiplicity-based veto logic.

The comparison can therefore be summarized in four points. First, the experimental data confirm that waveform-level veto observables are operationally useful and can be applied to real Micromegas candidate events without a large loss of calibration efficiency. Second, the prompt-veto response behaves as expected for muon-induced backgrounds,

giving the largest reduction in the measured cut flow. Third, the advanced veto selection removes an additional population of events with delayed or multiplicity-rich scintillator activity, supporting the design motivation of the scintillator–cadmium veto. Finally, the limited size of the published data set and the differences between the commissioned prototype and the final simulated geometry prevent a direct extraction of an absolute neutron-tagging efficiency.

### Thesis extension: neutron-enriched score comparison

For this thesis, the published comparison was extended with a higher-statistics reanalysis designed to answer a more specific question: which events in the experimental background sample are most compatible with the veto signatures expected from cosmic neutrons? The goal is not to assign a unique physical origin to every background event. Instead, the analysis constructs a controlled neutron-like score using calibration data, muon simulations, and neutron simulations, and then studies the experimental events that populate the neutron-enriched tail of that score. This approach makes the comparison more transparent than a single hard multiplicity cut because it keeps the calibration-derived accidental response, the prompt muon component, and the neutron-like delayed component as separate ingredients.

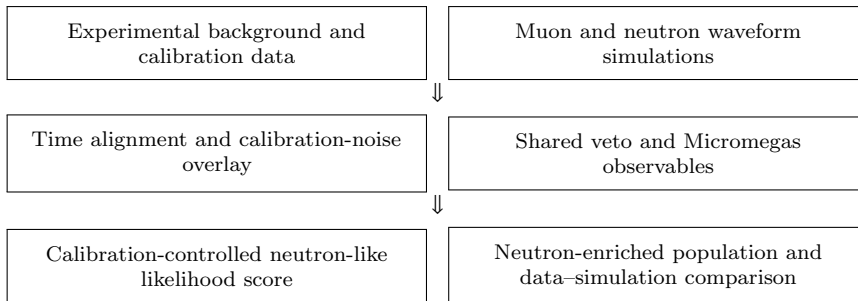


Figure 5.43: Workflow used for the extended experimental comparison. The calibration data serve two roles: they define the accidental-veto environment added to the simulations and they set the accidental acceptance of the neutron-like score.

Three terms are used throughout the rest of this section. A *neutron-like* event is an event whose observables resemble the neutron+noise template more than the calibration or muon+noise control templates. A *neutron-enriched* population is a statistical sample with a larger expected neutron contribution than the parent sample, but with possible contamination from other backgrounds. A *neutron-enriched candidate* is an experimental event selected by the neutron-like score; this label does not imply event-by-event proof that the primary particle was a neutron. This terminology is intentionally conservative because the veto observes correlated scintillator activity and

capture-related signatures rather than the incoming neutron directly.

The experimental input consists of the same prototype data used for the first veto publication. Two flat analysis trees were used: a calibration sample with 139454 events and a background sample with 777752 events. The calibration events are essential because they contain signal-like Micromegas triggers together with accidental veto coincidences produced by the real detector environment. They therefore provide an empirical model of random veto activity, including channel correlations, nonuniform channel noise, and occasional multi-peak structures. The background file contains the surface physics data in which most raw events are expected to be muon related, with a smaller population of events produced by other external or detector-internal backgrounds. Both files correspond to the commissioned 57-signal veto implementation, while the simulation templates described below use the 59-volume design geometry. This mismatch is a known systematic limitation, but it is acceptable for the present purpose because the comparison is based mainly on global timing, multiplicity, and energy-proxy observables rather than on a channel-by-channel efficiency extraction.

The simulation input consists of two dedicated 20,000-event production campaigns run with the same waveform-level analysis chain used for the veto studies. The muon production yielded 21569 analyzed events and the neutron production yielded 20274 analyzed events. In both cases the analysis tree includes the reconstructed veto peak vectors and the Micromegas observables needed to compare the simulated events with the experimental flat trees. Because the REST simulation stores the peak times in a physical time convention while the experimental analysis stores the veto waveform in binned acquisition coordinates, the simulated peak times were aligned according to

$$t_{\text{bin}} = \frac{t_{\text{REST}}}{0.2 \mu\text{s}} + 231, \quad (5.7)$$

where  $t_{\text{REST}}$  is the REST peak time in microseconds and  $t_{\text{bin}}$  is the experimental waveform-bin convention. This alignment was validated with the muon sample: the prompt-muon tag selected 97.65% of the aligned muon simulation, consistent with 97.84% in the earlier reference sample. The corresponding fraction in the experimental background file is 91.13%, while the calibration file has an accidental prompt-tag probability of 11.00%. These numbers show both that the simulated prompt timing is correctly aligned and that the experimental background sample is indeed dominated by prompt muon-like activity before veto rejection.

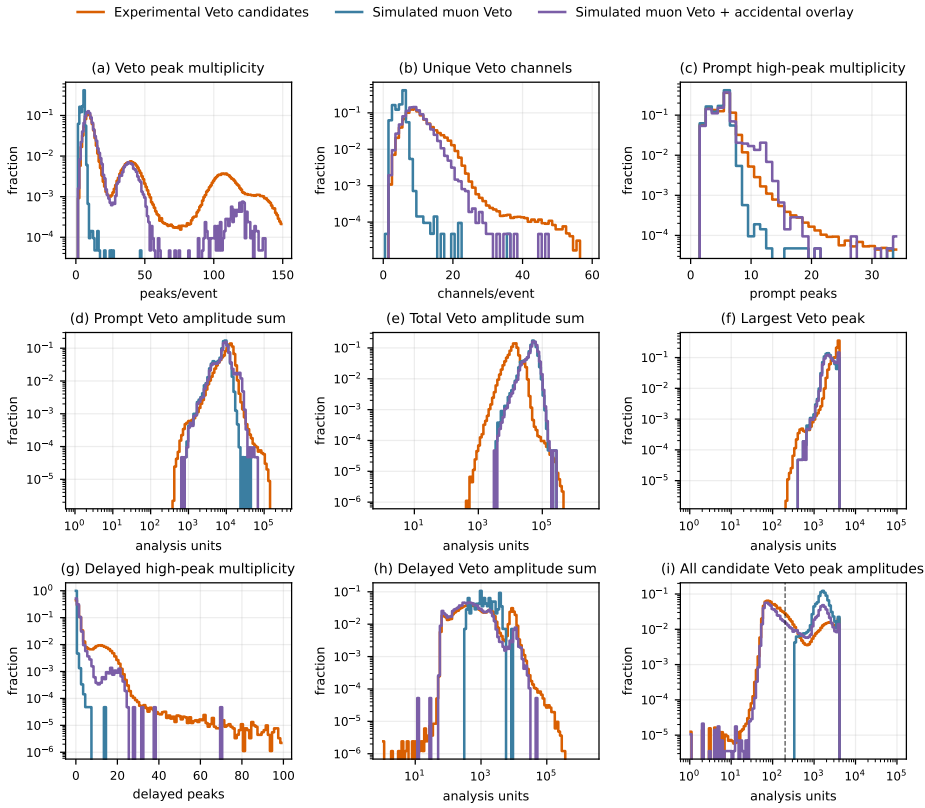


Figure 5.44: Prompt-veto candidate observables in the experimental background data, the simulated muon sample, and the same simulated muons after overlaying accidental veto activity sampled from calibration data. The prompt amplitudes and largest peaks are broadly simulation-like, while the experimental candidates have richer total multiplicity, unique-channel, and delayed-activity tails.

The same prompt definition can be used to inspect the experimental channel topology directly. Figure 5.45 shows conditional channel-coincidence matrices built from veto peaks above 200 analysis units. Events are split according to whether they satisfy the prompt-muon tag  $N_{\text{prompt}} \geq 2$ . The muon-like background population has broad, high-probability channel coincidences, as expected for through-going or showering cosmic activity crossing several veto panels. The non-muon-like and calibration samples are less uniformly correlated and retain visible channel-dependent accidental structure. This comparison is not used as an efficiency calibration for individual scintillators, but it is a useful check that the prompt tag separates a topologically distinct experimental population before the delayed neutron-sensitive score is applied.

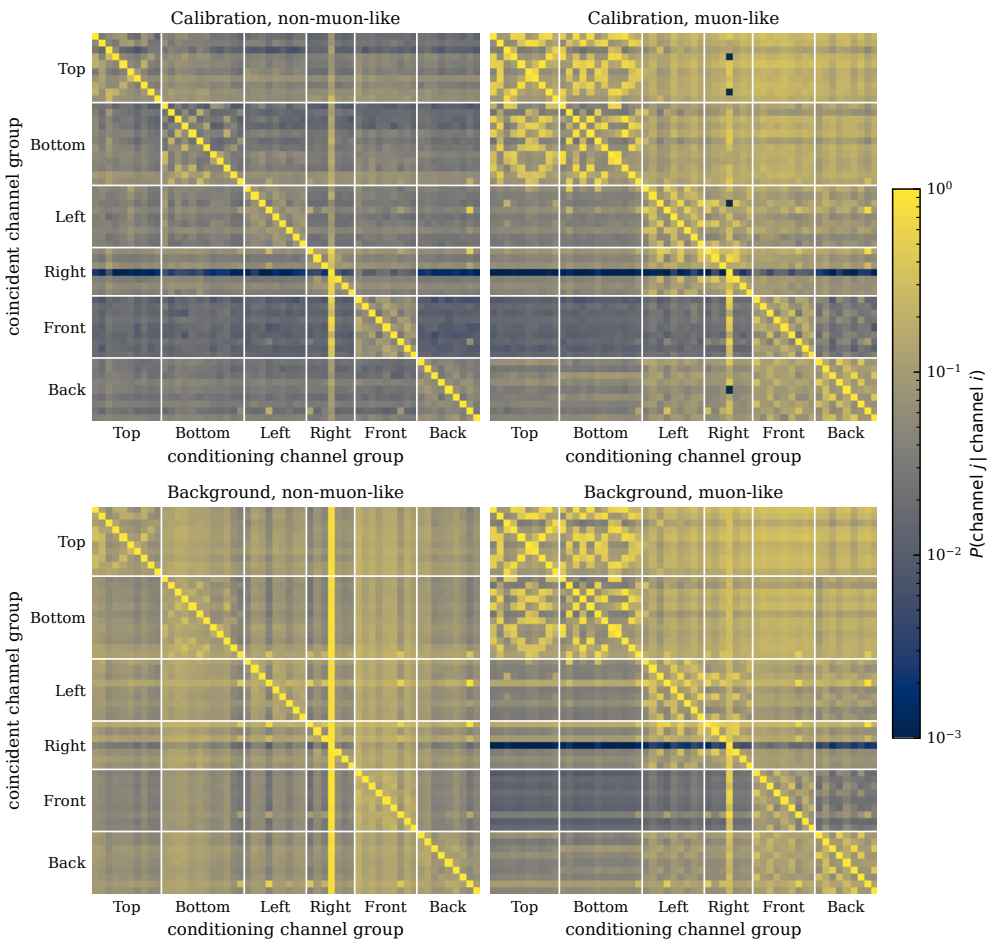


Figure 5.45: Experimental channel-level conditional trigger probability matrices for the IAXO-D0 veto data. The calibration and background samples are split into non-muon-like and muon-like populations using the prompt criterion  $N_{\text{prompt}} \geq 2$  for veto peaks above 200 analysis units in the prompt window. The matrices use the same conditional probability definition as the simulation channel-correlation figures, but only active channels in the commissioned paper readout are shown.

The prompt-candidate comparison makes the calibration role explicit. The prompt definition selects 902484 of 995271 experimental background events, 21062 of 21569 simulated muon events, and 21134 of 21569 simulated muon+noise events. Adding calibration accidentals brings the simulated multiplicity closer to the data, with a median of ten reconstructed veto peaks in both samples, but it does not fully reproduce the delayed tails. This supports the interpretation that the experimental prompt population is dominantly muon-like while still containing real detector accidentals and non-muon activity that must be treated separately in the neutron-enriched comparison.

The calibration sample was also used to add an empirical accidental-veto component to the simulations. For each simulated event, one calibration event was sampled with replacement and its veto-peak vectors were concatenated with the simulated veto peaks. This procedure does not attempt to model the electronics from first principles. Rather, it preserves the measured accidental veto environment, including the real distribution of random channels and amplitudes, while retaining the simulated correlation between the TPC event and the physical muon or neutron shower. The resulting templates are denoted below as muon+noise and neutron+noise. This step is important because a pure simulation sample has too little random veto activity compared with the experimental data, especially in the delayed windows used for neutron-sensitive selections.




Figure 5.46: Aligned veto peak-time distributions for the 20,000-event muon and neutron simulation campaigns after waveform-level analysis. Muons concentrate in the prompt region, while neutron events show a broader delayed component associated with secondary moderation and capture-related activity.

Figure 5.46 illustrates the central distinction exploited by the comparison. The muon simulation has a compact prompt response, with a median of five veto peaks, five unique veto channels, and essentially no delayed activity after the prompt window. The neutron simulation is less sharply prompt and more heterogeneous: the median event has three veto peaks and three unique channels, while the delayed window contains a median of two peaks and a median delayed energy proxy of about 986 in the analysis units. The absolute veto-energy scale should not be overinterpreted, but the timing pattern is robust. Muon-like events are identified by large prompt coincidences; neutron-like events are enriched by delayed and moderate-multiplicity activity.

### Likelihood score and candidate definition

The event classification was performed with a deliberately simple likelihood score. For each event, an observable vector  $\vec{x}$  was constructed from six global quantities:

$$\vec{x} = (\log E_{\text{nwin}}, N_{\text{nwin}}, \log E_{\text{delayed}}, N_{\text{delayed}}, N_{\text{unique}}, \log E_{\text{TPC}}), \quad (5.8)$$

where  $E_{\text{nwin}}$  and  $N_{\text{nwin}}$  refer to a neutron-sensitive delayed window,  $E_{\text{delayed}}$  and  $N_{\text{delayed}}$  refer to the full post-prompt delayed window,  $N_{\text{unique}}$  is the number of unique veto channels with peaks, and  $E_{\text{TPC}}$  is the reconstructed Micromegas energy observable. The score is the log-likelihood ratio

$$S_n = \log p(\vec{x} | \text{neutron} + \text{noise}) - \log p(\vec{x} | \text{calibration/muon} + \text{noise}), \quad (5.9)$$

where the second term represents the non-neutron control population constructed from calibration events and muon+noise simulation. The score is not a proof that an event was caused by a neutron. It is a ranking variable: large positive values identify events whose veto and TPC observables are more compatible with the neutron+noise template than with calibration-like or muon-like activity.

Before applying this score, two quality selections were imposed. Events with a prompt-muon tag were removed using the condition  $N_{\text{prompt}} \geq 2$  for veto peaks with amplitude above 200 analysis units in the prompt window  $190 \leq t_{\text{bin}} \leq 210$ . Events with extreme veto activity, defined as  $N_{\text{peaks}} \geq 100$  or  $N_{\text{unique}} \geq 50$ , were treated as burst-like and excluded from the clean comparison sample. The neutron-enriched candidate threshold was then set by the calibration control sample: the default requirement is that an event lies above the upper 1% tail of the clean calibration score distribution. This gives

$$S_n > 4.668479. \quad (5.10)$$

By construction, this threshold accepts approximately 1% of clean calibration events and therefore provides a direct empirical estimate of accidental acceptance for a signal-like X-ray control population.




Figure 5.47: Distribution of the neutron-like log-likelihood score used in the extended comparison. The threshold is defined from the 1% upper tail of the clean calibration sample, so the selected experimental population is calibrated against the measured accidental-veto environment rather than against simulation alone.

Table 5.16 summarizes the resulting populations. After prompt-muon suppression and burst removal, 65990 background events remain from the original 777752 experimental events. Applying the 1% calibration-tail threshold selects 3877 events. This corresponds to 0.50% of all experimental background events, 5.62% of the prompt-muon-suppressed background sample, and 5.88% of the clean prompt-muon-suppressed sample. The same threshold selects only 12 events in the muon+noise simulation, corresponding to 0.056% of all simulated muon events and 2.77% of the clean prompt-muon-suppressed muon+noise subset. In contrast, it selects 1722 neutron+noise events, or 8.49% of all simulated neutron events and 17.8% of the clean prompt-muon-suppressed neutron+noise subset.

Sample	Events	Prompt tag	Clean after prompt	Median $S_n$	Selected
Calibration	139454	11.00%	123441	-3.80	1235
Experimental background	777752	91.13%	65990	-2.10	3877
Muon+noise simulation	21569	97.98%	433	-2.47	12
Neutron+noise simulation	20274	51.95%	9697	2.70	1722

Table 5.16: Population summary for the extended neutron-like event comparison. The “Selected” column applies the 1% calibration-tail threshold of Eq. 5.10 after prompt-muon suppression and burst removal. The calibration row quantifies the empirical accidental acceptance by construction.

The accidental acceptance is therefore not an abstract estimate; it is measured from the calibration data themselves. At the chosen threshold, a signal-like calibration event has a 1% probability of entering the neutron-like tail after the prompt-muon and burst selections. Equivalently, the neutron-enriched candidate definition would reject or tag about 1% of calibration-like X-ray events because of random veto coincidences or calibration-event structures that resemble neutron-like activity. This value is the appropriate accidental scale for the score-based selection. The muon+noise leakage is smaller in the full simulated muon sample because most muon events are removed by the prompt tag, but it is not zero in the clean prompt-suppressed subset. This residual leakage corresponds to muons with atypically weak prompt signatures, muon-induced secondaries that survive the prompt cut, or accidental-noise configurations that populate the delayed windows.

It is useful to compare this result with the published advanced-veto reduction. The published cut flow removed 7 of 56 events after the prompt muon veto, an additional reduction of 12.5% in the final 2–7 keV focal-spot sample. The score-based study quoted above is broader: it is applied to the full available flat background sample and uses a likelihood threshold anchored to the calibration tail rather than the exact published focal-spot cut flow. With a looser threshold defined by the upper 5% calibration tail, the score selects 9039 of 65990 clean prompt-muon-suppressed background events, or 13.7%. This is close to the scale of the published 12.5% post-muon-veto reduction, although the two numbers should not be interpreted as a direct reproduction of the same selection. Their agreement in scale is nevertheless encouraging: both analyses indicate that a delayed/multiplicity-rich veto observable identifies an additional background population after prompt muons have been removed, while preserving most calibration-like events.

### Properties of the neutron-enriched experimental population

The selected experimental events are not simply high-energy muons that escaped the prompt veto. Their veto signatures are characterized by substantial delayed activity and moderate neutron-window activity. The median selected experimental candidate

has 47 veto peaks, 12 unique veto channels, one peak in the neutron-sensitive delayed window, and seven delayed peaks over the full delayed window. The median delayed energy proxy is 6822, while the median prompt-count observable is zero. This pattern is qualitatively different from the prompt muon template and is the reason these events are assigned large values of  $S_n$ .

The comparison with selected neutron+noise simulations is more nuanced. The neutron+noise events selected by the same score have a median of 9 veto peaks, 8 unique veto channels, one neutron-window peak, and four delayed peaks. Their median delayed energy proxy is 2707, lower than in the selected experimental population, while their median Micromegas energy and hit multiplicity are larger than in the experimental candidates. Thus, the experimental candidates are neutron-like according to the global score, but they are not perfectly described by the present neutron simulation. The data show stronger delayed veto multiplicity, whereas the simulation shows more energetic and more extended TPC topologies.

Median observable	Experimental candidates	Selected neutron+noise simulation
Veto peaks	47	9
Unique veto channels	12	8
Total veto-energy proxy	11249	13513
Prompt-window peak count	0	0
Neutron-window peak count	1	1
Neutron-window energy proxy	999	800
Delayed-window peak count	7	4
Delayed-window energy proxy	6822	2707
Micromegas energy observable	3025	10566
Micromegas hit multiplicity	6	51
$xy$ topology variable	2.31	14.19
$z$ topology variable	2.70	7.59

Table 5.17: Median properties of the experimental events selected by the neutron-like score compared with selected neutron+noise simulation events. The agreement in delayed/neutron-window veto structure supports the neutron-enriched interpretation, while the differences in total delayed multiplicity and Micromegas topology identify the main detector-response and background-composition systematics.

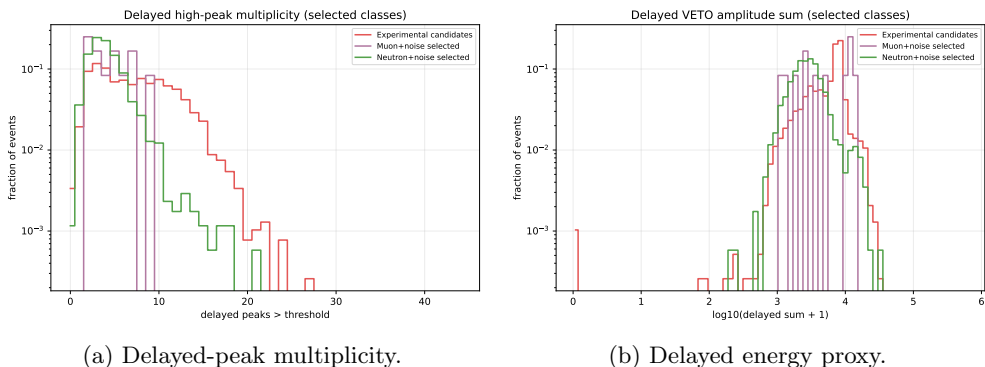


Figure 5.48: Comparison of delayed-veto observables for the experimental neutron-enriched candidates and the selected neutron+noise simulation. The selected data and simulation occupy the same broad neutron-sensitive region, but the experimental population has a stronger delayed multiplicity tail, pointing to either missing detector noise structure, additional background components, or differences between the prototype response and the simulated geometry.

Representative event rasters for the neutron-enriched experimental sample and selected neutron+noise simulation are shown in Appendix 6.12, Fig. 6.52. They are kept out of the main text because they are useful visual diagnostics, but the quantitative argument is already captured by the score distribution, population table, and delayed-observable comparison.

Several conclusions follow from these comparisons. First, the score-based candidate population is plausibly enriched in cosmic-neutron-induced activity, because it is selected by delayed and neutron-window veto features that are inefficient for ordinary calibration events and strongly suppressed in prompt muon simulations. Second, the population is not a pure neutron sample. The present templates include calibration-like accidentals, muons, and neutrons, but the experimental background may also contain gamma-induced events, radon- or activation-related activity, residual proton or hadronic secondaries, muon-induced secondaries with weak prompt veto response, unstable channels, and DAQ artifacts. Any of these components can be partially absorbed into the neutron-like likelihood tail if their observables resemble delayed veto activity. Third, the mismatch in Micromegas topology between selected experimental candidates and selected neutron simulation suggests that the current neutron production does not yet fully reproduce the response of the experimental prototype. This may arise from the gas mixture and analysis configuration, differences between the 57-signal prototype and the 59-volume simulation, threshold and gain effects, or incomplete modeling of the full environmental background composition.

The corresponding Micromegas-energy and hit-multiplicity diagnostics are also deferred to Appendix 6.12, Fig. 6.53. Their role is to document the main residual

mismatch: the neutron template reproduces the delayed-veto character better than the full TPC topology.

The safest interpretation is therefore statistical rather than categorical. At the 1% calibration-tail threshold, approximately 0.50% of all experimental background events, and 5.9% of the clean prompt-muon-suppressed background, form a neutron-enriched candidate population. This fraction is not the cosmic-neutron background fraction in the detector. It is the fraction of events that pass a particular neutron-like selection whose accidental acceptance is fixed from calibration data. The true neutron contribution could be smaller if other backgrounds populate the same delayed/multiplicity-rich region, or larger if a significant fraction of neutron-induced events have weak veto signatures and remain outside the selected tail. Nevertheless, the result is consistent with the published observation that delayed and multiplicity-based veto cuts remove a measurable fraction of the post-muon background while preserving most calibration events.

## 5.9 Systematic limitations and remaining work

### 5.9.1 Design-level systematics

The veto-performance values quoted above should be interpreted as design-level estimates rather than as final absolute background predictions. Several systematic effects remain relevant, both in the source term and in the detector response.

Source of uncertainty	Main impact on the veto studies
Cosmic-neutron normalization	The nominal flux is now tied to the outdoor HENSA 10 GeV spectrum; the remaining uncertainty is the transfer from the Zaragoza measurement to DESY and the comparison with CRY, EXPACS, and literature parameterizations.
Site dependence	Zaragoza/DESY latitude, building overburden, and indoor/outdoor conditions affect the spectrum and normalization of the surface neutron component.
Geant4 hadronic modeling	Neutron multiplication and secondary production in lead depend on the hadronic model and cross-section data used in the transport.
Quenching models	Birks-type scintillator quenching and Lindhard-type gas quenching affect the visible-energy scale for neutron-induced signals.
Light attenuation and PMT gain	Position-dependent attenuation and channel-to-channel gain variations modify the effective threshold for a given deposited energy.
Threshold and timing settings	The chosen prompt/delayed windows and peak thresholds determine the balance between neutron efficiency, accidentals, and dead time.
Geometry details	Mechanical gaps, support structures, cadmium placement, and the beam-pipe opening break the idealized symmetry of the shield and veto.
Finite Monte Carlo statistics	The tails of the rate distributions, especially for gamma and electron survivors, are limited by the available simulation statistics.

Table 5.18: Main systematic limitations affecting the interpretation of the veto-simulation results.

The dominant uncertainty for absolute neutron-background rates is therefore no longer the existence of a measured source term, but the transfer of the HENSA outdoor spectrum to the DESY site and the differences between measured and generator-level descriptions. At fixed source term, the dominant uncertainty for veto efficiency is the effective threshold, which combines quenching, light attenuation, PMT gain calibration, and peak-finding settings into a single observable-level response. Geometry gaps and finite Monte Carlo statistics mainly affect the tails of the residual distributions and the precision with which small surviving components can be quoted.

A full accidental-veto and dead-time estimate requires the measured single-channel rates together with the final online veto logic. In general, for a representative veto-trigger rate  $R_{\text{veto}}$  and coincidence window  $\Delta t$ , the accidental probability scales approximately as  $P_{\text{acc}} \simeq 1 - \exp(-R_{\text{veto}}\Delta t)$ . This quantity is not quoted here because the relevant single-channel rates evolved during commissioning and were not yet fixed in the configuration

used for the first prototype analysis. Empirically, however, the advanced experimental veto cuts preserve about 97% of the calibration events relative to the prompt-veto stage, which provides a first bound on the live-time penalty associated with delayed and multiplicity-based veto observables.

## 5.9.2 Limitations of the neutron-enriched comparison

The main limitations of the present data–simulation comparison are summarized in Table 5.19. They are narrower than the general systematics listed in Table 5.18: the question here is not the absolute uncertainty of the full veto design, but the degree to which the experimental neutron-enriched population can be interpreted with the currently available templates.

Limitation	Impact on the neutron-enriched interpretation
Prototype channel mapping	Experimental data contain up to 57 veto signals, while the simulation uses the 59-volume design geometry; global observables are robust, but channel-resolved efficiencies are not yet directly comparable.
Run-dependent thresholds and gains	The calibration-noise overlay captures the measured accidental environment in aggregate, but does not yet model threshold drifts, PMT gain changes, or channel-by-channel stability over time.
Broken and unstable vetoes	Missing or unstable channels can reduce prompt-muon tagging or produce delayed-window structures that are not present in the idealized simulation.
Gas and analysis configuration	The 20,000-event productions provide useful veto templates, but small differences in gas settings and analysis configuration can affect Micromegas energy and topology observables.
Calibration-noise overlay	Sampling one calibration event per simulated event preserves empirical accidental correlations, but it assumes the calibration noise population is representative of the background periods being compared.
Missing background templates	Protons, gammas, environmental radioactivity, radon-related activity, activation, and DAQ artifacts are not yet included in the likelihood model, so they can be absorbed into the neutron-like tail.
Source normalization	The candidate fraction is a data-selection fraction, not a cosmic-neutron flux measurement; the HENSA-based neutron source term is needed to convert it into an absolute background prediction.

Table 5.19: Comparison-specific limitations affecting the interpretation of the neutron-enriched experimental population. These are the dominant effects that must be resolved before the score-based selection can be converted into a quantitative neutron-background measurement.

The extended comparison also identifies the measurements needed to convert this neutron-enriched selection into a quantitative neutron-background estimate. The most important improvement is a fully matched experimental and simulated analysis chain, including the exact 57-channel prototype mapping, run-dependent thresholds, channel gains, and broken-channel configuration. The second requirement is a broader simulation library that includes not only cosmic muons and neutrons but also protons, gammas, environmental radioactivity, and muon-induced secondaries under the same waveform analysis. The third requirement is a run-by-run accidental model rather than

a single calibration-overlay template, because the delayed-window candidate rate is sensitive to channel noise and burst-like periods. Finally, the HENSA-based neutron source term must be transferred to the DESY site conditions before the candidate fraction can be translated into a BabyIAXO background prediction.

The same list defines what would be required for a publication-quality neutron-population measurement. A robust result should use the exact experimental geometry and readout map, a blinded validation period independent of the data used to tune the score, a run-by-run veto calibration and accidental model, and a complete template set including protons and environmental backgrounds. The neutron template should then be normalized with the measured or otherwise constrained cosmic-neutron source term, and the resulting candidate rate should be propagated through the same Micromegas region-of-interest and topological selections used for the final background model. Under those conditions, the neutron-enriched score could become not only a tagging variable but also a control region for constraining the cosmic-neutron component of the surface background.

In conclusion, the surface data validate the main design logic of the active shield and provide the first experimentally grounded handle on neutron-like veto signatures. The veto behaves as a conventional prompt muon veto and, in addition, provides a measurable delayed/multiplicity-based rejection channel that is consistent with the expected response to neutron-like residual backgrounds. The extended reanalysis shows that a calibration-controlled neutron-like score selects a small but significant experimental population with delayed veto activity, while also revealing systematic differences between data and the current neutron simulation. This result supports the use of a multi-layer scintillator-cadmium veto for BabyIAXO and motivates the next step: a larger and more stable data set, with fixed thresholds and a fully matched simulation geometry, to convert the present neutron-enriched population into a quantitative neutron-veto efficiency and background estimate. The neutron-enriched population defined here will also be useful in the background-model chapter as a validation target for the simulated cosmic-neutron and muon-induced components, because it provides an experimentally selected control sample anchored to the real veto response.

## 5.10 Summary and outlook

The shielding and veto studies lead to a clear design conclusion. Lead remains necessary for suppressing external photon backgrounds, but it is not a complete surface-background solution: for high-energy neutrons it acts partly as a converter, producing a softer and more distributed secondary shower. The veto concept must therefore be coupled to the passive shield rather than treated as a separate muon counter.

The final design response is a three-stage scintillator–cadmium veto surrounding the lead shield. The hydrogen-rich scintillator provides prompt and moderated-neutron sensitivity, the cadmium sheets convert the thermalized component into capture-gamma cascades, and the waveform analysis retains prompt, delayed, multiplicity, and topology information. The HENSA-driven layer scans support the three-layer cadmium configuration as the practical compromise: most of the available neutron-tagging gain is obtained before the fourth layer, while the additional layer increases channel count and accidental-veto exposure. The default physics configuration for this interpretation is high-precision neutron transport with strict isotope handling; `PhotonEvaporation` is used as the nominal cadmium-capture cascade model when prompt-gamma energy partition matters, with the default `ParticleHP` cascade retained as a systematic comparison.

The commissioned IAXO-D0 prototype validates the same strategy in data. The prompt veto gives the dominant muon-like rejection, and delayed/multiplicity-rich selections provide an additional measurable reduction with only a small calibration-efficiency penalty. The extended neutron-enriched comparison strengthens this interpretation by selecting a calibration-controlled population with delayed veto activity, while also showing where the present simulation is incomplete. The next improvements are therefore well defined: a matched 57-channel prototype geometry, run-dependent threshold and noise modeling, broader background templates, and transfer of the measured HENSA neutron source term to the BabyIAXO site conditions.

---

## Chapter 6

# Background Model

6.1	Radiopurity Measurements . . . . .	189
6.2	Background Measurements . . . . .	192
6.3	Simulation and Analysis Methodology . . . . .	192
6.3.1	Workflow overview . . . . .	192
6.3.2	Source-specific simulations . . . . .	194
6.3.3	Background-model component catalogue . . . . .	197
6.3.4	Event types and detector-response chain . . . . .	199
6.3.5	Analysis processes and observables . . . . .	204
6.3.6	Reference samples for cuts and normalization . . . . .	211
6.4	External Background . . . . .	221
6.4.1	Cosmic Rays . . . . .	222
6.4.2	Cosmic-ray event classes and veto survival . . . . .	223
6.4.3	Environmental Background . . . . .	224
6.5	Intrinsic Background . . . . .	234
6.5.1	Gas Contamination . . . . .	234
6.5.2	Radon Contamination . . . . .	235
6.5.3	Detector Materials . . . . .	239
6.5.4	Front-end Electronics . . . . .	243
6.5.5	Shielding . . . . .	248
6.6	Background-model status and uncertainty roadmap . . . . .	253
6.7	Supplementary interaction plots for detector and veto design . . . . .	262
6.8	Supplementary veto simulation campaign metadata . . . . .	265
6.9	Supplementary active-veto design diagnostics . . . . .	267
6.10	Supplementary HENSA veto layer-scan diagnostics . . . . .	270
6.11	Supplementary neutron-tagging diagnostics . . . . .	272
6.12	Supplementary neutron-enriched data–simulation diagnostics . . . . .	272

6.13 Supplementary passive-shielding scans . . . . .	273
6.14 Supplementary environmental-radioactivity plots . . . . .	276
6.15 Supplementary Micromegas system diagrams . . . . .	281
6.16 Supplementary UV-light calibration R&D . . . . .	283

## Introduction

As a rare-event search experiment, IAXO requires a detailed background model to distinguish signal candidates from background events. The purpose of the model is to translate measured activities and external particle fluxes into reconstructed, X-ray-like survivor rates in the Micromegas region of interest. This translation is non-trivial because the same source activity can produce different accepted rates depending on geometry, detector response, topology cuts, veto response, and normalization uncertainty.

This chapter describes the background model developed for IAXO-D0 and BabyIAXO, while noting that many of its components are also relevant to the future IAXO experiment. The model remains a collaborative effort that continues to evolve as new measurements and simulations become available. The work presented here builds on the previous iteration of the IAXO-D0 background model [48], which itself was informed by the CAST background-model studies. The chapter should therefore be read as a definition of the current source catalogue, detector-response chain, and component maturity status, rather than as the final absolute background-rate budget for BabyIAXO. The latter requires the final master rate table, with source normalizations, geometry version, exposure, selection window, surviving rate, and uncertainties fixed for all components.

Source class	Normalization input	Reference analysis	Current status and uncertainty	Thesis contribution
Cosmic muons	Surface cosmic-muon source term and detector-facing generation	2–7 keV for veto validation; low-energy ROI for model studies	Mature simulation and experimental validation; limited mainly by site normalization, geometry matching, and veto thresholds.	Production workflow, detector-response analysis, and veto-observable comparison.
Cosmic neutrons	HENSA/CRY/EXPACS-informed surface neutron spectra	2–7 keV reference for current analysis	Thesis production, still normalization-limited; limited by DESY transfer, hadronic modeling, quenching, and statistics.	HENSA source integration, multilayer-veto optimization, and neutron-history diagnostics.
Environmental radiation	Laboratory gamma/neutron spectra and concrete/floor models	Low-energy ROI	Provisional source study; limited by site-dependent normalization and geometry boundary conditions.	Source construction and detector-response propagation for comparison with cosmic components.
Internal radioactivity	HPGe/LSC material activities or upper limits	Low-energy ROI, with 2–7 keV as final-analysis reference	Partly mature; several components remain upper limits; limited by activities, masses, and source location.	Source-specific restG4 configurations and common reconstruction/cut application.
Gas/radon/surfaces	Gas activity, radon history, surface assumptions	Low-energy ROI	Source studies and upper bounds; limited by time-dependent gas handling and plate-out history.	Separation of gas-borne, plated-out, and material-contamination hypotheses.

Table 6.1: High-level status dashboard for the background-model components. The table is intended to make the maturity of each source class explicit before the detailed methodology and source-specific simulations are discussed.

Window	Role in this thesis	Usage
1–10 keV	Detector-design region	Broad IAXO Micromegas X-ray region used when discussing detector requirements and comparison with earlier Micromegas background goals.
0.1–10 keV	Simulation diagnostic range	Wide low-energy range used in some source-construction plots to verify spectral shapes, leakage mechanisms, and detector-response behavior.
2–7 keV	Reference analysis window	Default window for the final topology-selection studies in this chapter and for comparison with the published IAXO-D0 surface veto analysis.

Table 6.2: Energy-window conventions used in the background-model and veto-analysis chapters. Quantitative comparisons should use the 2–7 keV reference window unless another range is explicitly stated.

## 6.1 Radiopurity Measurements

Radiopurity screening provides the activity normalization for the intrinsic-background part of the model. The measurements summarized in this section are collaboration inputs, principally from the Zaragoza detector and radiopurity teams and from work performed at the Canfranc Underground Laboratory (LSC); they are not claimed here as an original contribution of this thesis. They are nevertheless an essential ingredient of the simulation chain, because every activity limit or measured contamination level

must ultimately be translated into an expected event rate in the Micromegas region of interest.

The screening strategy follows the low-background Micromegas program developed for CAST, TREX-DM, and the IAXO prototypes [118–120]. Material samples are measured mainly with ultra-low-background high-purity germanium (HPGe) detectors at LSC, where the underground overburden suppresses the cosmic-ray component of the counting background. The relevant gamma-emitting isotopes and decay-chain segments include  $^{40}\text{K}$ ,  $^{60}\text{Co}$ , the  $^{232}\text{Th}$  chain, and the  $^{238}\text{U}/^{235}\text{U}$  chains. For the background model, the outcome of these measurements is a set of specific activities, or upper limits, assigned to detector volumes such as the Micromegas readout, field cage, cathode, chamber, shielding, calibration hardware, cables, and electronics. Upper limits are retained explicitly because many selected materials are sufficiently clean that no statistically significant peak is observed in the screening spectrum.

Recent collaboration-meeting updates clarify how these inputs enter the current IAXO-D1 background model. In the 23rd IAXO Collaboration Meeting report on the IAXO-D1 Micromegas setups, M. Jimenez Puyuelo presented the status of the intrinsic-radioactivity model used for the LSC detector studies [121]. The normalization of most detector materials was taken from the TREX-DM and Micromegas material-screening campaigns, with many entries still treated as upper limits. The report also identified the few components that were not merely limits in the current model, notably  $^{39}\text{Ar}$  in the gas,  $^{210}\text{Pb}$  associated with the lead shielding, and  $^{40}\text{K}$  in the readout. This distinction is important: a simulated component normalized to an upper limit should be interpreted as a conservative bound on the corresponding background contribution, not as a measured central value.

The radiopure-electronics program is another collaboration input that affects the material inventory close to the detector. E. Picatoste reported the status of the Micromegas radiopure electronics at the 21st IAXO Collaboration Meeting, including the production and testing of front-end flex circuits, limandes, and FEC/BEC interconnect elements intended to reduce the amount of non-radiopure material inside the shielded volume [122]. This work is relevant for the background model because the electronics are geometrically close to the gas volume and can therefore contribute through compact low-energy deposits or through secondary radiation produced in nearby materials. For this reason, the final detector model must distinguish between material located inside the radiopure boundary and services or back-end electronics that are farther away or shielded differently.

Input source	Detector elements	Use in the background model
Published Micromegas radiopurity program	Microbulk readouts, vessel materials, field cage, calibration components, shielding samples	Establishes the baseline material-screening methodology and provides activity measurements or limits for materials already used in CAST, TREX-DM, and IAXO-related Micromegas detectors [118–120].
LSC and Zaragoza IAXO-D1 screening inputs	Readout, Mylar and cathode materials, lead shielding, copper and structural elements	Provides the activity normalization for the intrinsic-radioactivity simulations in the IAXO-D1 model; most entries are currently upper limits, with specific measured contributions such as $^{39}\text{Ar}$ , $^{210}\text{Pb}$ , and readout $^{40}\text{K}$ treated separately [121].
Radiopure-electronics development	Front-end flex circuits, limandes, FEC/BEC connection elements, nearby service materials	Defines which electronics components can be placed near the detector and which volumes should be represented explicitly as possible internal-contamination sources [122].
IAXO-D1 operation at LSC	Shielded detector, gas system, radon-suppression configuration, calibration hardware	Provides validation data and operational constraints for the model, especially for separating intrinsic radioactivity from gas-borne radon, surface contamination, and residual environmental backgrounds [121, 123, 124].

Table 6.3: Main radiopurity-related inputs used to construct the intrinsic-background model. The activity measurements and detector-operation studies listed here were performed by the IAXO detector and radiopurity teams and are included in this thesis as external collaboration inputs to the simulation work.

Radiopurity measurements alone do not fully determine the low-energy background, because the detector response depends on where the isotope is located, on the geometry between the source and the gas volume, and on the subsequent reconstruction cuts. The procedure used in the model is therefore to simulate each relevant isotope–volume pair with `Geant4`, process the surviving events with the same analysis chain used for experimental data, and only then scale the accepted rate by the measured activity or upper limit. The corresponding intrinsic-contamination simulations are described later in this chapter as part of the full production inventory.

The LSC IAXO-D1 data sets also demonstrate why the material model must be complemented by operational background studies. The detector was operated inside a 20 cm lead shield with a calibration source, nitrogen or radon-free-air flushing inside the shielding, and evolving gas recirculation and buffer configurations [121, 123, 124]. Those measurements showed that changes in gas handling and radon suppression can affect the alpha and low-energy backgrounds substantially, even when the solid materials have been selected for radiopurity. Consequently, the background model separates intrinsic material radioactivity from radon-related and surface-contamination components, rather than absorbing all observed low-energy events into a single material-activity term.

## 6.2 Background Measurements

The background model uses measurements in two different ways. First, radiopurity measurements provide normalization inputs for intrinsic material contaminations, as summarized above. Second, environmental and surface-cosmic measurements define or constrain external source terms: atmospheric neutrons and muons, environmental gammas from laboratory materials, and neutron fields modified by the surrounding shielding and floor. The Zaragoza laboratory measurements and HENSA-based neutron spectra are therefore used as source-term anchors for the detector simulations, while a future DESY-specific characterization is needed before the same catalogue can be converted into a final BabyIAXO site prediction.

In the present version of the model these measurements are treated as inputs with explicit maturity labels rather than as a single closed data set. The surface-veto validation uses the published IAXO-D0 prototype data set [111]; the HENSA-driven simulations are used to study the neutron response of the final multilayer veto geometry; and environmental gamma or material-produced neutron studies are retained as source-specific checks until their DESY normalization is fixed. This separation avoids mixing measured detector performance, simulated source transport, and site-dependent absolute fluxes into one number before the final master rate table is available.

## 6.3 Simulation and Analysis Methodology

The background model simulations and the subsequent event reconstruction were carried out with `REST-for-Physics` [78], using the `restG4` and `restManager` applications described in the software chapter. The Monte Carlo transport stage was defined through source-specific `simulation.rml` files, while the detector-response emulation and event reconstruction were performed with a common `analysis.rml` configuration. This same reconstruction chain is also used for the analysis of experimental data, since the simulated detector response is converted into the same reconstructed event format as real detector acquisitions. This common data model allows the same observables, selection criteria, and background-discrimination procedures to be applied consistently to both simulated and measured events.

### 6.3.1 Workflow overview

The methodology followed for the background model is summarized in Figure 6.1. The first stage consists of defining a source term and an appropriate GDML detector geometry. The source term depends on the origin of the background contribution under study: radioisotope contamination in a detector component, environmental radiation entering from outside the shielding, or cosmic-ray secondaries. These inputs

are then propagated with `restG4`, which performs the `Geant4` transport and stores the event-level truth information as a `TRestGeant4Event`.

In a second stage, `restManager` processes the simulated event through a detector-response chain designed to reproduce the experimental readout observables as closely as possible. This response chain includes both the Micromegas detector readout and the active-veto readout, which are described in `REST-for-Physics` through dedicated readout definitions. Although both systems are handled within the same analysis framework and are ultimately stored in a common event structure, they represent physically different detector subsystems: the Micromegas readout reconstructs the charge signal produced in the gaseous TPC, while the veto readout reconstructs scintillation signals produced in the surrounding veto modules. Consequently, their channel mapping, signal formation, timing, shaping, and reconstructed observables are treated with subsystem-specific parameters and processes.

The reconstructed output is then used to derive the observables employed for the background studies and for the X-ray-like event selection. In this way, the simulated events can be compared with experimental data at the level of reconstructed quantities, rather than only at the level of idealized energy depositions.

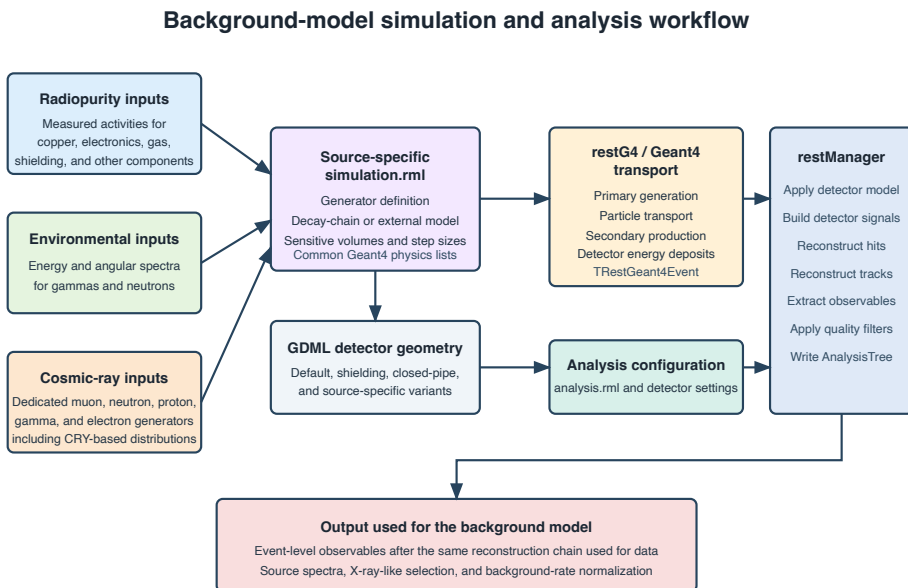


Figure 6.1: Simulation and analysis workflow used for the IAXO background model. Source-specific Monte Carlo configurations are combined with the corresponding GDML detector geometry and propagated with `restG4`, producing `Geant4` truth-level events. The resulting events are then processed with `restManager` through the detector-response and reconstruction chain. Dedicated `REST-for-Physics` readout descriptions are used for the two instrumentally different subsystems: the Micromegas TPC readout and the active-veto readout. The final output is a set of reconstructed observables that can be analyzed with the same selection logic used for experimental data.

This separation between transport and reconstruction was particularly important in the present work. It allowed the same reconstruction chain to be applied consistently to different background sources, while at the same time making it possible to update detector-response parameters, electronics settings, veto thresholds, or analysis cuts without repeating the full `Geant4` transport stage. The same strategy also ensured that simulated events and experimental data were compared using the same reconstructed observables and selection criteria, while preserving the distinct detector-response models required by the Micromegas and veto readout systems.

### 6.3.2 Source-specific simulations

The Monte Carlo transport stage was configured with a family of source-specific `simulation.rml` files. Each file defined the source term for one background component, but the structure of the simulations was kept common: a GDML geometry, a primary-

event generator, the sensitive detector volumes, and the physics lists required for electromagnetic interactions, radioactive decays, and, when needed, hadronic transport. This organization isolates the physical origin of each background contribution while preserving a common interface to the reconstruction chain.

For internal contaminations, the simulated source was defined directly from the GDML volume corresponding to the material under study. This category includes detector materials such as copper, electronics, gas, kapton, mylar, shielding elements, and radon or plated-out progeny. The primary positions were sampled uniformly inside the selected volume, or on its surface when the measured contamination was given as a surface activity.

The activity assigned to each source was obtained from radiopurity measurements performed by the collaboration at the Laboratorio Subterráneo de Canfranc. Depending on the sample, these measurements provide either activities per unit mass or per unit surface. To avoid inconsistencies in the normalization, the geometrical volume was extracted from the GDML geometry using ROOT; together with the material density, this gives the simulated mass of the component. The measured activity and the simulated mass or surface were then used to compute the total activity assigned to the source. This activity provides the normalization factor used to convert the simulated event sample into a background rate.

When the measured activity referred to a parent isotope assumed to be in secular equilibrium, the full radioactive decay chain was simulated. In this way, the generated events preserve both the spatial distribution of the contamination in the detector geometry and the correct normalization to the measured activity of the material.

Environmental backgrounds were also considered in the simulation campaign, although they became less central to the final background model. At the time when these studies were developed, the experimental layout under consideration assumed that the detector would operate inside a laboratory environment, with at least partial underground overburden. Under those assumptions, radiation from the surrounding laboratory materials, such as concrete gammas and neutrons, was a relevant contribution to evaluate.

These backgrounds were treated as external fluxes incident on the detector and shielding system. When a direct simulation from the laboratory boundaries was inefficient, the source term was factorized into two stages. First, energy and angular distributions were obtained for particles emerging from the surrounding materials. These distributions were then used as input spectra for the detector-level simulation, avoiding the need to repeat the expensive laboratory-scale transport for each detector-response or selection study.

In the final experimental scenario, where the detector is expected to operate essentially at surface level and not inside a deep underground laboratory, these

environmental contributions are less important than the cosmic-ray-induced and intrinsic-radioactivity backgrounds discussed in the following sections. Nevertheless, they are included here for completeness, since they were part of the background-model development and provide a useful cross-check of the simulation workflow.

Cosmic-ray-induced backgrounds required a separate set of source generators because their normalization and event topology depend strongly on the particle species, angular distribution, energy spectrum, and exposed surface of the apparatus. Dedicated `simulation.rml` configurations were prepared for the main secondary components at surface level: muons, neutrons, protons, gammas, and electrons. Depending on the study, the primary distributions were defined either analytically or from histograms produced with `CRY`. For the largest production campaigns, geometry-aware sampling strategies developed during this work [84] were used to avoid spending most of the computation time on particles whose trajectories would not intersect the detector geometry. The detailed modelling of the cosmic-ray sources is discussed later in this chapter. The relevant point for the present methodology is that, once the transport stage is completed, cosmic-ray and non-cosmic samples are treated with the same detector-response and reconstruction philosophy.

The detector description used in these simulations was derived from detailed GDML models of the BabyIAXO/IAXO-D0 Micromegas detector, passive shielding, active-veto system, and auxiliary volumes. Different geometry variants were selected depending on the purpose of each sample. Full geometries including shielding and veto volumes were used for cosmic-ray and shielding studies; reduced geometries were used for isolated material-contamination studies; and chamber-focused geometries were used for X-ray calibration, detector-response validation, or cut-efficiency studies. The active gas volume above the Micromegas readout was the main sensitive region used to compute the detector background rate. Additional sensitive volumes, such as plastic scintillators or neutron-capture layers, were included when the goal was to study veto response, neutron-tagging performance, or the history of particles contributing to a reconstructed event.

The `Geant4` transport configuration was kept as uniform as possible across the different source classes, so that differences between samples were driven mainly by the source term and geometry rather than by changes in the physics list. Electromagnetic interactions were described with `G4EmLivermorePhysics`, which is appropriate for low-energy photon and electron transport. Hadronic interactions were handled with the high-precision neutron and binary-cascade models required to describe neutron transport, inelastic interactions, and secondary production in the shielding. For internal-contamination samples, radioactive-decay processes were enabled, including the associated atomic de-excitation mechanisms such as internal conversion, fluorescence, and Auger-electron emission. Production cuts were generally kept at the millimetre

scale for charged particles and photons, while the gas volume was assigned a smaller maximum step size to preserve the topology of low-energy depositions in the X-ray region of interest.

### 6.3.3 Background-model component catalogue

The background model was built as a catalogue of independent source components rather than as a single combined Monte Carlo sample. Each component describes one physical origin of background: cosmic-ray secondaries, environmental radiation, intrinsic radioactivity of detector materials, or radon-related activity. For each source, the simulation defines where the particles are generated, which spectrum or decay chain is used, and how the resulting event sample is normalized to an expected rate.

This organization keeps the model modular. A new material-screening result, an updated cosmic-ray flux, or a revised detector geometry can be incorporated by updating only the affected source component. The simulated events from all components are then passed through the same detector-response and reconstruction chain before being compared at the analysis level.

Table 6.4 summarizes the main source classes considered in the model. The first four entries correspond to physical background-rate components. The final entry is an auxiliary sample used to define the detector response and the X-ray-like selection efficiency applied when the rate components are combined.

Catalogue entry	Simulated components	Normalization input	Role in the model
Cosmic rays	Muons, neutrons, gammas, protons, and electrons generated with sea-level spectra	Differential particle fluxes and generated phase space, folded with the exposed detector geometry	Describes the contribution of cosmic-ray secondaries at surface level, including both direct interactions and secondaries produced in the shielding or detector materials.
Environmental radiation	External gammas and neutrons from laboratory materials, including concrete-wall and floor contributions	Measured or simulated environmental spectra and detector-facing flux	Describes radiation entering the shielding from the surrounding experimental environment. This contribution was studied mainly for completeness and for earlier layout assumptions involving a laboratory setting.
Internal material radioactivity	Copper, electronics, gas, kapton, mylar, teflon, shielding layers, and telescope-side materials	Material-screening activities, component masses or surfaces, and isotope branching ratios	Converts measured radioactivity of detector and shielding components into source-by-source background rates in the analysis window.
Radon and surface progeny	Gas $^{222}\text{Rn}$ , cathode $^{218}\text{Po}$ , cathode $^{210}\text{Pb}$ , and vessel or shielding-surface $^{210}\text{Pb}$	Radon activity, exposed surface area, or assumed plated-out activity density	Accounts for airborne and surface-deposited activity close to the sensitive gas, where compact low-energy events can mimic X-ray-like topologies.
Calibration and cut efficiency	$^{55}\text{Fe}$ calibration and uniform low-energy X-ray samples matched to the detector conditions	Calibration exposure or flat simulation weights; used for efficiency rather than as a background rate	Defines the X-ray-like signal reference, energy response, and selection efficiency applied to the background components.

Table 6.4: Catalogue of simulated components used to build the background model. The first entries correspond to physical background sources normalized to material activities, environmental fluxes, or cosmic-ray fluxes. The auxiliary calibration samples define the detector response and X-ray-like selection efficiency used when the source components are combined.

The individual source components are combined only after detector-response emulation, event reconstruction, and event selection. In practice, each simulated sample is first reduced to the number of events that survive the analysis cuts in the reference energy window. That number is then scaled by the appropriate physical normalization: the measured activity of a material, the activity density of a surface contamination, the radon activity in the gas, or the incident particle flux for external and cosmic-ray sources.

The final background estimate is therefore assembled from reconstructed and selected events, not from idealized energy depositions. This is important because the same source can have a very different impact depending on where the interaction occurs, how the charge or veto signal is reconstructed, and whether the event passes the X-ray-like topology selection. Keeping the source generation, detector response, and normalization as separate steps also makes the calculation easier to update as new measurements or improved detector descriptions become available.

### 6.3.4 Event types and detector-response chain

The output of the transport stage is not analyzed directly from the `Geant4` truth information. Instead, it is converted step by step into event representations that progressively resemble the information delivered by the experimental readout. Figure 6.2 summarizes this chain.

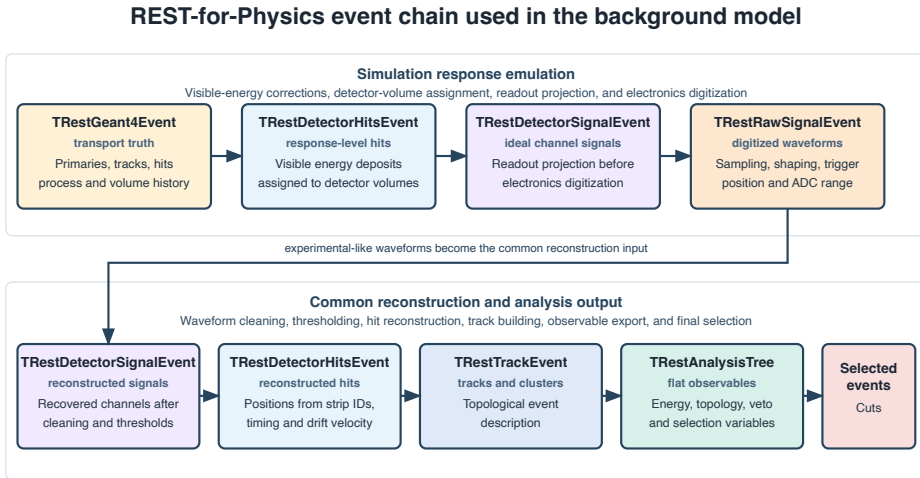


Figure 6.2: Event-type evolution in the `REST-for-Physics` reconstruction chain used for the background model. The chain transforms `Geant4` truth information into digitized raw waveforms and then into reconstructed hits, tracks, flat observables, and selected events, allowing simulated samples to be analyzed with quantities directly comparable to experimental data.

The initial container is the `Geant4` event, implemented as `TRestGeant4Event`. In the following discussion this object is referred to simply as the truth event. It preserves the essential transport information: the primary particle, the hierarchy of secondaries, and the hits produced in the relevant detector volumes. At this stage the event still corresponds to an ideal Monte Carlo description and is therefore not directly comparable to the output of the data acquisition system.

Figure 6.3 shows a representative cosmic-muon event at this stage, pairing a compact print-style event summary with the corresponding geometrical event display.




Figure 6.3: Representative Geant4 event information for a simulated cosmic muon. The left panel shows a compact `TRestGeant4Event` summary, with the primary particle, selected secondary tracks, and sensitive-volume hits. The right panel shows the same type of event as a REST event display, where the muon trajectory crosses the detector and veto geometry before any detector-response emulation or reconstruction is applied.

The first conversion produces a detector-hits event, `TRestDetectorHitsEvent`. At this response-level hit stage, the Geant4 hits are rewritten as detector hits associated with a detector category, in particular TPC or veto. This is the appropriate point to apply processes that modify the deposited energy before digitization, such as quenching in nuclear recoils or attenuation of the light collected in the veto scintillators. It is also the stage where geometrical corrections, such as the rotation between the Micromegas sensitive volume and the readout coordinates, are applied.

The detector hits are then projected onto readout channels, producing a detector-signal event, `TRestDetectorSignalEvent`. These signals are still idealized: they describe the charge or light response assigned to detector channels before the electronics response has been applied. The next conversion produces a raw-signal event, `TRestRawSignalEvent`, which emulates the digitized waveform format delivered by the readout electronics. This step incorporates shaping, sampling, trigger position, dynamic range, and channel-dependent calibration factors. The raw waveform stage is therefore the key bridge between Monte Carlo truth and experimental-like data.

Figure 6.4 shows detector-event views for the same simulated cosmic-muon event after the response has been projected onto the detector readouts. The Micromegas panel uses the readout-channel geometry to show the charge pattern in the TPC, while the veto panel shows the corresponding scintillator modules with visible energy deposits.

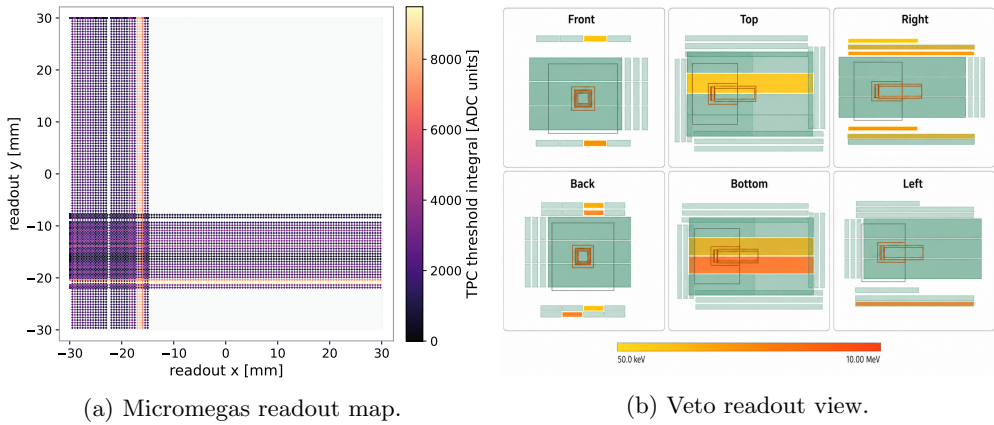


Figure 6.4: Complementary detector-event views for the same representative simulated cosmic muon shown at truth level in Fig. 6.3. The Micromegas display maps the TPC channel-integral response onto the IAXO-D1 readout geometry, with the color scale given in ADC units. The veto display shows the active scintillator panels in the surrounding veto geometry and their associated visible-energy scale.

The waveform representation itself is shown in Fig. 6.5. At this stage the event is stored as a `TRestRawSignalEvent`, with separate digitized traces for the Micromegas and veto channels. The native waveform coordinates are the ADC amplitude and acquisition bin, while the same information can be expressed in physical units using the sampling periods, trigger delays, and calibration factors defined in the analysis configuration.




Figure 6.5: Representative raw waveforms for a simulated cosmic-muon event, read directly from a `TRestRawSignalEvent`. The left panel shows the native digitized representation in ADC units as a function of acquisition bin, with Micromegas X and Y strips shown in different colors and veto channels overlaid on a secondary axis. The right panel shows the same waveforms after applying the timing offsets, sampling periods, and channel calibration factors used in the detector-response analysis. In physical time, the prompt veto response precedes the delayed Micromegas charge signal.

The raw-signal stage is also the point at which simulated and experimental data can be compared most directly at waveform level. Figure 6.6 shows such a comparison for an experimental muon candidate and a simulated cosmic-muon event processed through the same raw-signal representation. The two events are not expected to be identical, since they correspond to different physical particles and trajectories. The relevant validation point is instead that the shaped Micromegas and veto pulses have comparable time ordering, widths, and amplitude scales once the detector-response parameters have been applied.

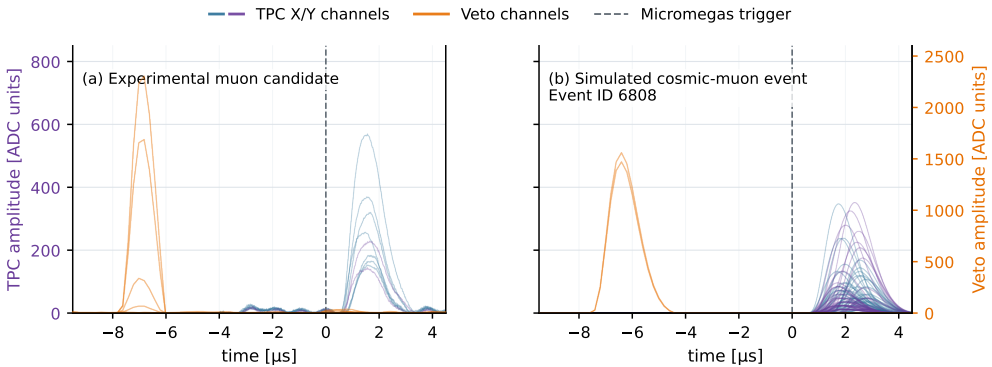


Figure 6.6: Waveform-level comparison between an experimental muon candidate and a simulated cosmic-muon event after detector-response processing. In both panels the Micromegas strip waveforms are shown on the left axis, separated into X and Y readout channels, while veto-channel waveforms are shown on the right axis. The comparison illustrates that the raw-signal emulation produces TPC and veto pulses with realistic relative timing, peak scale, and pulse width before the event is reduced to reconstructed hits and tracks.

Once the raw waveforms have been generated, the reconstruction proceeds in the same direction as the analysis of experimental data. The raw signals are converted back into reconstructed detector signals after thresholding and zero suppression, and then into reconstructed detector hits, where the hit coordinates are recovered from strip identities, timing information, and drift velocity. These hits are grouped into a track event, `TRestTrackEvent`, which stores the topological information used for background discrimination. For the strip readout used here, this representation is naturally expressed as two projected track views rather than as the original three-dimensional Monte Carlo trajectory.

Figure 6.7 shows the reconstructed track representation for the same event sequence. The two panels correspond to the XZ and YZ projections stored in the track event. The colored points show the reconstructed hits and their associated energy, while the solid line is an energy-weighted projected fit added only as a visual guide to the reconstructed topology.




Figure 6.7: Reconstructed track-event representation for the same simulated cosmic-muon event. The Micromegas strip readout produces two projected topological views, XZ and YZ, after raw-waveform reconstruction and hit building. The points are the reconstructed track hits stored in the `TRestTrackEvent`, colored by their reconstructed hit energy, and the solid lines indicate energy-weighted projected track fits used only to guide the eye.

The corresponding scalar observables are finally written to the analysis tree, `TRestAnalysisTree`, where the analysis cuts are applied and the accepted events are used for the background-rate estimate.

Detector signals and detector hits therefore appear twice in the chain, but with different physical meanings. The first occurrence corresponds to the ideal detector response predicted by the Monte Carlo before electronics digitization. The second occurrence corresponds to the reconstructed information that would be available to the analyst after the same signal-processing and reconstruction stages used for experimental data. Distinguishing between these two stages is essential when validating the detector-response model and when interpreting the effect of the analysis cuts.

### 6.3.5 Analysis processes and observables

The common analysis chain was encoded in a single `analysis.rml` configuration. Its purpose was to define one reconstructed observable space for both simulated and measured events. Simulation-only stages emulate the detector response before digitization, while experimental-input stages import measured DAQ waveforms. After these input-specific branches, both event classes follow the same raw-waveform reconstruction, peak finding, hit reconstruction, and topology analysis.

For simulated events, the first stages apply visible-energy corrections, record Monte Carlo truth observables for validation, and convert Geant4 energy deposits into detector hits assigned to the TPC or veto volumes. The roles of these Geant4-level processes are not equivalent. `TRestGeant4QuenchingProcess` is the only process in this block intended to modify the detector response itself. It converts nuclear-recoil energy

deposits into a visible-energy estimate before the event is interpreted as detector response. The other **Geant4**-level analysis processes provide truth-level observables, such as the primary particle, the direction and energy of the source particle, the first interaction in the sensitive volume, or auxiliary veto-truth diagnostics. These quantities are useful for debugging, validating the simulation chain, and understanding individual events, but they are not used as final selection variables. They are too idealized for that purpose, because equivalent information is not available in experimental data.

In the current configuration the quenching process is applied to both the Micromegas gas volume and the scintillator volumes. A visible-energy correction is needed because the energy deposited in **Geant4** is not generally equal to the detector signal. In the Micromegas gas, nuclear recoils lose part of their energy to atomic motion rather than ionization, motivating a Lindhard-type ionization yield [114]. The process uses the hadronic target information stored by **restG4** and evaluates this yield from the recoil energy and from the target isotope mass and atomic numbers. The visible energy assigned to the hit is

$$E_{\text{vis}} = Q_{\text{L}}(E_{\text{R}}, A, Z)E_{\text{R}}, \quad Q_{\text{L}} = \frac{kg(\epsilon)}{1 + kg(\epsilon)}, \quad (6.1)$$

with

$$\epsilon = 11.5E_{\text{R}}Z^{-7/3}, \quad k = 0.133Z^{2/3}A^{-1/2}, \quad g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon. \quad (6.2)$$

Here  $E_{\text{R}}$  is expressed in keV,  $A$  is the target mass number, and  $Z$  is the target atomic number. For the veto scintillators, the relevant non-linearity is scintillation saturation: highly ionizing recoil protons and light fragments produced during neutron moderation yield less light per deposited MeV than electron-like deposits. The same process therefore applies a first-order Birks-law correction [112, 113]. In differential form, the scintillation light yield is approximated as

$$\frac{dE_{\text{vis}}}{dx} = \frac{dE/dx}{1 + k_{\text{B}}dE/dx}, \quad (6.3)$$

or, for an individual **Geant4** hit with deposited energy  $E_{\text{dep}}$  over an effective step length  $\Delta x$ ,

$$E_{\text{vis}} \simeq \frac{E_{\text{dep}}}{1 + k_{\text{B}}E_{\text{dep}}/\Delta x}. \quad (6.4)$$

The production configuration uses a configurable Birks constant initialized to  $k_{\text{B}} = 0.126 \text{ mm/MeV}$ , and estimates  $\Delta x$  from the local spacing of neighboring **Geant4** hits when an explicit step length is not stored. This value is used as a nominal organic-plastic-scintillator setting, adequate for a first-order veto light-yield model and consistent with values commonly adopted in **Geant4**-style plastic-scintillator simulations [125]. It

should not be interpreted as a dedicated calibration of the BabyIAXO scintillators. Direct measurements in plastic scintillators report material-dependent values, including a somewhat larger value for BC-408, so  $k_B$  remains configurable and should be varied in detector-response systematic studies [126]. Using a smaller  $k_B$  corresponds to weaker quenching and therefore larger visible energies for recoil-proton-like deposits, which is relevant for veto-threshold studies. The corrected visible energy is written back to the copied Geant4 hit collection before the detector-hit and raw-waveform stages, so the downstream response simulation receives the quenched detector energy rather than only an auxiliary Geant4 diagnostic summary.

The size of this correction was checked with a representative cosmic-neutron validation sample. In that sample, the energy-weighted visible energy in the Micromegas gas was reduced by only about 0.22% on average, because Lindhard quenching is applied only to the small subset of gas hits identified as nuclear-recoil-like hadronic deposits; the median event was unchanged, although individual argon recoil hits could be reduced by factors of a few. In the veto scintillators the effect was much larger: the Birks correction reduced the energy-weighted visible signal by about 25.7%, consistent with the high ionization density of recoil protons and other charged secondaries. Photons, including neutron-capture gamma-cascade hits recorded as photon deposits, are explicitly left unquenched by this process; in the same validation sample, the total energy carried by those hits was unchanged within numerical precision.

After this visible-energy correction, the event is converted from the truth-event representation into the detector-hit representation. The corresponding conversion and the subsequent projection onto detector channels depend on the detector readout metadata described in the software chapter and illustrated in Fig. 6.4. For this reason the conversion is not described here as a separate physics model: its role is to make the detector geometry, sensitive-volume aliases, and readout-channel geometry consistent before the response processes act on the event. The detector-response stage then introduces the effects that separate an ideal energy deposit from a measurable detector response: scintillator light attenuation, coordinate alignment, electron diffusion in the gas, finite resolution, and readout-plane diagnostics before digitization.

The first of these effects is the attenuation of light in the active-veto scintillators. A veto hit in `Geant4` is a local energy deposit in a plastic panel, but the measured veto signal is produced only after scintillation light has propagated through the panel and light-guide system to the photomultiplier. The amount of light collected therefore depends not only on the deposited energy, but also on where the interaction occurred within the scintillator. This point is particularly important for the IAXO-D0/BabyIAXO veto geometry because many panels are long bars read out from one end. The same recoil-proton or capture-gamma energy deposition can produce a different waveform amplitude depending on whether it occurs close to the photomultiplier or

near the far end of the panel. Consequently, a fixed peak threshold in the analysis corresponds to a position-dependent deposited-energy threshold in the scintillator.

This effect is implemented with `TRestDetectorLightAttenuationProcess`. The process acts only on hits identified as veto hits; TPC hits are passed through unchanged. For each veto hit, the readout metadata is used to determine the veto readout plane and channel associated with the hit position. The distance  $d$  from the hit to the effective readout end of the corresponding panel is then used to attenuate the hit energy according to

$$E_{\text{att}} = E_{\text{vis}} \exp\left(-\frac{d}{\lambda_{\text{att}}}\right), \quad (6.5)$$

where  $E_{\text{vis}}$  is the visible energy after quenching and  $\lambda_{\text{att}}$  is the effective light-attenuation length. The same geometrical path length also gives a propagation-time correction,

$$t_{\text{att}} = t + \frac{d}{v_{\text{eff}}}, \quad (6.6)$$

with  $v_{\text{eff}}$  the effective light speed used for the veto channel. The nominal attenuation length used in the veto studies is of order 400 cm, consistent with the measured and datasheet-scale attenuation behavior discussed in the shielding and veto chapter. This is not a full optical-photon simulation; reflections, surface treatments, coupling details, and channel-to-channel gain variations are absorbed into effective parameters. Nevertheless, it captures the leading geometrical effect that controls whether neutron-related activity in a long scintillator panel remains above threshold after light collection.

The second major response effect is electron diffusion in the Micromegas gas. Whereas light attenuation modifies the veto amplitude, diffusion modifies the spatial and temporal distribution of the TPC charge cloud before it is projected onto the strip readout. Ionization electrons created far from the Micromegas plane drift for a longer time and undergo a larger random walk than electrons created close to the readout. This broadens the charge distribution both transversely, in the readout plane, and longitudinally, along the drift-time direction. Diffusion is therefore directly connected to the reconstructed width, cluster multiplicity, hit spread, and track-shape observables that enter the X-ray-like selection. It also provides the practical link between the detector-response chain and the gas-transport calculations performed with `Garfield++/Magboltz`, summarized in the software chapter.

In the analysis chain this is handled by `TRestDetectorElectronDiffusionProcess`. The process reads the gas pressure, drift field, work function, Fano factor, and longitudinal and transverse diffusion coefficients from the `TRestDetectorGas` metadata when those values are available. These gas parameters are precomputed with `Garfield++` for the selected gas mixture and operating field, rather than simulated microscopically inside each Geant4 event. For a detector hit with energy  $E$ , the process

converts the hit into an effective number of primary electrons,

$$N_e \simeq \frac{E}{W}, \quad (6.7)$$

where  $W$  is the mean energy required to create one electron-ion pair in the gas. Depending on the configuration, the number of electrons may include Poisson or Fano fluctuations, and it may be capped for computational efficiency in high-energy events. Each electron is then displaced with Gaussian diffusion widths that scale with the square root of the drift distance  $z_d$ :

$$\sigma_T = D_T \sqrt{z_d}, \quad \sigma_L = D_L \sqrt{z_d}, \quad (6.8)$$

where  $D_T$  and  $D_L$  are the transverse and longitudinal diffusion coefficients expressed in the units used by the gas metadata. The transverse diffusion spreads charge between neighboring readout strips, while the longitudinal diffusion broadens the arrival-time distribution and therefore the reconstructed coordinate along the drift direction. The process also includes an optional attachment probability, allowing electrons to be lost during drift if an attachment model is enabled.

This response model is deliberately intermediate in complexity. It is more realistic than simply moving the full hit energy to the readout without broadening, because it produces drift-distance-dependent cloud sizes and finite statistical charge granularity. At the same time, it avoids the cost of a full microscopic electron-transport simulation for every event in a large background campaign. For the background model this is the appropriate compromise: the simulated X-ray-like observables remain tied to measured or Garfield-derived gas properties, while the same reconstruction chain can still be applied to millions of source-specific Monte Carlo events and to experimental data.

The diffusion model does not, by itself, reproduce every contribution to the measured detector resolution. Residual effects such as gain non-uniformity, avalanche fluctuations, electronics noise, imperfect calibration, and channel-to-channel response variations can broaden the reconstructed energy and hit distributions beyond the Garfield-derived diffusion expectation. For this reason the response chain also includes `TRestDetectorHitsSmearingProcess`. This process applies an empirical smearing to detector hits before digitization, with separate settings for the TPC and veto subsystems when required. Its role is not to replace the physical diffusion model, but to absorb the remaining resolution needed to match calibration data, in particular the width and shape of the reconstructed  $^{55}\text{Fe}$  peak. In practice, the diffusion parameters fix the gas-transport scale, while the smearing parameters provide a controlled detector-response tuning step when calibration runs show that diffusion alone gives an overly narrow or otherwise idealized response.

The conversion from detector hits to raw waveforms is central to the realism of the

method. It uses separate sampling and shaping parameters for the TPC and the veto system, a fixed waveform length of 512 bins, and trigger delays chosen to reproduce the position of the signal within the acquisition window. This step also applies the calibration factors that place simulated signals on the same ADC scale as experimental data. For the Micromegas readout, the response is anchored with the  $^{55}\text{Fe}$  calibration peak: the conversion from deposited energy to shaped waveform amplitude is tuned so that the reconstructed 5.9 keV peak appears at the same position in simulation and data, after the same raw-signal reconstruction is applied. For the veto system, the corresponding reference is the through-going muon response. Each scintillator panel is allowed to have its own effective calibration factor because the measured muon peak is not universal: it depends on panel length, light-collection geometry, orientation, optical quality, coupling, and the individual readout channel response. The veto calibration therefore aligns the simulated and measured muon-peak positions module by module, rather than imposing a single global scintillator scale. In other words, the simulated raw signals are not simply proportional to the total deposited energy; they are calibrated and time-structured waveforms that can be processed in the same way as measured data.

This transformation also depends on the detector readout metadata described in the software chapter. That metadata provides the mapping between DAQ identifiers and physical detector elements, including the Micromegas strip geometry and the veto-panel aliases. This layer makes it possible to interpret simulated and measured waveforms in the same detector-centered coordinate system rather than as anonymous channel numbers.

Once the raw waveform representation has been reached, the process chain reproduces the standard experimental pre-processing. It attaches the readout metadata, removes masked channels, applies baseline and common-noise corrections, computes channel-level observables, and searches for TPC and veto peaks. The veto peak finder is particularly important because it extracts the timing, multiplicity, amplitude, and channel information used later for veto selections. The subsequent reconstruction recovers detector signals and spatial hits from the waveform representation, after which hit and track analyses provide the topological observables used for X-ray-background discrimination.

Analysis stage	Processes in the final configuration	Role in the reconstructed observable space
Simulation truth and conversion	<code>Geant4QuenchingProcess</code> , <code>Geant4AnalysisProcess</code> , <code>Geant4ToDetectorHitsProcess</code>	Apply the Geant4-stage visible-energy correction, store truth diagnostics for validation, and map Geant4 deposits into detector-hit objects. The truth observables are not used for the final selection.
Detector-response emulation	<code>DetectorLightAttenuationProcess</code> , <code>DetectorHitsRotationProcess</code> (before), <code>DetectorElectronDiffusionProcess</code> , <code>DetectorHitsSmearingProcess</code> , <code>DetectorHitsReadoutAnalysisProcess</code> (before)	Model light losses, coordinate alignment, charge diffusion, finite resolution, and pre-digitization readout-plane quantities.
Digitization and response calibration	<code>DetectorHitsToSignalProcess</code> , <code>DetectorSignalToRawSignalProcess</code>	Convert detector hits into strip and veto waveforms with the chosen shaping, sampling, trigger delay, dynamic range, and TPC/veto calibration factors.
Experimental DAQ input	<code>RawMultiFEMINOSToSignalProcess</code> , <code>RawFeminosRootToSignalProcess</code>	Provide the measured-data entry points before joining the common raw-waveform reconstruction chain.
Readout metadata and channel masking	<code>RawReadoutMetadataProcess</code> , <code>RawSignalRemoveChannelsProcess</code>	Attach the detector-channel mapping and remove inactive or noisy channels before waveform analysis.
Raw waveform conditioning	<code>RawSignalRangeReductionProcess</code> , <code>RawBaseLineCorrectionProcess</code> , <code>RawCommonNoiseReductionProcess</code>	Emulate the finite ADC range for simulation and apply baseline and common-noise corrections to TPC and veto channels.
Raw waveform observables and peaks	<code>RawSignalChannelActivityProcess</code> , <code>RawSignalAnalysisProcess</code> , <code>RawPeaksFinderProcess</code>	Extract channel activity, amplitudes, integrals, threshold observables, peak times, peak multiplicities, and veto-channel information.
Signal and hit reconstruction	<code>RawToDetectorSignalProcess</code> , <code>DetectorSignalChannelActivityProcess</code> , <code>DetectorSignalToHitsProcess</code>	Recover detector signals and convert them into reconstructed spatial hits.
Reconstructed-hit observables	<code>DetectorHitsReadoutAnalysisProcess</code> (after), <code>DetectorHitsAnalysisProcess</code> , <code>DetectorHitsGaussAnalysisProcess</code> , <code>DetectorHitsRotationProcess</code> (after)	Fill readout-plane, spatial, and Gaussian-width observables after the experimental-like reconstruction.
Track observables	<code>DetectorHitsToTrackProcess</code> , <code>Track2DAnalysisProcess</code>	Build the track representation and compute the two-dimensional topological variables entering the X-ray-like selection.

Table 6.5: REST-for-Physics process groups used in the final analysis configuration. The `TRest` class prefix is omitted in the process column for readability. Simulation-only and experimental-input branches are shown together because both are projected into the same reconstructed observable space before the final selection.

The final observables employed in the background model are therefore reconstructed quantities rather than ideal transport quantities. This distinction is important because many backgrounds are rejected not only by their total deposited energy, but also by their temporal and topological signatures: channel multiplicity, hit dispersion, track topology, veto coincidences, and delayed signals. Treating these effects at the waveform and reconstruction level provides a more realistic estimate of the residual background than a simple energy-in-sensitive-volume approach.

The observables generated along this chain are stored in the `TRestAnalysisTree`, which becomes the common input for selection optimization, control plots, and source-by-source background estimates. This provides an important practical advantage of the framework: once the reconstruction chain has been executed, threshold scans, topology studies, or veto-timing optimizations can be performed directly on the stored observables without re-deriving the low-level quantities from the event containers. In this sense, the common `analysis.rml` file defines not only the reconstruction logic, but also the common observable space in which the different background sources can be compared and combined.

### 6.3.6 Reference samples for cuts and normalization

In addition to source-specific background simulations, dedicated reference samples were used to characterize the response to X-ray-like events and to validate the analysis cuts. These samples are conceptually different from the background source productions. Their purpose is not to estimate a physical source rate, but to define the signal acceptance of the reconstruction and to provide a controlled reference against which background events can be selected or rejected. They therefore occupy a central position between the detector-response model and the final source-normalized background table.

Two classes of signal-like simulation samples were used for this purpose. The first class consists of  $^{55}\text{Fe}$ -like calibration simulations. The dominant manganese  $K_\alpha$  line at 5.9 keV lies inside the axion-search energy interval and produces compact photoelectric conversions in the gas. This sample is used to validate the response near the standard calibration energy and to compare the simulated reconstructed spectrum with measured calibration data, as discussed in Section 3.6.1. The same sample also provides useful examples of topologies that are X-ray-like in energy but not in spatial structure. One such event is shown in Fig. 6.8. The simulated  $^{55}\text{Fe}$  photon undergoes photoelectric absorption, but the atomic relaxation produces an argon fluorescence photon that travels several millimetres before being reabsorbed. The result is a two-site charge topology, rather than a single compact conversion. REST reconstructs this as two separated charge clusters in both strip projections, so the event is rejected by observables that test whether most of the energy is concentrated in a single dominant track. This example illustrates why track observables provide a particularly powerful route for rejecting background-like events while preserving genuinely point-like X-ray candidates.

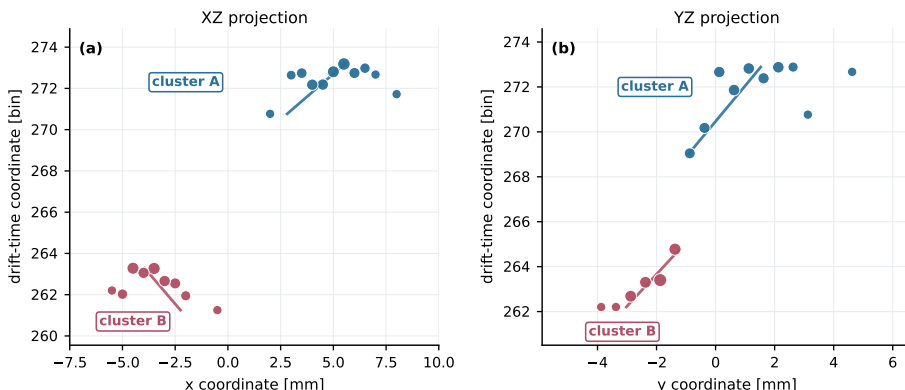


Figure 6.8: Simulated  $^{55}\text{Fe}$  calibration event reconstructed as two separated charge clusters. The event is not a Compton scatter: the primary 5.90 keV photon undergoes photoelectric absorption, and an argon fluorescence photon is reabsorbed at a second site. The reconstruction finds two clusters in both the XZ and YZ projections. Since the dominant reconstructed cluster carries only about half of the total track energy, the event fails the point-like topology selection.

The second class consists of uniform low-energy gamma simulations, in which primary photons are generated over a flat 0–12 keV energy range. This sample is used to map the cut efficiency across the full low-energy region rather than only at the  $^{55}\text{Fe}$  line. The remaining reference samples in Table 6.6 provide the corresponding background rejection and accidental-veto information.

Sample	Role	Monte Carlo configuration	Use in the analysis
$^{55}\text{Fe}$ calibration	Detector-response validation at 5.9 keV	Point-like X-ray source, gas matched to the target background sample, chamber-focused geometry	Energy calibration, peak shape, response comparison with calibration data
Uniform 0–12 keV gamma	Signal-like acceptance across the low-energy region	Flat low-energy photon spectrum, same gas and reconstruction chain as the calibration sample	Efficiency of energy, fiducial, and topology cuts as a function of reconstructed energy
Background source samples	Source-specific residual background	Cosmic, environmental, radon, and contamination sources in the relevant detector geometries	Raw and cut-surviving background populations to be multiplied by source normalizations
Experimental calibration data	Empirical accidental veto model	Real calibration events with no physical correlation to simulated X-ray photons	Veto-noise overlay and accidental signal-loss estimate

Table 6.6: Dedicated samples used to define the X-ray-like selection and connect reconstructed simulation observables to the final background-rate calculation. The first two rows define signal acceptance, while the background source samples define rejection factors for each physical source.

The production campaign keeps the calibration and background samples gas-matched. For the final cosmic-ray comparison presented below, the  $^{55}\text{Fe}$ , muon, and

neutron samples were reconstructed with the same argon–isobutane detector-response settings used by the IAXO-D1 simulation chain. Equivalent Xe-Ne- $i\text{C}_4\text{H}_{10}$  calibration productions are also part of the broader BabyIAXO program, but they are combined only with background samples reconstructed with the same gas, pressure, drift field, and readout configuration. Keeping the reconstruction stage identical is essential: the efficiencies must be expressed in terms of the same observables that are later applied to the background components. In practice, the most relevant reconstructed quantities are the fiducial or readout-plane energy, the number of reconstructed hits, the spread of the charge cloud in the readout plane, the balance between the two strip directions, and the track observables produced after hit clustering.

The X-ray-like selection can be written schematically as a sequence of cuts,

$$C_x = C_E \cap C_{\text{fid}} \cap C_{\text{topo}}, \quad (6.9)$$

where  $C_E$  selects the reconstructed energy region of interest,  $C_{\text{fid}}$  removes events outside the active or well-reconstructed detector region, and  $C_{\text{topo}}$  rejects events with extended or track-like topology. The uniform X-ray sample is the natural dataset for measuring the corresponding signal efficiency,

$$\varepsilon_x(E) = \frac{N_{\text{uniform}}(E | C_x)}{N_{\text{uniform}}(E)}. \quad (6.10)$$

This definition can be used either to produce reconstructed-energy efficiency curves or, after applying the appropriate event weights, to estimate the efficiency for an assumed axion signal spectrum. The  $^{55}\text{Fe}$  sample provides a complementary check at the calibration energy: it tests whether the Monte Carlo response produces the same compact event population, reconstructed peak position, and peak width observed in measured calibration data.

### Machine-learning definition of the X-ray-like topology cut

The refined X-ray-like topology selection used in this thesis is based on a boosted decision tree (BDT) trained on reconstructed `TRestAnalysisTree` observables. The purpose of the BDT is not to replace the physical interpretation of the Micromegas cuts, but to implement a multivariate version of the same idea: an axion-induced X-ray candidate should be a compact, localized event in the gas, while many background candidates contain extended charge, multiple fragments, asymmetric strip projections, or veto-correlated activity. The classifier is therefore treated as an analysis cut defined in observable space, not as a truth-level background tagger.

All events first pass the reference energy requirement

$$2 \text{ keV} \leq E_{\text{rec}} \leq 7 \text{ keV}, \quad (6.11)$$

which is the energy interval used for the final analysis in this work. The reconstructed energy is then removed from the classifier inputs. This prevents the BDT from learning a source-specific energy spectrum and forces it to use topological information within the analysis window. The signal-like training population is defined by  $^{55}\text{Fe}$  calibration events or gas-matched  $^{55}\text{Fe}$  simulations processed with the same reconstruction chain. The background-like population is constructed from a mixture of simulated background components, with the relative event weights chosen from the background-model normalization. This mixture prescription is important. During development, several per-source classifiers were trained against individual samples such as muons or neutrons to understand the available rejection power, but such classifiers would be too specialized for the final background model. The production selector is instead trained once against the combined background mixture and then evaluated source by source. The source identity, particle type, `Geant4` history, event origin, run label, and any simulation-only classification variable are excluded from the input feature list. They are used only for weighting, bookkeeping, and post-selection interpretation.

The BDT score is computed with the `HistGradientBoostingClassifier` implementation in `scikit-learn` [127]. For a reconstructed event with observable vector  $\mathbf{x}$ , the score is

$$s_{\text{BDT}}(\mathbf{x}) = P(^{55}\text{Fe-like} \mid \mathbf{x}), \quad (6.12)$$

as returned by the classifier probability estimator. With this convention the score is bounded between zero and one; larger values indicate a more X-ray-like topology. The score is nevertheless not interpreted as an absolute physical probability, because it depends on the training mixture, class weights, and finite simulation statistics. Only the ordering of events and the selected operating threshold are used in the final cut.

The topology cut is

$$C_{\text{BDT}}(s_0) : s_{\text{BDT}} > s_0. \quad (6.13)$$

The threshold  $s_0$  is chosen by maximizing the expected sensitivity figure of merit in the 2–7 keV window,

$$\mathcal{F}(s_0) = \frac{S(s_0)}{\sqrt{B(s_0)}}, \quad (6.14)$$

where  $S(s_0)$  is the number of signal-like calibration or simulated X-ray events surviving the threshold and  $B(s_0)$  is the weighted number of background-mixture events surviving it. This choice supersedes the earlier convention of fixing the threshold at 80%  $^{55}\text{Fe}$  acceptance. The 80% point remains useful as a reference operating point and for comparison with previous cut definitions, but the final analysis uses the threshold that

maximizes Eq. 6.14.

Two feature sets were compared. The baseline set used global hit and readout quantities: hit multiplicities, charge-cloud widths,  $x/y$  balance variables, and energy-balance observables. The final background-signal discrimination (BSD) selector uses the track-level feature set adapted from the TMVA BDT workflow developed for the IAXO-D1 argon analysis and translated here to the common `REST-for-Physics` observable export. In the thesis text these variables are referred to by their physics meaning, not by their full branch names. The exact branch mapping is kept in the analysis scripts, while the analysis itself can be summarized as the topology vector

$$\mathbf{x}_{\text{topo}} = (f_{\text{trk}}, \sigma_z, \sigma_x, \sigma_y, N_x, N_y, b_E, b_\sigma). \quad (6.15)$$

Symbol	Observable	Physics interpretation	Role
$f_{\text{trk}}$	Dominant-track energy fraction	Fraction of the reconstructed event energy assigned to the largest track or cluster. Compact X-ray conversions are expected to be close to one.	Dominant
$\sigma_z$	Longitudinal width	Width of the dominant cluster along the drift-time direction. Extended or multi-site deposits broaden this coordinate.	Important
$\sigma_x, \sigma_y$	Transverse widths	Widths of the dominant cluster in the two Micromegas strip projections. They distinguish point-like charge clouds from elongated or split deposits.	Important
$N_x, N_y$	Projected hit multiplicities	Number of active strips or reconstructed hits in each projection for the dominant track. Large values usually indicate extended topology.	Important
$b_E$	Projection energy balance	Balance of the reconstructed charge between the two strip projections. It checks whether the event is reconstructed consistently in $x$ and $y$ .	Supporting
$b_\sigma$	Projection-width balance	Balance of the reconstructed transverse widths in the two projections. It is sensitive to asymmetric or poorly reconstructed topologies.	Supporting

Table 6.7: Physics-level observables used by the BSD/BDT topology selector. The reconstructed energy  $E_{\text{rec}}$  defines the 2–7 keV analysis window but is not used as a classifier input.

These observables are directly tied to the physical difference between an X-ray conversion and a background residual event. A  $^{55}\text{Fe}$  event is expected to be compact and dominated by a single reconstructed track or cluster. Cosmic-ray, neutron-induced, radon-related, or material-radioactivity events can fall in the same reconstructed energy window, but they more often contain extended charge, multiple fragments, or an

uneven distribution of energy between the main cluster and the rest of the event. The learned ranking confirms this picture. The dominant observable is  $f_{\text{trk}}$ , the energy fraction contained in the largest reconstructed track. Its permutation importance in the experimental validation training is more than an order of magnitude larger than that of the transverse widths, which are the next most useful variables. This is physically expected: isolated X-ray conversions tend to place nearly all reconstructed energy in one compact object, whereas many residual backgrounds enter the low-energy window through partial deposits, secondary fragments, or split charge patterns. The width and hit-multiplicity variables then define how compact that dominant object is. The balance variables  $b_E$  and  $b_\sigma$  are retained as supporting consistency checks, but their lower importance shows that the main separation is not simply an  $x/y$  calibration imbalance.

The BDT was evaluated against two transparent reference methods. The first is a binned log-odds classifier built from the same observables, which approximates the product of one-dimensional signal-to-background likelihood ratios. The second is a greedy interval-cut procedure that sequentially chooses one-dimensional rectangular intervals using  $S/\sqrt{B}$  as the optimization criterion. Each step searches an allowed interval in a single observable, applies the interval that gives the best figure of merit, and then repeats on the remaining observables. A single decision tree and an extra-trees ensemble were also tested as interpretability checks. The single tree produced readable nested rules but was less competitive as a final selector, while the extra-trees ensemble performed close to the gradient-boosted classifier. These methods are less flexible or less stable than the BDT, but they provide useful cross-checks and help identify which observables carry most of the rejection power.

Selector	$^{55}\text{Fe}$ eff.	$B$ kept	$B$ acc.	$\mathcal{F}$
<i>Measured argon calibration/background split</i>				
BDT, max. $\mathcal{F}$	79.66%	8/5102	0.157%	$2.13 \times 10^4$
BDT, 80% ref.	80.00%	9/5102	0.176%	$2.02 \times 10^4$
Log-odds, 80% ref.	80.00%	10/5102	0.196%	$1.91 \times 10^4$
Manual cuts	80.42%	16/5102	0.314%	$1.52 \times 10^4$
Adaptive intervals	80.13%	29/5102	0.568%	$1.12 \times 10^4$
<i>Gas-matched <math>^{55}\text{Fe}</math>/cosmic-neutron simulation</i>				
BDT, max. $\mathcal{F}$	69.85%	3/3495	0.086%	$1.38 \times 10^5$
BDT, 80% ref.	80.00%	7/3495	0.200%	$1.04 \times 10^5$
Log-odds, 80% ref.	80.00%	38/3495	1.09%	$4.45 \times 10^4$
Simple X-ray cuts	93.42%	53/3495	1.52%	$4.40 \times 10^4$

Table 6.8: Representative validation-split performance of topology selectors developed for the X-ray-like event selection. The measured rows use the July 2024 argon calibration/background optimization files after the common preselection around the  $^{55}\text{Fe}$  calibration peak. The simulation rows use gas-matched  $^{55}\text{Fe}$  and cosmic-neutron samples after reconstructed-energy and feature preselection. The final background-model selector follows the same BDT prescription but trains against the source-weighted background mixture in the 2–7 keV analysis window.




Figure 6.9: Row-normalized BSD/BDT selection matrices for the  $S/\sqrt{B}$ -optimized operating point in the measured argon validation split and in the gas-matched  $^{55}\text{Fe}$ /cosmic-neutron simulation validation. Each cell gives the raw number of events and the corresponding fraction of its true validation class. The upper-left cell is the retained X-ray-like efficiency, while the lower-left cell is the background leakage that enters the sensitivity figure of merit.

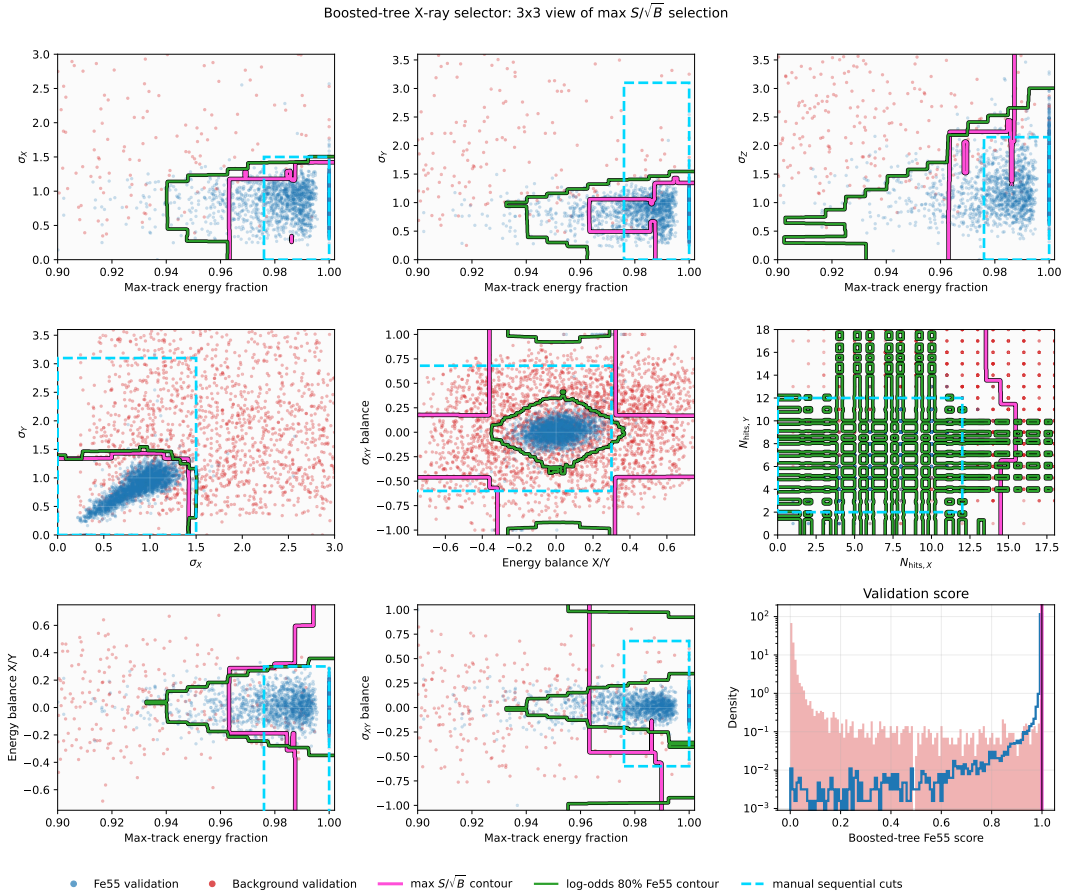


Figure 6.10: Two-dimensional projections of the measured argon calibration/background validation sample used to compare topology selectors. Blue points correspond to  $^{55}\text{Fe}$ -like calibration events and red points to background validation events. The magenta contour shows the BDT boundary at the  $S/\sqrt{B}$ -optimized threshold, the green contour shows the binned log-odds selector at 80%  $^{55}\text{Fe}$  acceptance, and the dashed cyan lines show the manual sequential cuts. The panels are projections of a higher-dimensional selection, so a point apparently inside one two-dimensional boundary can still fail another observable.

Table 6.8 shows the behavior that motivated the adoption of the track-informed BDT. On measured calibration and background optimization data, the  $S/\sqrt{B}$ -optimized BDT keeps nearly the same calibration efficiency as the historical 80% working point while reducing the number of validation background events by roughly a factor of two relative to the manual sequential cuts. At a fixed 80% calibration efficiency, the BDT and binned log-odds selectors perform similarly, whereas the adaptive interval cuts are less effective because a sequence of one-dimensional rectangles cannot reproduce the

useful correlations between observables. On the neutron simulation comparison, the BDT provides stronger rejection than the log-odds and simple X-ray-like cuts. The gain is important because neutron-induced residuals that survive the energy preselection are already the difficult, compact tail of the neutron population. The BDT is therefore used as the common topology selector, while the log-odds, manual cuts, and adaptive rectangular cuts are retained as transparent robustness checks. Figure 6.9 gives the same operating point in confusion-matrix form, emphasizing that the relevant off-diagonal entry is the small background leakage into the X-ray-like sample rather than the overall classification accuracy. Figure 6.10 makes the corresponding observable-space behavior visible: most of the separation is already organized by  $f_{\text{trk}}$ , widths, and hit multiplicities, but the BDT boundary can follow their correlations more naturally than rectangular cuts.

The same selection definition is also applicable to experimental data. This is an important conceptual point rather than a technical afterthought. The BDT is built on scalar observables stored in the `TRestAnalysisTree`, not on Monte Carlo truth quantities. Experimental calibration and background runs processed through the corresponding `REST-for-Physics` chain therefore produce the same type of energy, hit, balance, and track variables as the simulated samples. During this thesis, the experimental work consisted mainly of analyzing existing IAXO-D0 prototype data rather than performing new laboratory campaigns. Those data were nevertheless essential: the  $^{55}\text{Fe}$  calibration runs provide the experimental analogue of the signal-like population, background runs test the behavior of the selectors on real surface data, and calibration-triggered events provide the empirical veto-noise sample used in the accidental-coincidence studies. In this sense, the final BDT selection is a simulation-trained, calibration-anchored, and data-applicable selector. It can be applied to both simulated and measured `REST-for-Physics` analysis trees, provided that the gas, readout, calibration, and reconstruction configuration are matched. At this stage of the analysis, no veto observables enter the BDT. The experimental-data application shown in this section should therefore be interpreted as a Micromegas-only topology comparison: the published IAXO-D0 surface analysis selected the 2–7 keV focal-spot sample with conventional Micromegas cuts before applying the prompt and delayed veto selections [111], while the BDT provides an alternative topology selector built from the reconstructed Micromegas observables alone. The appropriate comparison with the published result is thus a pre-veto cut-flow comparison: number of events and  $^{55}\text{Fe}$  calibration efficiency after the published Micromegas cuts versus after the BDT topology selection, evaluated on the same experimental background and calibration data set. The extension to veto information is a separate analysis layer. It is introduced in the shielding and veto chapter, where prompt, delayed, multiplicity, and veto-pattern observables can be compared directly with the published prompt-muon and

advanced-veto reductions.

The active veto introduces a second selection, denoted here as  $C_{\text{veto}}$ . For cosmic muons and many neutron-induced events, veto activity is a real physical part of the event, produced by prompt charged particles, electromagnetic secondaries, moderation, or neutron capture products. For genuine X-ray signal events, however, the external veto is not physically correlated with the X-ray conversion in the gas. Any signal rejection caused by the veto must therefore be interpreted as an accidental-coincidence loss. This distinction is important enough that signal and background samples must be treated differently.

For signal-like X-ray simulations, the veto branches are not left in an unrealistically empty state when veto cuts are studied. Instead, the simulated X-ray event is overlaid with veto activity sampled from experimental calibration events. Calibration events provide an empirical measurement of random veto peaks, channel occupancy, multiplicity, and timing structure in the detector environment. After this overlay, the veto survival probability for signal-like events is

$$\epsilon_{\text{sig}}^{\text{veto}} = 1 - P_{\text{acc}}(C_{\text{veto}}), \quad (6.16)$$

where  $P_{\text{acc}}(C_{\text{veto}})$  is the accidental probability that unrelated veto noise satisfies the veto rejection condition. This is the correct quantity to apply to an axion-like X-ray signal. It is also the quantity that determines the dead-time-like cost of using a given veto selection.

For background components with physical veto correlations, the treatment is different. The veto decision is applied to the reconstructed veto response produced by the same detector simulation. If appropriate, calibration-like veto noise can be added on top of the simulated veto response in order to reproduce the accidental component present in the experimental data. This is particularly relevant for comparisons with measured background runs, where the observed veto pattern is the superposition of true particle-induced activity, electronic noise, random coincidences, and the effect of missing or unstable channels. The prompt muon cuts and delayed neutron-sensitive observables discussed in the shielding and veto chapter therefore enter the background model as reconstructed veto selections, not as truth-level labels.

The final source contribution is obtained by combining the raw source normalization with the event-level survival of the X-ray-like and veto selections. For a source  $i$ , the most general expression is a weighted sum over reconstructed events,

$$R_i^{\text{after}} = \sum_{j \in i} w_{ij} I_j (C_x \cap \bar{C}_{\text{veto}}), \quad (6.17)$$

where  $w_{ij}$  contains the source normalization and any geometrical, spectral, live-time,

activity, surface, or phase-space weight assigned to event  $j$ , and  $I_j$  is one when the event passes the X-ray-like selection and is not rejected by the veto. For an unweighted sample generated directly from the physical source distribution, Eq. 6.17 reduces to the familiar expression  $R_i^{\text{raw}} N_i(C_x \cap \bar{C}_{\text{veto}}) / N_i^{\text{generated}}$ . Depending on the source class, the raw normalization is derived from a material activity, a measured environmental flux, a cosmic-ray flux model, or an experimental live-time normalization.

The total predicted background rate in the region of interest is then

$$R_{\text{tot}}^{\text{after}} = \sum_i R_i^{\text{after}}, \quad (6.18)$$

with analogous sums before cuts and after X-ray-like cuts alone. This three-stage presentation is useful because it separates the physics source model from the detector selection:

1. the raw rate before analysis cuts tests the source normalization and geometry;
2. the rate after X-ray-like cuts tests whether the component can mimic a compact low-energy Micromegas event;
3. the rate after X-ray-like plus veto cuts tests the residual background after the active-shielding strategy is applied.

The final background tables are therefore organized, for each source class, around the number of simulated events, the number of analyzed events, the efficiency of the X-ray-like selection, the veto survival fraction, the raw normalized rate, and the residual rate after all cuts. This structure makes it possible to update the rate budget without changing the conceptual organization of the chapter: new simulations or improved source normalizations modify only the source-specific entries in the sum, while the definitions of the X-ray-like and veto selections remain common.

## 6.4 External Background

External contributions to the background are those that come from outside the detector, such as cosmic rays and environmental radiation. IAXO and BabyIAXO are planned at DESY in Hamburg, Germany, which is not an underground low-background laboratory. At the start of the background-model work the reference BabyIAXO scenario was the HERA South Hall, where hall structure and access geometry could have provided a site-dependent modification of the cosmic-ray field. Recent collaboration planning has shifted the reference implementation toward an outside, on-surface DESY location [50, 51]. The cosmic-ray-induced background must therefore be treated as a significant contribution to the total background rate, and a model without meaningful underground overburden is the relevant baseline.

### 6.4.1 Cosmic Rays

The main particles produced by the cosmic rays that would reach the detector are muons, gammas, protons, electrons and neutrons.

#### Cosmic Muons

Muons are the most abundant cosmic ray particle that reaches the detector. They are highly penetrating particles that can traverse the entire detector volume and produce a signal in the detector.

#### Cosmic Gammas

The spectrum of photons produced by the cosmic rays has contributions from very high energies. These very high energy photons cannot be stopped by the lead shielding and produce a cascade of secondary particles when they interact with the shielding.

#### Cosmic Protons

Protons produced by the cosmic rays can contribute to the detector background. The high energy protons behave in a similar way to the high energy neutrons, producing a cascade of secondary particles via interactions in the shielding.

#### Cosmic Electrons

Electrons and positrons produced by the cosmic rays also contribute to the background rate. Their contribution is expected to be small, but it is still necessary to consider them in the background model.

#### Cosmic Neutrons

High-energy cosmic neutrons reaching ground level constitute one of the most challenging background sources for IAXO. Unlike muons, neutrons are electrically neutral and highly penetrating. When traversing the lead shielding, high-energy neutrons undergo inelastic scattering and spallation reactions, producing complex hadronic showers of secondary particles—including lower-energy evaporation neutrons, protons, and gammas. These secondary particles can deposit energy in the Micromegas TPC in the region of interest. Because the primary incident neutrons do not produce a continuous ionization track, they easily evade standard active muon veto systems, thereby requiring a dedicated, heavily optimized multilayer veto system alongside specific passive shielding moderator materials.

### 6.4.2 Cosmic-ray event classes and veto survival

The source-by-source background rates do not fully describe how cosmic-ray events appear after reconstruction. For the background model, however, the role of the detailed event-history studies is deliberately limited: they justify the source-dependent survival factors applied after the Micromegas selection and the active-veto selection. The full detector interpretation of the veto response, including prompt muon tags, delayed neutron-sensitive observables, neutron-history categories, and the comparison with experimental IAXO-D0 veto data, is given in the shielding and veto chapter.

The relevant distinction is that cosmic-ray sources do not survive the analysis for the same physical reason. Muon-induced events usually carry a prompt, high-amplitude, multi-panel scintillator signature. Their residual contribution after the prompt veto is therefore expected to be dominated by atypical cases: inefficient regions, unstable or disabled channels, weak prompt deposits, or secondary particles produced by the muon in the surrounding materials. Neutron-induced events are less direct. The neutron-history study discussed in the shielding and veto chapter shows that the largest class of TPC-depositing neutron events is produced by electromagnetic descendants rather than by a primary neutron scattering elastically in the gas. This explains why a neutron-initiated event can look x-ray-like in the Micromegas while still leaving delayed, multiplicity-rich, or spatially diffuse veto activity in the scintillator-cadmium system.

Source class	Route to a low-energy TPC event	Veto-survival implication	Main validation
Muons	Direct ionization, bremsstrahlung, or muon-induced secondaries in the shielding.	Suppressed mainly by the prompt multi-panel veto; residuals are controlled by prompt-tag inefficiency and secondary production.	the shielding and veto chapter.
Neutrons	Hadronic cascade in lead, copper, scintillator, and cadmium; the gas deposit is often produced by a gamma or electron descendant.	Requires delayed, multiplicity, and broad veto-pattern observables; no-veto events define the irreducible tail.	the shielding and veto chapter.
Protons and hadrons	Secondary cascades similar to neutron-induced events, but with a charged primary or charged descendants.	Mixture of prompt and delayed veto logic; topology and timing are source dependent.	Veto simulations.
Gammas and electrons	Electromagnetic interactions in the shielding or detector materials.	Controlled mostly by passive shielding and Micromegas topology; veto correlations are weaker and usually indirect.	Source-specific simulations.

Table 6.9: Cosmic-ray event classes as used in the background model. The table summarizes how each source class enters the low-energy Micromegas sample and how the active-veto survival factor should be interpreted.

This separation keeps the background-model chapter focused on the final source catalogue and rate construction. The veto observables enter here as reconstructed selections and efficiencies, not as truth-level labels. Consequently, the final rates should

be read as source-normalized predictions after a common Micromegas analysis and a source-dependent veto survival factor. The same event-history diagnostics remain essential, but their proper home is the veto-system chapter, where they can be compared directly with the waveform-level veto response and with the experimental prompt and neutron-enriched control samples.

### 6.4.3 Environmental Background

Environmental background refers here to radiation that originates outside the detector from the laboratory or experimental-hall surroundings. This component must be interpreted with some care. The environmental-gamma model described below was originally developed for an enclosed laboratory configuration, where the dominant source was assumed to be natural radioactivity from nearby concrete walls, floor, and ceiling. That assumption is useful for IAXO-D0 studies and for generic shielding estimates, but it is not by itself a final site model for BabyIAXO. BabyIAXO will operate at ground level, and the relevant external field will combine cosmic-ray secondaries, radiation from the floor and surrounding structures, and any local material near the detector. The concrete-source simulations are therefore retained in this thesis as a validated methodology and as a conservative environmental-radiation component, while the final BabyIAXO normalization should be tied to site-specific measurements or to a site-specific radiation map.

The dominant non-cosmic external sources in the original model are the concrete walls, floor, and ceiling, together with other nearby construction materials. Concrete contains natural radioisotopes from the  $^{40}\text{K}$ ,  $^{232}\text{Th}$ ,  $^{235}\text{U}$ , and  $^{238}\text{U}$  families, whose decay products emit photons, electrons, and, through secondary processes, neutrons. These particles can enter the detector volume directly, or they can interact in the shielding and produce lower-energy secondaries that mimic an x-ray-like event in the gas.

The simulation strategy separates the environmental source calculation from the detector transport. Directly simulating radioactive decays throughout the full laboratory volume would be inefficient, because most decay products are absorbed before reaching the detector. Instead, the decay products escaping from representative concrete volumes are first simulated and stored as energy and angular distributions. Those distributions are then used as primary-particle inputs for the detector simulation. This two-step approach preserves the relevant spectral and angular information while avoiding the cost of transporting a very large number of particles that never leave the surrounding material.

## Environmental Gammas

Environmental gammas are modeled as photons generated by the natural radioactivity of the surrounding concrete. This should be understood as the environmental radioactivity component, distinct from the cosmic-gamma component discussed in Section 6.4.1. For BabyIAXO, the concrete-wall model should be revisited once the relevant detector position, surrounding structures, and site measurements are fixed. Until then, it provides a useful estimate of how MeV-scale laboratory photons propagate through the shielding and which detector-entry paths remain open. The relevant parent chains are simulated independently for  $^{40}\text{K}$ ,  $^{232}\text{Th}$ ,  $^{235}\text{U}$ , and  $^{238}\text{U}$ . For each isotope, radioactive decays are generated uniformly in a concrete slab of thickness 1 m and effectively infinite transverse size. The particles are transported until they either stop in the material or cross the surface facing the detector. For each escaping particle, the simulation records its type, kinetic energy, and exit angle with respect to the surface normal. The photon component of these distributions defines the environmental-gamma source term used in the detector simulation.

Two detector-level implementations are kept in the simulation repository. The first is a compact spherical source in which gamma primaries are launched from a 5 m radius surface surrounding the detector. Their energies are sampled from the `EnvironmentalGammas` distribution and their incidence angle follows the analytical  $\sin(2\theta)$  form expected for particles escaping an extended medium. This configuration is useful for fast rate estimates and for studying the detector response to an approximately isotropic laboratory gamma field. The second implementation is a wall-source model. In that case, photons are generated from planar surfaces corresponding to the laboratory walls and their energy and zenith-angle distributions are read from isotope-specific histograms, `U238_CONCRETE_1000mm.root`, `Th232_CONCRETE_1000mm.root`, `K40_CONCRETE_1000mm.root`, and `U235_CONCRETE_1000mm.root`. This model preserves directional information and allows the individual isotope contributions to be combined with different assumed activities.

The absolute normalization of this component cannot be obtained from the transport simulation alone. It must be derived either from measured gamma fluxes in the laboratory or from the activity concentrations assumed for the surrounding concrete. The simulations therefore provide the detector acceptance per generated environmental gamma, while the final background contribution is obtained by multiplying this acceptance by the measured or assumed laboratory gamma flux. For BabyIAXO, where the detector orientation changes during solar tracking, the final rate should use an orientation-averaged flux or a conservative envelope over the relevant pointing range. The current production inventory includes the environmental-gamma configurations, but the detector-level environmental-gamma samples still need to be produced with the same reconstruction chain used for the cosmic and intrinsic background samples

before a final numerical rate can be quoted.

### Concrete Decay Simulation




Figure 6.11: Detector-level implementation of the environmental-background source. Particles are launched from an enclosing surface around the detector. Their energies are sampled from the environmental source spectrum and their incidence angles are sampled either from the  $\sin(2\theta)$  approximation or from the concrete-emission histograms obtained in the preliminary decay simulation.

The decay chains of the radioisotopes present in concrete are long and generate many secondary particles. Only a small fraction of these secondaries escape the material, and an even smaller fraction can reach the detector after passing through the passive shielding. For this reason, an auxiliary `Geant4` application was used to simulate the decay chains in a simplified concrete geometry before running the full detector model. Each isotope was generated uniformly in depth and the exiting photons and electrons were histogrammed in kinetic energy and exit angle. The resulting histograms define a source term that can be reused by the main `restG4` detector simulations without repeatedly simulating the full decay cascade inside the concrete.




Figure 6.12: Auxiliary decay simulation used to construct the environmental source term. Radioisotopes are sampled uniformly in depth inside a concrete volume with effectively infinite transverse size. Particles that escape through the surface are recorded with their particle type, production depth, kinetic energy, and exit angle with respect to the surface normal. These distributions are then used as inputs to the detector-level environmental simulations.

The auxiliary simulation tracks photons, electrons, positrons, and neutrons produced by the radioactive decay chains. In practice, photons and electrons dominate the escaping population. The radiogenic neutron component is treated separately because accurate production by  $(\alpha, n)$  reactions and spontaneous fission requires additional nuclear data and validation. For the present background model, the concrete-decay histograms are therefore used for the environmental gamma and electron source terms, while environmental neutrons are generated from literature-based spectra as described below.




Figure 6.13: Photons and electrons produced by the decay of  $^{238}\text{U}$  in a concrete block of 1 m depth.




Figure 6.14: Exit angle for photons and electrons produced by the decay of  $^{238}\text{U}$  in a concrete block of 1 m depth. The angle is defined as the angle between the normal to the surface and the direction of the particle.

Besides the energy and type of the exiting particles, the angle of the particles with respect to the normal to the surface is also recorded. The angular distribution is also important, as it will affect the number of particles that reach the detector. The

secondary particles inside the medium are produced isotropically, but due to geometrical effects, the exit-angle distribution is not isotropic. Instead, it approximately follows a  $\sin(2\theta)$  law, as shown in Figure 6.14. The secondary particles produced, especially photons for which the statistics are highest, do not follow the  $\sin(2\theta)$  form exactly because particles with larger exit angles are more strongly attenuated, producing a distribution slightly skewed toward lower angles. For all radioisotopes and materials considered, the angular distribution of secondaries follows the  $\sin(2\theta)$  expectation closely enough that the analytical form would already be a reasonable approximation. Even so, the full simulated distributions were used in order to preserve the best available accuracy. The corresponding spectra for the other concrete radioisotopes and the decay-chain diagrams are kept in Appendix 6.14, since they document the source construction but do not change the detector-level argument.

Measurements of the environmental gamma radiation are needed to normalize the detector simulation. Dedicated measurements at the detector site provide the most direct normalization; measurements with a NaI detector are also useful because the prominent lines from natural radioactivity constrain the relative contributions of the main decay chains.




Figure 6.15: Background spectrum measured in the Zaragoza laboratory with a NaI detector. The range between 0.25 MeV and 2.75 MeV was selected to compute the background rate. The spectrum contains the continuum and gamma-line structure expected from naturally occurring radioactivity in the laboratory environment.




Figure 6.16: Monte Carlo simulation of a NaI detector with same geometry as the one used in the measurements placed in the center of a sphere with a diameter of 5 m. Photons energies are sampled from a uniform distribution, labeled as “Input” in the legend. For reference, simulations labeled as “Frontal” and “Lateral” show the detector response when photons are launched directly toward the center of the circular face or perpendicular to the cylinder axis and passing through the center of the cylinder, respectively. The simulation using the sphere is rescaled by a factor of 10000 to improve visibility. A difference is observed between the detector spectrum obtained with the sphere simulation and those from the frontal/lateral simulations because of geometrical effects. For the sphere simulation, the photons were launched with an angular distribution following  $\sin(2\theta)$ , which is the expected distribution for photons produced by the decay of natural radioisotopes in a material, as shown by the previous simulations.




Figure 6.17: The image on the left shows the input energy distribution used in the simulation of the NaI detector. The angular distribution used is  $\sin(2\theta)$ . The image on the right shows the comparison between the energy spectrum measured in the laboratory with the NaI detector and the one measured in the **Geant4** simulation. While the peak structure is not reproduced, the overall shape of the spectrum is similar.




Figure 6.18: Sample events produced by the environmental gamma background simulation. On the left image, a 1.25 MeV photon enters through the shielding opening and interacts with the vessel via Compton scattering, then enters the gas volume and ionizes the gas via photoelectric effect. On the right image, a 0.9 MeV photon interacts with the copper pipe outside of the shielding and produces a Bremsstrahlung electron, which then enters through the vacuum line and cathode window, directly ionizing the gas volume.

### Environmental Neutrons

Radioactive decays in the materials surrounding the detector also produce neutrons via  $(\alpha, n)$  reactions or spontaneous fission, such as in the case of  $^{238}\text{U}$ . These neutrons, with energies in the MeV range, can reach the detector and produce a signal. Since the detector and the materials close to it are made mostly of copper with very high

radiopurity, the number of radiogenic neutrons produced locally is expected to be very low. However, because neutrons can be very penetrating, the neutrons produced by radioactive decays in the surrounding materials must still be considered.

Initially, an approach similar to the one used for the gamma background was followed, in which the production of neutrons by the natural radioisotopes in the concrete was simulated. The simulation was performed using **Geant4**, using an identical methodology to the one used for the gamma background. This approach was later abandoned due to issues with the accurate simulation of the neutron production via  $(\alpha, n)$  reactions and spontaneous fission.

Instead, a rough approximation of the neutron flux was adopted from the literature [128], which provides the neutron flux produced by the natural radioisotopes in concrete. The energy distribution roughly follows an exponential decay law.

A generic neutron evaporation spectrum, centered around 1 MeV, was used as a first-order approximation for the primary-neutron generation in these simulations.

A dedicated campaign of measurements of the neutron rate and spectrum in the laboratory is planned for the future and will provide a more accurate estimate of the neutron background rate in the detector.




Figure 6.19: Primary neutron energy distribution used in the environmental neutron background simulation (red, labeled as “input”), approximated by an exponential decay law. Also showing is the primary energy distribution of the neutrons that deposited energy above the readout in the simulation, showing that higher energy neutrons have a higher probability of depositing energy in the detector (blue, labeled as “simulation”).




Figure 6.20: Sample events produced by the environmental neutron background simulation. On the top image, a 5.89 MeV neutron penetrates the shielding and interacts with the copper vessel producing a photon via  $(n, \gamma)$  process. The photon then interacts via Compton scattering near the gas volume and the resulting electron ionizes the gas. On the bottom image, a 4.01 MeV neutron interacts with the lead shielding and the shielding inner copper wall via elastic scattering, reducing its energy and is then captured by the copper, producing a photon which in turn produces a positron which ionizes the gas volume.




Figure 6.21: Detector background for environmental neutrons as a function of the deposited energy in the detector. The copper K- $\alpha$  line can be seen at 8.04 keV. The background level in the energy range of interest of (0.1 - 10) keV is  $1.51 \times 10^{-6} \text{ keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$ .

## 6.5 Intrinsic Background

### 6.5.1 Gas Contamination

The IAXO Micromegas detector is designed to work with an argon-based gas mixture or a xenon-neon-based gas mixture. Given that the gas mixture is the medium where the interactions take place, it is crucial to have a good understanding of the isotopic composition of the gas and the possible contaminants that could be present. The presence of common gas contaminants, such as oxygen and water, has an effect on the detector performance. These contaminants, due to their low concentration and natural isotopic composition, do not have a measurable effect on the background rate. Radon gas, which inevitably leaks into the detector, is a more serious concern and will be described in a separate section.

Of all the possible radioisotopes that could be present in the gas, the most relevant for the IAXO Micromegas detector is  $^{39}\text{Ar}$ , which will only be present in the argon-based gas mixture.  $^{39}\text{Ar}$  is a beta emitter with  $Q_\beta = 565 \text{ keV}$  and a half-life of  $T_{1/2} = 269 \text{ y}$ .

Atmospheric argon, which is the most common source of argon, has a  $^{39}\text{Ar}$  activity of approximately  $1.0 \text{ Bq/kg}^3$  [129]. This means that the activity of  $^{39}\text{Ar}$  in the detector chamber will be approximately  $5 \times 10^{-6} \text{ Bq}$  or one decay every 55 days on average.




Figure 6.22: Ionizations produced by the decay of  $^{39}\text{Ar}$  in the gas volume of the detector for 15 different events. Each event is represented by a different color. Data obtained from `Geant4` simulation on a gas mixture of argon with 1% isobutane at a pressure of 1.4 bar.

Figure 6.22 shows some examples of the ionizations produced by the decay of  $^{39}\text{Ar}$  in the gas volume of the detector. The electrons produced by the decay of  $^{39}\text{Ar}$  leave a usually long track in the gas volume, meaning that the topological cuts used to reject the background events will be very effective in rejecting these events.

The accepted contribution of this component is therefore expected to be small after the X-ray-like topology selection, but the final normalized rate is kept as a separate catalogue entry because it depends on the chosen gas inventory and on whether atmospheric or underground argon is assumed. In the current background-model status table it is treated as a gas-intrinsic source study rather than as a closed rate contribution.

## 6.5.2 Radon Contamination

Radon is a radioactive noble gas produced in the decay chains of uranium and thorium. It is one of the most studied sources of background in low background experiments, as it can easily leak into the detector and it is very difficult to completely eliminate it. Because it is a gas, radon can escape from materials and surroundings, effectively breaking the secular equilibrium of the chain. As a result, its concentration cannot be reliably inferred from the activity of its parent isotopes.

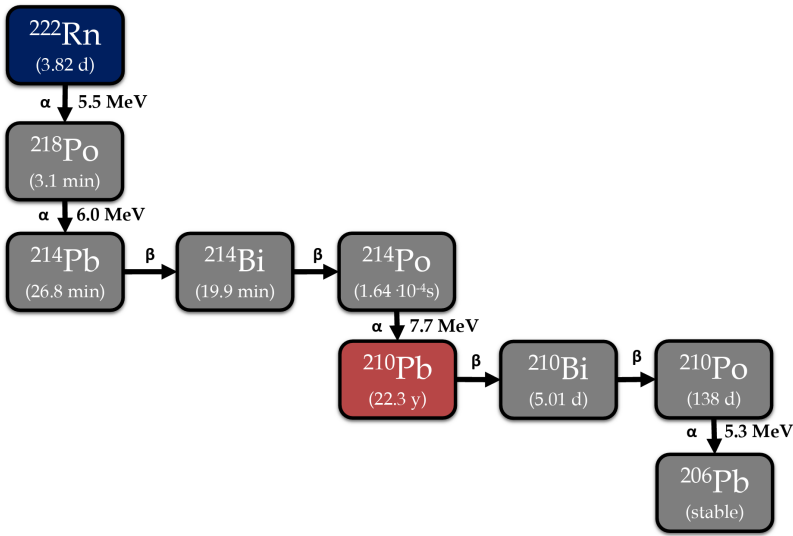


Figure 6.23: Decay chain of  $^{222}\text{Rn}$ . Image from [130].

Out of all the radon isotopes,  $^{222}\text{Rn}$  is the most abundant and long-lived isotope, with a half-life of  $T_{1/2} = 3.82$  days. In the context of the IAXO micromegas detector,  $^{222}\text{Rn}$  can produce background events in different ways. It is present in the gas volume, acting as a uniformly distributed source of background, similar to  $^{39}\text{Ar}$  discussed in the previous section, with the difference that  $^{222}\text{Rn}$  is a  $\alpha$  emitter as seen in its decay chain in Figure 6.23.

The progeny of  $^{222}\text{Rn}$ , in particular  $^{210}\text{Pb}$  due to its long half-life of  $T_{1/2} = 22.3$  years, can accumulate on the surfaces of the detector and produce background events when they decay.  $^{210}\text{Pb}$  produced by the decay of  $^{222}\text{Rn}$  is a  $\beta$  emitter and is treated as a distinct source of background, since its activity cannot be directly inferred from the activity of  $^{222}\text{Rn}$ .  $^{210}\text{Pb}$  is present in every surface of the detector, but its more significant contribution to the background rate is expected to come from the surfaces that are in contact with the gas volume.

Another relevant isotope in the decay chain of  $^{222}\text{Rn}$  is  $^{218}\text{Po}$ , which is a  $\alpha$  emitter with a half-life of  $T_{1/2} = 3.1$  minutes. This isotope is produced uniformly in the gas volume, as it is the first decay product of  $^{222}\text{Rn}$ . Because it is positively charged after the decay, it is affected by the drift field and tends to accumulate in the cathode window. This contamination will also be treated as a distinct source of background. Due to this effect, the concentration of  $^{210}\text{Pb}$  in the cathode will be different from that on the other surfaces; therefore, it is also treated as a distinct source of background.




Figure 6.24: Ionizations produced by the decay of  $^{222}\text{Rn}$  in the gas volume of the detector for 15 different events. Each event is represented by a different color. Data obtained from **Geant4** simulation on a gas mixture of argon with 1% isobutane at a pressure of 1.4 bar.




Figure 6.25: Ionizations produced by the decay chain of  $^{210}\text{Pb}$  in the internal surface of the detector vessel for 15 different events. Each event is represented by a different color. Data obtained from **Geant4** simulation on a gas mixture of argon with 1% isobutane at a pressure of 1.4 bar.




Figure 6.26: Ionizations produced by the decay chain of  $^{218}\text{Po}$  in the internal surface of the negatively charged cathode for 15 different events. In the simulation, the decay chain of  $^{218}\text{Po}$  was simulated only up to  $^{210}\text{Pb}$ , which was then simulated separately. Each event is represented by a different color. Data obtained from `Geant4` simulation on a gas mixture of argon with 1% isobutane at a pressure of 1.4 bar.




Figure 6.27: Ionizations produced by the decay chain of  $^{210}\text{Pb}$  in the internal surface of the negatively charged cathode for 15 different events. Each event is represented by a different color. Data obtained from `Geant4` simulation on a gas mixture of argon with 1% isobutane at a pressure of 1.4 bar.

The radon-related figures in this section are diagnostic source studies rather than final rate predictions. A quantitative radon contribution requires the gas activity, flushing conditions, surface exposure history, and plate-out model to be fixed for the detector configuration being normalized. For this reason, the radon and surface-progeny entries are kept as explicit upper-bound or scenario components in the source catalogue.

### 6.5.3 Detector Materials

The materials used for the construction of the detector are carefully selected to minimize the background rate, giving closer attention to the materials that are closer to the gas volume. Here is a list of the most relevant detector materials from the point of view of the background model:

- **Copper:** It is the main material used in the construction of the detector. It is also present in the readout electronics and in the detector readout plane.
- **Lead:** The shielding is made out of lead, which is used to shield the detector from the environmental radiation. General purpose lead is less radiopure than the copper used in the detector.
- **Stainless Steel:** It is the material of choice for most of the structural components of the detector and platform. It is not considered radiopure so the detector will be designed to minimize the amount of stainless steel in the detector. The beamline connecting the detector to the magnet is made out of stainless steel and is the most significant contribution to the background from this material.
- **Kapton:** A polymer used in the construction of the Micromegas readout plane. It can be produced with a very low level of radioactivity and its radiopurity is critical as it is part of the detector readout plane.
- **PTFE:** Also known as Teflon, it is used in the construction of the detector. It is a polymer that can be produced with a very low level of radioactivity and is easily machined. It is used for its insulating properties and low outgassing. It is treated explicitly in the material catalogue because some PTFE elements are close to the sensitive gas volume.
- **Aluminized Mylar:** A very thin layer of aluminum deposited on a Mylar film is used as the cathode of the detector. It must be as thin as possible to minimize the attenuation of the conversion photons while resisting the gas pressure. Due to its location above the central region of the detector where the axion signal is expected to be produced, it is critical to have a very low level of radioactivity.

- **Front-end electronics materials:** The front-end electronics must be very close to the detector in order to minimize signal loss. This PCB is manufactured using low-background materials, but it is not possible to control all the materials used in the PCB. For the purposes of the background model, the front-end electronics are treated as a single material.

All of these materials have a very well understood radiopurity profile, as they are commonly used in low background experiments and were also used in the CAST experiment. The radioactivity of these materials is measured using high purity germanium (HPGe) detectors, which have a very good energy resolution and can measure the activity of the most common radioisotopes present in the materials. Activity measurements were performed by the Zaragoza detector group in the facilities of the Canfranc Underground Laboratory (LSC) in Spain with samples of the materials used in the construction of the detector. Some other materials not mentioned in this list were also measured, such as epoxy resins and glues used in the construction of the detector.

### Copper Parts

Electroformed copper has an extremely high purity (over 99.9999%) but still contains some contaminants. The most relevant copper contaminants from the point of view of the background model are  $^{60}\text{Co}$ ,  $^{40}\text{K}$ ,  $^{238}\text{U}$  and  $^{232}\text{Th}$ . The most relevant contaminant of these isotopes is  $^{60}\text{Co}$ , which is produced by cosmic ray activation of the copper. The concentration of  $^{60}\text{Co}$  in the copper depends on the cosmic ray flux and the time of exposure to the cosmic rays. Because IAXO will not be located in an underground laboratory, the copper will be continuously exposed to the cosmic rays and the concentration of  $^{60}\text{Co}$  will be in a state of equilibrium with the cosmic ray flux, resulting in a much higher concentration than in the case of underground experiments, where the activity of  $^{60}\text{Co}$  reduces over time due to the relatively short half-life of  $^{60}\text{Co}$  ( $T_{1/2} = 5.27$  years) and the low level of activation of the copper.




Figure 6.28: Energy of photons produced by the decay of different isotopes present in the copper parts. Only showing those photons that exit the volume with the energy at which they exit the volume.




Figure 6.29: Energy of electrons produced by the decay of different isotopes present in the copper parts. Only showing those electrons that exit the volume with the energy at which they exit the volume.




Figure 6.30: Depth of parent isotope for photons produced by the decay of different isotopes present in the copper parts. Only showing those photons that exit the volume. The depth is defined as the distance from the surface to the parent isotope.




Figure 6.31: Depth of parent isotope for electrons produced by the decay of different isotopes present in the copper parts. Only showing those electrons that exit the volume. The depth is defined as the distance from the surface to the parent isotope.

For practical purposes, copper can be treated as a source of photons and electrons produced by the decay of the isotopes present in the material. By performing an

initial simulation of these decays, the energy distributions of the photons and electrons emanating from the copper can be obtained, as shown in figures 6.28 and 6.29. These energies are of the order of 1 MeV.

Figures 6.30 and 6.31 show the depth of the parent isotope for the photons and electrons produced by the decay of the isotopes present in the copper. Most of the photons and electrons are produced by decays close to the surface, although a significant number of photons are still produced by decays at greater depths. For the simulations, the decays are assumed to be uniformly distributed throughout the copper volume and are sampled over the full volume of the copper parts. The detailed depth and escaping-particle spectra are retained as validation plots for the transport model; the final source ranking should be quoted from the normalized survivor-rate table rather than from these intermediate distributions.

#### 6.5.4 Front-end Electronics

The front-end electronics constitute a special intrinsic-background source because they must be placed close to the Micromegas readout in order to preserve signal quality, while at the same time containing materials that are harder to radiopurify than bulk copper, kapton, or PTFE. This contribution was therefore treated separately from the generic copper and readout-plane contaminations. The simulation described here corresponds to the four front-end cards represented in the IAXO-D1 geometry, without including the flat cables or other service elements. The activity model follows the component-level radiopurity study of the BabyIAXO electronics, used here as a collaboration input rather than as an original measurement of this thesis [54, 122].

The simulated source volumes were the four electronics-card volumes in the `IAXO-D1/Default.gdml` geometry. For each card and each isotope, radioactive decays were generated uniformly inside the corresponding card volume and transported with the same `Geant4` physics configuration used for the rest of the intrinsic-background model. The isotopes included in the present production were  $^{40}\text{K}$ ,  $^{60}\text{Co}$ ,  $^{137}\text{Cs}$ , the  $^{232}\text{Th}$  chain, the  $^{235}\text{U}$  chain, and the  $^{238}\text{U}$  chain. The activities assigned to the four cards are summarized in Table 6.10. The optional contribution from the 10  $\Omega$  resistors was not included in the nominal normalization because those components were not part of the original low-background board design considered in the reference electronics study.

Isotope or chain	Activity for four cards [Bq]	Activity per card [Bq]
$^{40}\text{K}$	$1.30044 \times 10^{-1}$	$3.25110 \times 10^{-2}$
$^{60}\text{Co}$	$1.03702 \times 10^{-2}$	$2.59255 \times 10^{-3}$
$^{137}\text{Cs}$	$1.54240 \times 10^{-3}$	$3.85600 \times 10^{-4}$
$^{232}\text{Th}$ chain	$2.56600 \times 10^{-1}$	$6.41500 \times 10^{-2}$
$^{235}\text{U}$ chain	$2.04800 \times 10^{-3}$	$5.12000 \times 10^{-4}$
$^{238}\text{U}$ chain	2.41443	$6.03608 \times 10^{-1}$

Table 6.10: Nominal activity vector used for the four electronics-card source volumes. The values exclude the optional  $10\ \Omega$  resistors and the flat cables.

The normalization was performed at the level of generated decays. For each isotope–card pair, the event weight was defined as

$$w_{i,c} = \frac{A_{i,c}}{N_{i,c}^{\text{gen}}},$$

where  $A_{i,c}$  is the activity assigned to isotope  $i$  in card  $c$ , and  $N_{i,c}^{\text{gen}}$  is the number of generated decays in that production. This choice is important because only a small fraction of the decays produce a stored detector event, and normalizing to saved or reconstructed entries would bias the rate estimate. The production consisted of 24 isotope–card groups, corresponding to the six isotopes or decay chains in Table 6.10 for each of the four cards. In total,  $8.48 \times 10^9$  decays were generated and  $2.71 \times 10^4$  detector events were saved before the flat analysis export. The  $^{235}\text{U}$  chain is the least efficient source in this geometry, and several  $^{235}\text{U}$  jobs produced no saved detector event; those jobs were nevertheless retained in the generated-event denominator.

The resulting spectrum in the 0–12 keV region is shown in Figure 6.32. The rates are expressed in the background units used throughout this chapter,  $\text{counts s}^{-1} \text{keV}^{-1} \text{cm}^{-2}$ , using the circular signal fiducial area  $r < 1$  cm for the area normalization. Before event-selection cuts, the electronics-card contribution in the simulated 0–12 keV range is  $5.09 \times 10^{-7} \text{ counts s}^{-1} \text{keV}^{-1} \text{cm}^{-2}$ . Within the plotted 0–12 keV window, the fiducial and x-ray preselection requirements reduce the contribution to  $3.93 \times 10^{-8} \text{ counts s}^{-1} \text{keV}^{-1} \text{cm}^{-2}$ . After the energy-binned x-ray-like selection, the residual electronics-card rate is  $3.32 \times 10^{-10} \text{ counts s}^{-1} \text{keV}^{-1} \text{cm}^{-2}$ . The simple veto cut does not further reduce this internal component, as expected, because the source is located inside the shielded detector volume and does not generally produce a correlated external-veto signature.




Figure 6.32: Activity-normalized 0–12 keV fiducial spectra for the four electronics cards, summed over all source volumes. Left: background level after fiducial preselection. Right: residual background level after the energy-binned x-ray-like selection and veto cut. Each bin is normalized by the isotope activity, the number of generated decays, the bin width, and the circular signal fiducial area  $r < 1$  cm. Error bars show the propagated Poisson uncertainty from the finite number of simulated events in each weighted bin.

The residual rate is dominated by the  $^{238}\text{U}$  chain. In the current high-statistics production,  $^{238}\text{U}$  accounts for  $3.56 \times 10^{-9}$  counts  $\text{s}^{-1} \text{cm}^{-2}$  integrated over the 0–12 keV window after x-ray-like cuts, corresponding to nearly 90% of the surviving electronics-card contribution. This behavior is consistent with the component-level electronics study, in which ceramic capacitors were identified as the dominant  $^{238}\text{U}$ -driven contribution to the electronics background. The second largest post-cut contribution is the  $^{232}\text{Th}$  chain, while  $^{40}\text{K}$ ,  $^{60}\text{Co}$ ,  $^{137}\text{Cs}$ , and  $^{235}\text{U}$  are subdominant after the x-ray-like topology cuts.




Figure 6.33: Left: mean 0–12 keV electronics-card background level per isotope after fiducial preselection and after the full x-ray-like selection, summed over all four cards. Right: simulation campaign statistics per isotope, including the total generated primaries, the number of saved `EventTree` entries, and the activity-normalized residual background level with propagated Poisson uncertainty. <sup>238</sup>U has the highest activity (2.41 Bq for four cards) and, although it produces fewer saved events than <sup>60</sup>Co, its larger per-decay weight makes it the dominant contributor to the residual background rate after cuts.




Figure 6.34: A  $^{238}\text{U}$  decay event from electronics card 2 that survives the full x-ray-like selection and veto cuts, displayed with the **REST-for-Physics Geant4** event viewer. The main view shows the detector chamber (copper vessel, gas volume, and readout plane) with particle tracks from the radioactive decay, while the inset panels show the compact veto projections. This event deposited approximately 3.7 keV in the fiducial volume and is one of two  $^{238}\text{U}$  events that passed all topological cuts in the combined 24-isotope-card high-statistics production.

These results should be interpreted as the card-only electronics contribution for the current geometry and activity vector. They do not yet include flat cables, possible connector materials, or a detailed separation of the individual board components within each card volume. Those refinements are relevant for the final BabyIAXO background model because the activity is not uniformly distributed among components: the  $^{238}\text{U}$  activity is concentrated mainly in ceramic capacitors, while resistors, diodes, and chips contribute different isotope mixtures. Nevertheless, the card-only production provides a useful first estimate of the scale of the electronics background and confirms that, under the present activity assumptions and x-ray-like selection, the front-end cards are not expected to dominate the final 1–10 keV background budget.

### 6.5.5 Shielding

The lead shielding sits around the detector and is the most massive component of the detector. The shielding is made out of lead, but it can also present some contamination from other materials. Most of the emissions of these contaminants are absorbed by the lead, but those produced in the region of shielding that is closer to the detector can reach the detector and contribute to the background rate. Neutrons produced in the lead by nuclear fission of  $^{238}\text{U}$  (a common contaminant) or  $(\alpha, n)$  reactions can reach the detector no matter where they are produced in the shielding, due to their high penetration power.

The contaminants considered for the lead shielding are:  $^{40}\text{K}$ ,  $^{232}\text{Th}$ ,  $^{238}\text{U}$ ,  $^{235}\text{U}$  and  $^{210}\text{Pb}$ . Out of these isotopes,  $^{210}\text{Pb}$  is the one with the most specific activity and a relatively short half-life of  $T_{1/2} = 22.3$  years. This means that the activity of  $^{210}\text{Pb}$  in lead is reduced over time, meaning that it is possible to obtain a lead shielding with a very low level of radioactivity by using lead that has been stored for a long time, preferably underground to reduce the cosmic ray activation of the lead. A particularly interesting type of lead is archeological lead that was extracted from shipwrecks, as it has been stored underwater (which shields it from cosmic rays) for a long time and has a very low level of radioactivity. This low-activity lead is very difficult to obtain and thus can only be used in the inner layers of the shielding, where the lower level of radioactivity is the most important.




Figure 6.35: Energy of photons and electrons produced by the decay of  $^{210}\text{Pb}$  in lead. Only showing those that exit the volume. The number of photons and electrons is similar, in contrast with the other simulations done in concrete, where there is a much bigger number of photons than electrons.




Figure 6.36: Geometry of the detector chamber and lead shielding used for the simulations. The region highlighted in red corresponding to the inner surface of the lead shielding and is the primary generation volume for the simulations. The inner region of the lead shielding is separated from the detector by a layer of copper, shown in the figure.

The contaminants in the lead shielding are expected to be uniformly distributed throughout the shield volume, but only the contaminants on the inner surface are simulated, since they are the only ones that can reach the detector in significant quantities after the shielding effect of the lead is taken into account. Figure 6.36 shows the geometry of the lead shielding used for the simulations. Highlighted in red is the region of the inner surface of the lead shielding, which is the primary generation volume for the simulations.




Figure 6.37: Angle with respect to the surface normal of photons and electrons produced by the decay of  $^{210}\text{Pb}$  in lead. Only showing those that exit the volume.




Figure 6.38: Depth of parent isotope for photons and electrons produced by the decay of  $^{210}\text{Pb}$  in lead. Only showing those that exit the volume. As expected, the closer to the surface the parent isotope is, the more likely it is to produce a photon or electron that exits the volume. For electrons, the dependence on depth is much more pronounced than for photons. This information can be used to estimate an upper limit for the depth of the primary-generation volume shown in Figure 6.36.




Figure 6.39: Energy of photons produced by the decay of different isotopes present in the lead shielding. Only showing those photons that exit the volume with the energy at which they exit the volume.




Figure 6.40: Energy of electrons produced by the decay of different isotopes present in the lead shielding. Only showing those electrons that exit the volume with the energy at which they exit the volume.




Figure 6.41: Depth of parent isotope for photons produced by the decay of different isotopes present in the lead shielding. Only showing those photons that exit the volume. The depth is defined as the distance from the surface to the parent isotope.




Figure 6.42: Depth of parent isotope for electrons produced by the decay of different isotopes present in the lead shielding. Only showing those electrons that exit the volume. The depth is defined as the distance from the surface to the parent isotope.

The lead-shielding activity is not quoted as a final numerical rate in this version of the chapter. The source-depth and escaping-particle distributions above document the

simulation strategy and show why the inner lead layers and  $^{210}\text{Pb}$  contamination are the physically relevant cases. The corresponding rate contribution should be inserted only once the production campaign has been normalized with the final shielding geometry, activity model, exposure, and common 2–7 keV selection.

### Neutrons from the fission of $^{238}\text{U}$




Figure 6.43: Energy spectrum of neutrons from spontaneous fission of  $^{238}\text{U}$  [128].

Neutrons can be produced by the fission of  $^{238}\text{U}$ , which is a common contaminant in the lead shielding. The fission of  $^{238}\text{U}$  produces neutrons with an energy spectrum as shown in figure 6.43. In a similar way to the environmental neutrons, these neutrons can reach the detector and produce background events.

Spontaneous fission of  $^{238}\text{U}$  is a very rare process. The number of fission neutrons produced per  $^{238}\text{U}$  decay is  $1.31 \times 10^{-6} \text{ s}^{-1}\text{Bq}^{-1}$ , obtained from the previous IAXO background model [48]. Using the shielding mass and  $^{238}\text{U}$  activity adopted for that model, the corresponding neutron activity is small compared with the externally generated cosmic-neutron source. This results in a total fission neutron activity of  $9.06 \times 10^{-7} \text{ s}^{-1}$  in the shielding.

## 6.6 Background-model status and uncertainty roadmap

The background model developed in this chapter should be interpreted as a framework for a quantitative source-by-source prediction rather than as a single closed numerical

result. The simulations share a common detector-response and reconstruction chain, but the maturity of the input normalizations is not the same for all sources. This distinction is important because the nominal BabyIAXO background target is quoted for a detector region of interest, while different validation studies use different analysis windows. In this thesis, the broad detector-modeling range is the low-energy X-ray region, typically 0.1–10 keV, whereas the experimental veto validation discussed in the shielding and veto chapter uses the 2–7 keV window of the IAXO-D0 surface analysis. Whenever rates are compared, the applied energy window and selection chain must therefore be stated explicitly.

Source	Input normalization	Geometry / setup	Analysis treatment	Dominant uncertainty	Status
Cosmic muons	Guan sea-level CosmicMuons model, with CRY cross-checks	Surface detector with shielding and veto	Micromegas cuts plus prompt veto	Veto threshold, live time, angular spectrum	Good for design studies
Cosmic neutrons	HENSA outdoor 10 GeV spectrum, with CRY/EXPACS cross-checks	Surface detector with lead and veto volumes	Micromegas cuts plus delayed/multiplicity-rich veto observables	DESY site dependence, angular model, hadronic modeling	Nominal source selected
Cosmic protons	CRY surface source term	Surface detector with shielding and veto	Same detector-response and veto chain as neutron studies	Source model and limited statistics	Preliminary
Environmental gammas	Laboratory spectra and concrete simulations	Room/laboratory model around detector	Transport to detector and Micromegas selection	Site dependence and material composition	Needs DESY update
Environmental neutrons	Literature/model inputs and local measurements	Laboratory neutron field around shielding	Transport through shielding and detector response	Source normalization and moderation model	Needs final normalization
Gas and radon	Gas composition, radon assumptions, screening inputs	Active gas volume and detector chamber	Internal-source simulation with Micromegas cuts	Concentration, plate-out, emanation history	Partially complete
Materials, electronics, and shielding	Radiopurity screening and component masses	Source-specific detector volumes	Source-specific Geant4 transport plus common reconstruction	Screening limits, geometry details, surface-depth approximation	In progress

Table 6.11: Current status of the source classes entering the BabyIAXO/IAXO-D0 background model. The table is intentionally a status table rather than a final rate budget, because several source normalizations and detector-response parameters remain open inputs.

Uncertainty	Current treatment	Next action
Cosmic-neutron normalization	The outdoor HENSA 10 GeV spectrum is used as the nominal measured source term, with CRY and EXPACS retained as cross-checks.	Propagate the HENSA-based source term to the residual neutron rate and scan the DESY site-transfer uncertainty.
Zaragoza/DESY site dependence	Zaragoza and DESY latitudes are considered in the cosmic-source setup, but not all final DESY boundary conditions are fixed.	Produce a DESY-specific source term once the site configuration is frozen.
Geant4 hadronic modeling	High-precision neutron and binary-cascade models are used consistently across source classes.	Compare key neutron observables across relevant physics-list choices.
Quenching and light attenuation	Included in the waveform-level veto simulation chain with nominal, configurable parameters.	Scan the quenching and attenuation parameters and calibrate them against channel-by-channel prototype data.
PMT gain, thresholds, and timing alignment	Treated approximately in the design-level waveform simulation.	Build a prototype-matched branch using measured gains, thresholds, attenuation, and timing offsets.
Geometry gaps and readout mapping	Final simulations use the optimized 59-panel concept, while the prototype records up to 57 veto signals.	Keep separate prototype-matched and final-design simulation branches.
Accidental vetoes and dead time	Identified as a requirement for the final online veto definition.	Estimate accidental rates using measured single-channel rates and realistic coincidence windows.
Finite Monte Carlo statistics	Production campaigns are sized according to source importance and computational cost.	Quote statistical uncertainties for the final source-by-source rate table.

Table 6.12: Roadmap of the dominant uncertainties that must be closed before the background model can be treated as a final absolute prediction for BabyIAXO.

The most important missing element is therefore a final master rate table. For each source, that table should state the input normalization, geometry version, exposure or number of primaries, analysis window, cuts, surviving rate, statistical uncertainty, dominant systematic uncertainty, and maturity level. The framework presented here defines how such a table should be produced; the remaining work is to replace first-pass assumptions with the final measured normalizations and validated detector-response parameters.



# Summary and Conclusions

This thesis has developed simulation infrastructure and background-modeling tools for the Micromegas detector line of IAXO-D0 and BabyIAXO. The work is motivated by the operating conditions of BabyIAXO as a surface-level helioscope, where a small keV X-ray signal must be identified above detector, environmental, and cosmic-ray-induced backgrounds. The central result is not a single background number, but a detector-response-level framework that connects Monte Carlo truth information, reconstructed Micromegas observables, veto observables, and experimental validation data.

## Software and simulation infrastructure

The software contribution of the thesis is centered on the use and development of `REST-for-Physics`, its `Geant4` interface `restG4`, and the associated production workflows required for large background simulations. The relevant achievement is the construction of a practical workflow in which source-specific `Geant4` simulations can be processed through common detector-response and reconstruction chains. This makes it possible to compare different background components using the same reconstructed observables rather than relying only on idealized deposited-energy quantities.

A specific contribution of this infrastructure is the geometry-generation workflow used to build the GDML descriptions of the IAXO-D0 detector, shielding, and veto configurations. The use of high-level, version-controlled geometry definitions made it possible to compare passive-shield scans, veto layer variants, tilted configurations, and final 59-panel layouts without losing traceability of materials, volume names, or channel mappings. This was essential for the later background-model and veto-system results: the simulations depended not only on the `Geant4` physics list, but also on the ability to regenerate and identify exactly which detector geometry had been transported.

The thesis also describes the computing and production logic needed to make these simulations feasible. This includes event-type transformations, source-generation strategies, cosmic-ray injection, track pruning, and high-throughput production with reproducible configuration files. Together with the geometry-generation workflow, these

developments provide the software basis for the background-model and veto-system studies discussed in the later chapters.

## **Cosmic-ray source modeling**

Surface operation makes cosmic-ray-induced backgrounds a first-order design constraint. The simulations identify muons, neutrons, and protons as the most relevant cosmic components after the Micromegas event selection. Muons dominate the raw surface rate but are efficiently tagged by prompt scintillator signals. High-energy cosmic neutrons are more difficult because they can interact in the lead shielding and produce secondary cascades that mimic low-energy detector events. Proton-induced residuals are smaller than the neutron component but remain non-negligible in the same simulation chain.

The cosmic-source treatment combines CRY-based generation, reference comparisons with EXPACS, and local neutron-flux information from the HENSA-related studies. The present status is adequate for detector-design and veto-optimization studies. The final absolute neutron-induced background prediction, however, still depends on the final source normalization and on a DESY-specific treatment of the surface environment.

## **Detector-response and background-model workflow**

The background model is organized around source-specific simulations followed by a common reconstruction and selection chain. This is the key methodological bridge in the thesis: the question is not only whether a simulated particle deposits energy in the gas, but whether the resulting event would survive the same energy, topology, timing, and veto selections applied to data. The approach is applied to external sources, cosmic-ray components, environmental radiation, gas and radon contamination, detector materials, electronics, and shielding contributions.

At the present stage, the background-model chapter defines the structure of the calculation more completely than it closes the final numerical budget. Several source classes still require final measured normalizations, updated DESY boundary conditions, or replacement of provisional tables with final activity and rate values. For that reason, the thesis now separates the background-model status from the final-rate target and identifies the missing inputs explicitly.

## **Active-veto findings**

The shielding and veto studies show that passive shielding alone is insufficient for the high-energy neutron problem. Lead shielding is necessary for gamma suppression, but it can also multiply or redistribute neutron-induced backgrounds through secondary

production. The active veto is therefore designed as more than a conventional muon veto: it combines prompt scintillator signals, delayed neutron-capture-related activity, and channel multiplicity in a multilayer plastic-scintillator and cadmium system.

The waveform-level simulations predict strong rejection of the residual muon component after Micromegas cuts and partial rejection of neutron- and proton-induced residuals. In the simulation chain discussed in the veto chapter, the simplified waveform-level veto reduces the residual neutron component by about 30% and the proton component by about 70% after Micromegas cuts. Prototype data taken with IAXO-D0 validate the same analysis logic: after the prompt muon veto, delayed and multiplicity-rich veto cuts remove 7 additional events from the 56-event post-veto sample and preserve about 97% of the calibration efficiency relative to the Micromegas-selected reference. The final measured background level in the 2–7 keV prototype analysis is  $(8.56 \pm 1.22) \times 10^{-7}$  counts keV<sup>-1</sup> cm<sup>-2</sup> s<sup>-1</sup>.

This result should be interpreted carefully. The prototype data validate the prompt/delayed/multiplicity discrimination strategy, not an absolute event-by-event neutron tag. The advanced veto selection is best described as a delayed/multiplicity-rich selection enriched in neutron-like activity. A quantitative neutron-veto efficiency measurement requires a simulation branch matched to the commissioned prototype geometry, thresholds, gains, attenuation, timing, and readout configuration.

## Remaining limitations and next steps

The main remaining uncertainties are now well identified. They include the absolute cosmic-neutron normalization, the extrapolation from Zaragoza and prototype conditions to DESY surface operation, **Geant4** hadronic modeling, quenching and light attenuation, PMT gains, veto thresholds, timing alignment, geometry gaps, accidental vetoes, dead time, and finite Monte Carlo statistics. These limitations do not invalidate the simulation program, but they define the work needed to convert the present design-level and validation-level studies into a final BabyIAXO background prediction.

Area	Result of this thesis	Open item
Software infrastructure	Common <code>REST-for-Physics/restG4</code> workflow for detector-response simulations.	Freeze final production configurations and preserve reproducible run metadata.
Cosmic-ray backgrounds	Surface cosmic-ray source terms and detector-level propagation used for shielding and veto design.	Finalize HENSA/DESY normalization and quote propagated uncertainties.
Background model	Source-specific simulations organized through common reconstruction and selection logic.	Complete the final source-by-source rate table.
Active veto	Multilayer scintillator-cadmium design validated at the level of prompt, delayed, and multiplicity observables.	Produce prototype-matched and final-design simulation branches.
Experimental comparison	Prototype data show additional delayed/multiplicity rejection beyond the prompt muon veto.	Increase statistics and quantify accidental veto/dead-time effects.

Table 6.13: Summary of thesis outputs and remaining work needed for the final BabyIAXO detector-background model.

In summary, this thesis establishes the simulation and analysis infrastructure needed to model the BabyIAXO Micromegas detector background at the level of reconstructed observables. It identifies the surface cosmic-neutron problem, develops the corresponding veto strategy, validates the qualitative prompt/delayed/multiplicity logic with prototype data, and defines the uncertainty roadmap required for a final absolute background prediction. The next step is to close the remaining normalizations and detector-response calibrations so that the framework developed here can be used as the quantitative background model for BabyIAXO operation.

# Appendix

```
1 #include <iostream>
2 #include <chrono>
3 #include <TFile.h>
4 #include <TTree.h>
5
6 using namespace std;
7
8 double read_write(const string &input_filename =
9     "/mnt/c/Users/lobis/git/FeminosDAQ/examples/iaxo-test.root", const string
10     &output_filename = "out.root")
11 {
12     const auto start = chrono::high_resolution_clock::now();
13
14     auto input_file = TFile::Open(input_filename.c_str(), "READ");
15     auto input_tree = input_file->Get<TTree>("events");
16
17     unsigned long long timestamp;
18     vector<unsigned short> *signal_ids = nullptr;
19     vector<unsigned short> *signal_values = nullptr;
20
21     input_tree->SetBranchAddress("timestamp", &timestamp);
22     input_tree->SetBranchAddress("signal_ids", &signal_ids);
23     input_tree->SetBranchAddress("signal_values", &signal_values);
24
25     auto output_file = TFile::Open(output_filename.c_str(), "RECREATE");
26     auto output_tree = new TTree("events", "events");
27
28     output_tree->Branch("timestamp", &timestamp);
29     output_tree->Branch("signal_ids", &signal_ids);
30     output_tree->Branch("signal_values", &signal_values);
31
32     for (int i = 0; i < input_tree->GetEntries(); i++)
33     {
34         input_tree->GetEntry(i);
35         output_tree->Fill();
36     }
37
38     output_tree->Write();
39     output_file->Close();
40
41     const auto elapsed_time_seconds =
42         chrono::duration_cast<chrono::milliseconds>(chrono::high_resolution_clock::now()
43             - start).count() * 0.001;
44     return elapsed_time_seconds;
```

## 6.7 Supplementary interaction plots for detector and veto design

The shielding and veto chapter includes only the interaction plots that directly support the active-veto design: fast-neutron reactions in lead, neutron moderation in low- $Z$  materials, and cadmium capture after moderation. Additional material-response plots that are useful for cross-checking the detector model, but less central to the veto narrative, are collected here. Figures 6.44 and 6.45 extend the neutron-data validation to the two active-veto materials most directly involved in neutron tagging: the BC408 plastic scintillator and the cadmium sheets. The evaluated reference curves are built from ENDF/B-VIII.0 MF=3 pointwise cross sections using the relevant isotope or atom-fraction weights. For BC408, the scintillator is approximated as polyvinyltoluene,  $C_9H_{10}$ , so the material curve is the atom-fraction-weighted sum of natural carbon and hydrogen. For cadmium, the curve is the natural-abundance sum of the stable Cd isotopes. The Geant4 points are extracted from the G4NDL4.6 high-precision neutron data library used by Geant4 11.0.3, with the same channel grouping as in the main lead validation.




Figure 6.44: Comparison of ENDF/B-VIII.0 and Geant4 high-precision neutron cross sections for the BC408 scintillator, approximated as  $C_9H_{10}$ . The agreement confirms that the Geant4 neutron data reproduce the evaluated elastic, capture, and inelastic channels relevant for neutron moderation and scintillator energy deposition over the thermal-to-fast energy range.




Figure 6.45: Comparison of ENDF/B-VIII.0 and Geant4 high-precision neutron cross sections for natural cadmium sheets. The large thermal and epithermal capture cross section is the nuclear-data basis for adding cadmium to the veto design: moderated neutrons reaching the Cd layer are efficiently captured, producing delayed gamma cascades that can be tagged by the surrounding scintillator.




Figure 6.46: Transmission of low-energy X-rays through the aluminized Mylar entrance window, calculated from photon attenuation data [55]. This signal-region effect belongs to the detector-efficiency model rather than to the veto rejection mechanism, and is therefore kept as supplementary material.

## 6.8 Supplementary veto simulation campaign metadata

The main veto-system chapter uses compact labels for the simulation campaigns in order to keep the design argument readable. The tables below record the source model, response level, and timing convention associated with each campaign. They are retained here as provenance for the comparison between historical design scans, final HENSA/CosmicMuons productions, and experimental prototype validation.

Study	Geometry and source model	Response and scope
<b>Inclination scan</b>	Full shielded detector with atmospheric neutrons from CRY.	<i>Response:</i> TPC background rate. <i>Scope:</i> demonstrates a variation of only about $\pm 3\%$ between extreme tracking angles.
<b>Lead-thickness scan</b>	Parameterized Pb shell with and without the pipe opening; all main cosmic secondaries.	<i>Response:</i> post-analysis background rate. <i>Scope:</i> quantifies the strong photon attenuation and the weak neutron benefit from additional lead.
<b>Pb/borated-HDPE/Pb study</b>	Fixed 20 cm total Pb with a scanned intermediate borated-HDPE layer; neutrons from CRY.	<i>Response:</i> post-analysis neutron background rate. <i>Scope:</i> tests passive moderation and confirms that it is insufficient on its own.
<b>Sandwich ordering scan</b>	Single 10 cm active scintillator slab combined with Pb, B-HDPE, and capture-material variants: BC408, BC408 with 5 mm Cd sheets, and EJ-254 5%; incident neutrons from 1 MeV to 100 GeV.	<i>Response:</i> Birks-quenched visible-energy threshold. <i>Scope:</i> isolates the effect of material ordering before the full segmented waveform-level veto is introduced.
<b>Layer optimization</b>	1-, 2-, 3-, and 4-layer veto geometries, with cadmium and capture-material variants.	<i>Response:</i> idealized deposited-energy scans followed by HENSA waveform-level validation. <i>Scope:</i> compares geometrical redundancy and tests that the 3-layer conclusion survives the full response chain.
<b>Full veto performance</b>	Final 59-panel geometry; CosmicMuons for muons, HENSA neutrons, and CRY cross-checks for the other atmospheric secondaries.	<i>Response:</i> waveform-level peaks after quenching, attenuation, shaping, and readout windows. <i>Scope:</i> provides the estimates most directly comparable to the experimental veto analysis.

Table 6.14: Main simulation configurations used throughout the veto-system chapter. This table makes explicit where the discussion refers to idealized deposited-energy studies and where it refers to waveform-level observables.

Study	Primary sample	Geometry / response level	Timing convention	Statistical / reproducibility note
Inclination scan	CRY neutrons in the detector-relevant energy range	Full shielded detector; comparison of TPC background rate only	Not applicable	Same source term and reconstruction settings used for the horizontal and two extreme tilted orientations; quoted result is the relative variation.
Lead and passive-neutron scans	CRY mixed secondaries or neutrons, as stated in each study	Parameterized shielding geometries; post-analysis TPC background rate	Not applicable	Common transport and analysis settings kept fixed within each geometry scan so that only the shielding layout changes.
Layer / cadmium comparison	High-energy neutron sample on simplified multi-layer veto layouts	Idealized deposited energy in scintillator volumes	Not applicable	Reported as cumulative rejection curves versus threshold; intended as comparative design indicators rather than final waveform-level efficiencies.
Full veto performance	CosmicMuons for muons, HENSA for neutrons, and CRY cross-check samples for the remaining atmospheric secondaries	Final 59-panel simulated geometry with quenching, attenuation, 500 ns sampling, 3000 ns shaping, and peak finding	100 $\mu$ s total acquisition window, Microegas trigger at 30 $\mu$ s	Residual gamma and electron entries are finite-statistics limits; waveform chain follows the common <code>Geant4/restG4</code> configuration described in the software and background-model chapters.
Experimental validation	Surface IAXO-D0 prototype data set (52.1 days)	Commissioned 57-signal implementation analyzed with waveform observables	100 $\mu$ s total acquisition window, Microegas trigger at 30 $\mu$ s	Statistical uncertainty and calibration efficiency are taken directly from the measured cut flow.

Table 6.15: Additional metadata for the simulation and validation campaigns discussed in the veto-system chapter. Where no explicit timing convention is given, the study is based on geometry-level or post-analysis rate comparisons rather than waveform-level veto observables.

## 6.9 Supplementary active-veto design diagnostics

The main text uses a simplified representative sandwich comparison to motivate the active-material choice. The full ordering scan and the quenching diagnostic are retained here because they document that the qualitative conclusion is stable across the scanned material orderings.




Figure 6.47: Energy-weighted visible fraction after applying the Birks correction in the simplified sandwich scans. The plotted quantity is  $\sum E_{\text{vis}} / \sum E_{\text{dep}}$ , aggregated over the material-ordering configurations for each scintillator option; shaded bands show the configuration-to-configuration range.

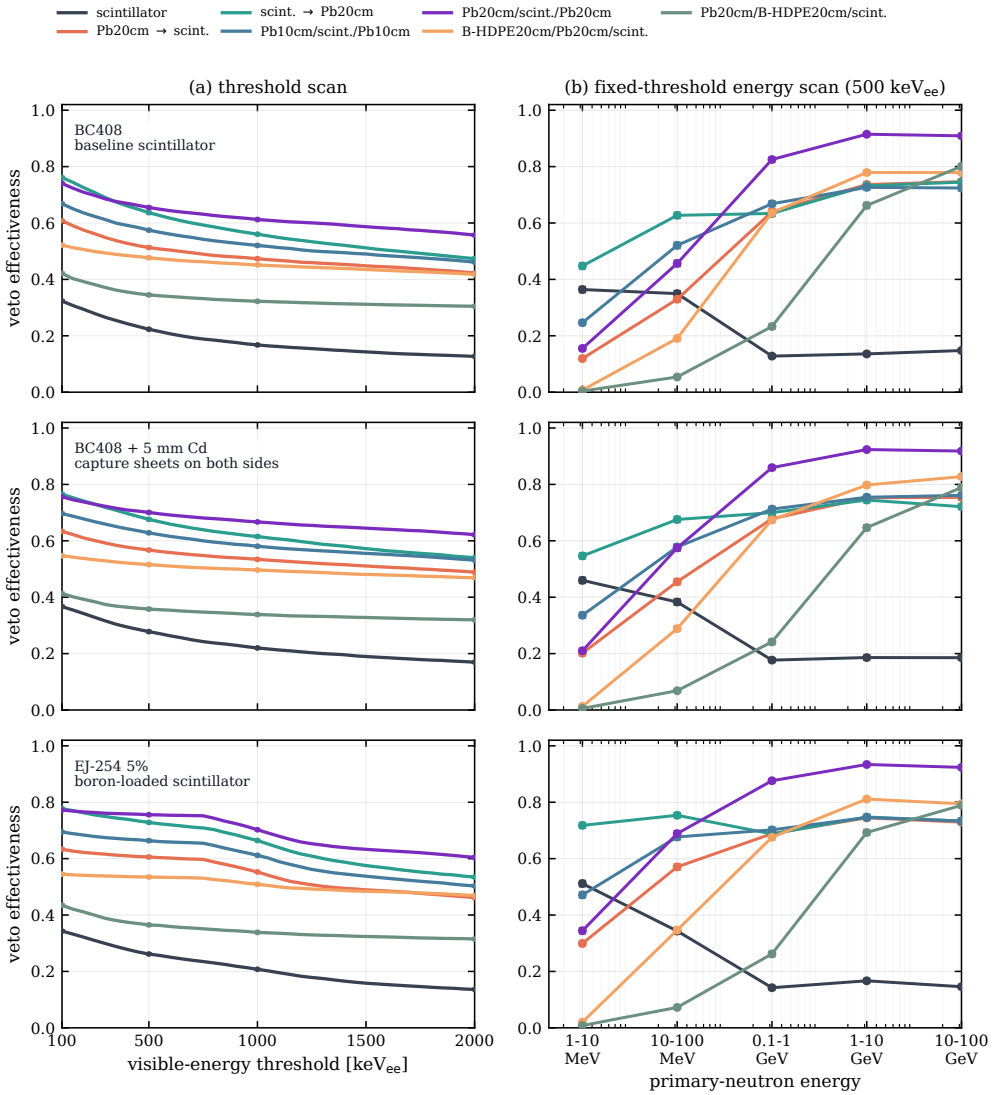


Figure 6.48: Full simplified sandwich scan for high-energy neutrons incident on the BC408-equivalent baseline, the same BC408 slab with 5 mm cadmium sheets on both sides, and an EJ-254 5% boron-loaded scintillator. Each row shows the scanned material orderings before light attenuation, timing-window selection, electronics shaping, and waveform-level peak reconstruction.

## 6.10 Supplementary HENSA veto layer-scan diagnostics

The main veto-system chapter uses a compact layer-design figure that combines threshold response, relative event rate, and capture-material comparison. The two original diagnostic projections are preserved here for traceability.




Figure 6.49: Neutron rejection as a function of reconstructed veto visible-energy threshold for the HENSA outdoor neutron layer scan. The curves use the detector-response chain rather than the older idealized deposited-energy observable.




Figure 6.50: Final HENSA outdoor neutron comparison for the 1-, 2-, 3-, and 4-layer cadmium veto geometries, including three-layer gadolinium and stainless-steel capture-material variants.

## 6.11 Supplementary neutron-tagging diagnostics

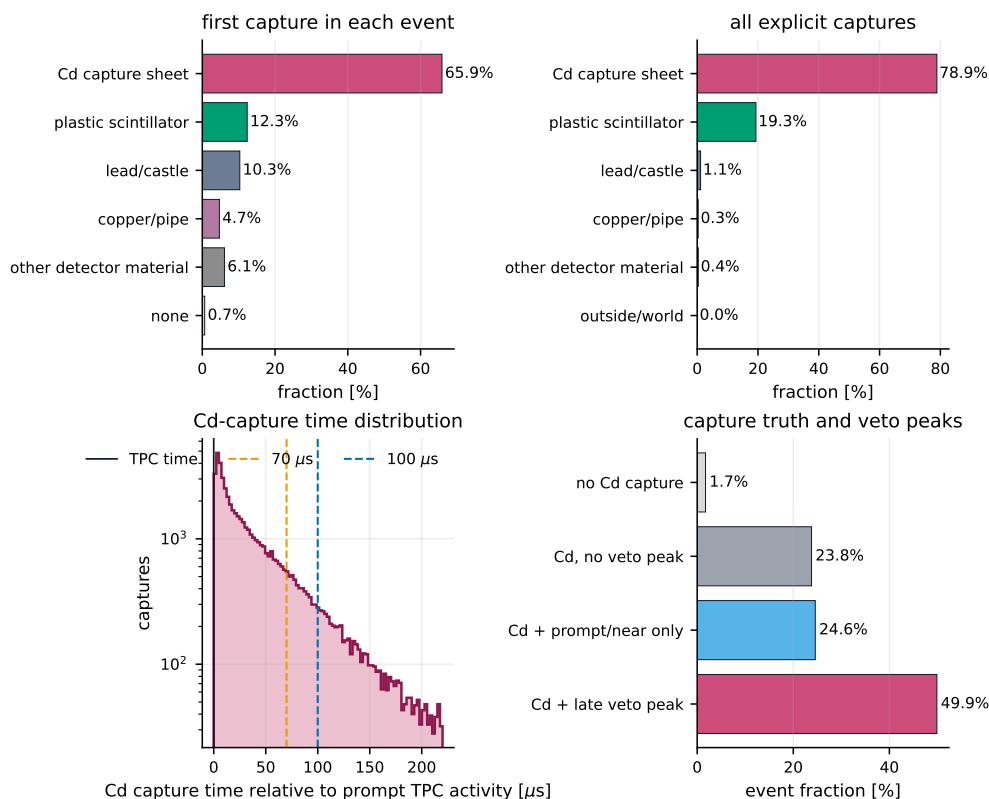


Figure 6.51: Truth-level neutron-capture diagnostics for the HENSA outdoor three-layer cadmium simulation after the common TPC selection. Top left: material region of the first neutron capture in each selected event. Top right: material region of all explicit neutron-capture processes. Bottom left: time distribution of cadmium-sheet captures for events with prompt TPC activity, with the 70  $\mu$ s and 100  $\mu$ s post-trigger reference windows indicated. Bottom right: event-level comparison between cadmium-capture truth and reconstructed veto-peak observables.

## 6.12 Supplementary neutron-enriched data—simulation diagnostics

The score distribution and delayed-veto comparison are kept in the main veto-system chapter. The event-raster and Micromegas-topology projections below document the residual differences between the selected experimental population and the selected neutron+noise simulation.

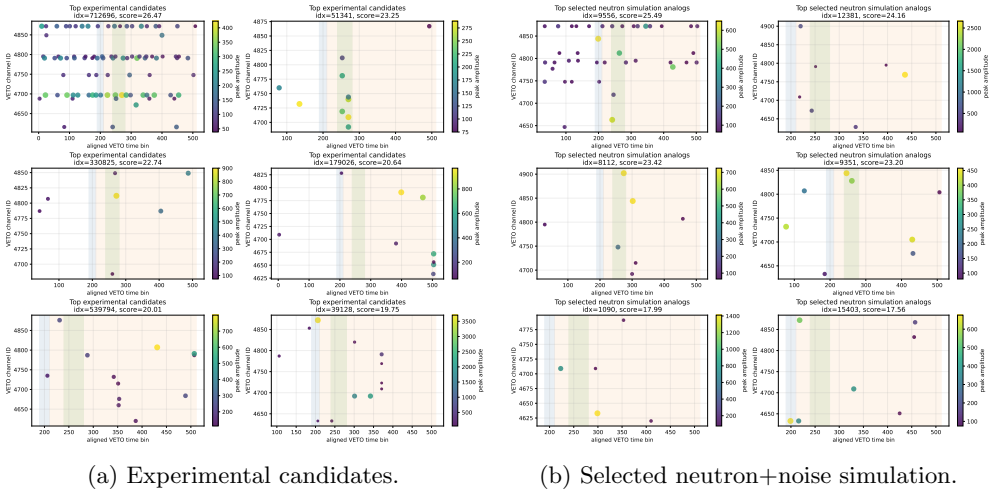


Figure 6.52: Representative event rasters from the neutron-enriched experimental sample and from selected neutron+noise simulation events. Each raster displays the veto peak pattern in channel and time coordinates for high-score events.

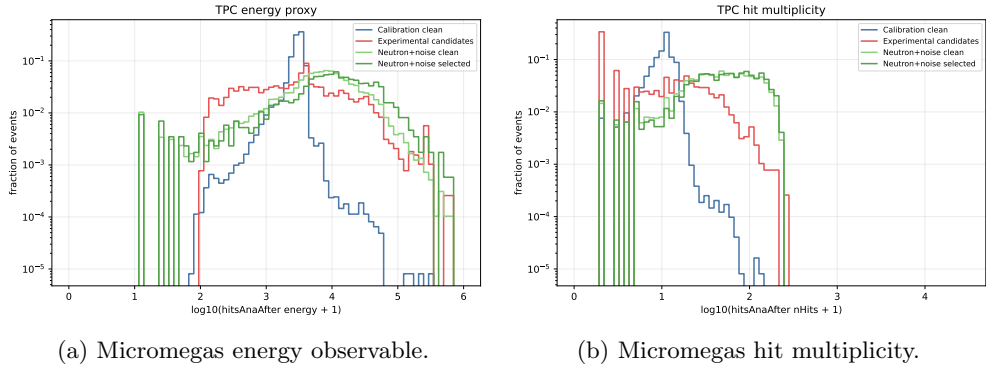


Figure 6.53: Micromegas-observable comparison for the neutron-enriched experimental candidates and selected neutron+noise simulation. The selected neutron simulation is generally more energetic and has larger hit multiplicity than the selected experimental candidates.

## 6.13 Supplementary passive-shielding scans

The main shielding and veto chapter uses a compact photon/neutron summary of the lead-thickness scan because those two components determine the passive-shielding design decision. The full per-particle scans are retained here as simulation provenance and as checks that the other cosmic-ray-induced components do not change the conclusion.




Figure 6.54: Muon-induced background rate as a function of lead shield thickness for the ideal  $4\pi$  and pipe-opening configurations. The weak dependence on lead thickness confirms that passive lead shielding is not a muon-mitigation strategy.




Figure 6.55: Proton-induced background rate as a function of lead shield thickness for the ideal  $4\pi$  and pipe-opening configurations. The scan is retained as a cross-check because proton-induced cascades can behave similarly to neutron-induced cascades after shielding interactions.




Figure 6.56: Electron-induced background rate as a function of lead shield thickness for the ideal  $4\pi$  and pipe-opening configurations. This component is not the limiting case for the veto design but is included for completeness.




Figure 6.57: Photon-induced background rate as a function of lead shield thickness for the ideal  $4\pi$  and pipe-opening configurations. This is the detailed version of the photon panel summarized in the main veto-system chapter.




Figure 6.58: Neutron-induced background rate as a function of lead shield thickness for the ideal  $4\pi$  and pipe-opening configurations. This is the detailed version of the neutron panel summarized in the main veto-system chapter.

## 6.14 Supplementary environmental-radioactivity plots

The background-model chapter keeps the representative  $^{238}\text{U}$  concrete-emission spectrum and the common exit-angle distribution in the main text. The remaining isotope spectra and decay-chain diagrams are collected here because they document the construction of the environmental-radioactivity source term, but their interpretation is the same as for the representative case discussed in the chapter.




Figure 6.59: Photons and electrons produced by the decay of  $^{235}\text{U}$  in a concrete block of 1 m depth.




Figure 6.60: Photons and electrons produced by the decay of  $^{232}\text{Th}$  in a concrete block of 1 m depth.




Figure 6.61: Photons and electrons produced by the decay of  $^{40}\text{K}$  in a concrete block of 1 m depth.

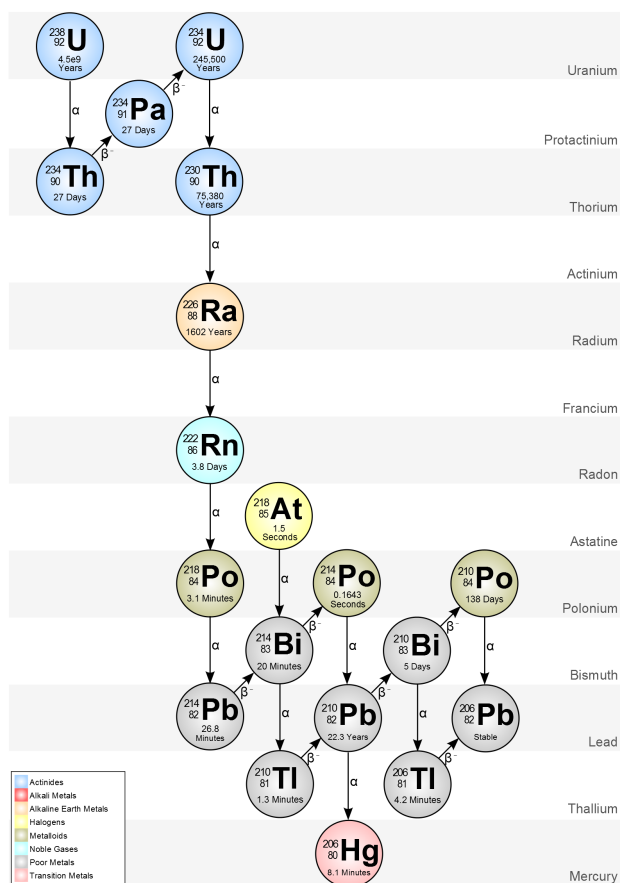


Figure 6.62: Decay chain of  $^{238}\text{U}$ , one of the natural radioactivity families contributing to the environmental gamma field.

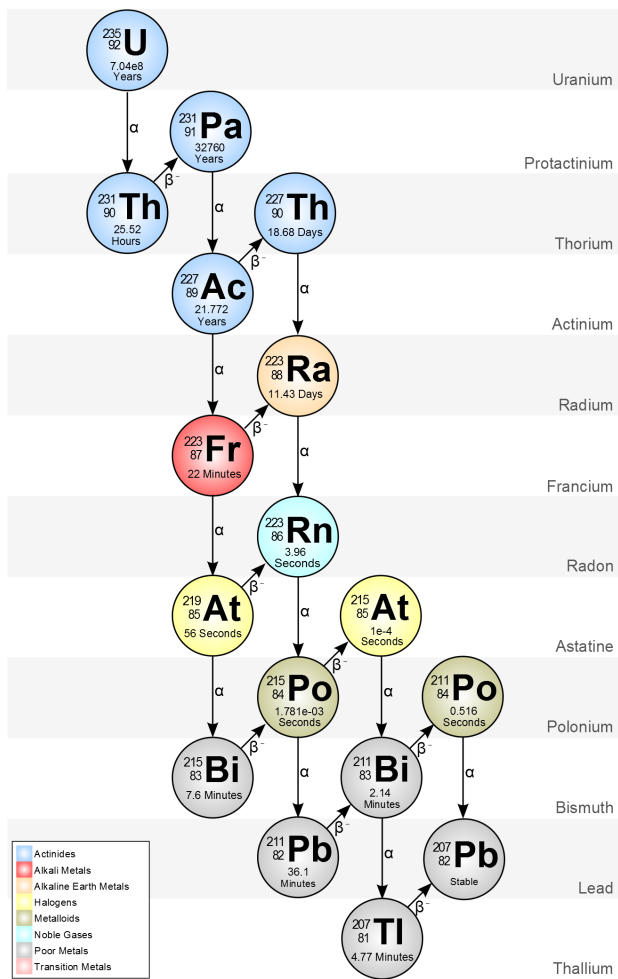


Figure 6.63: Decay chain of  $^{235}\text{U}$ , included in the concrete-radioactivity source model.

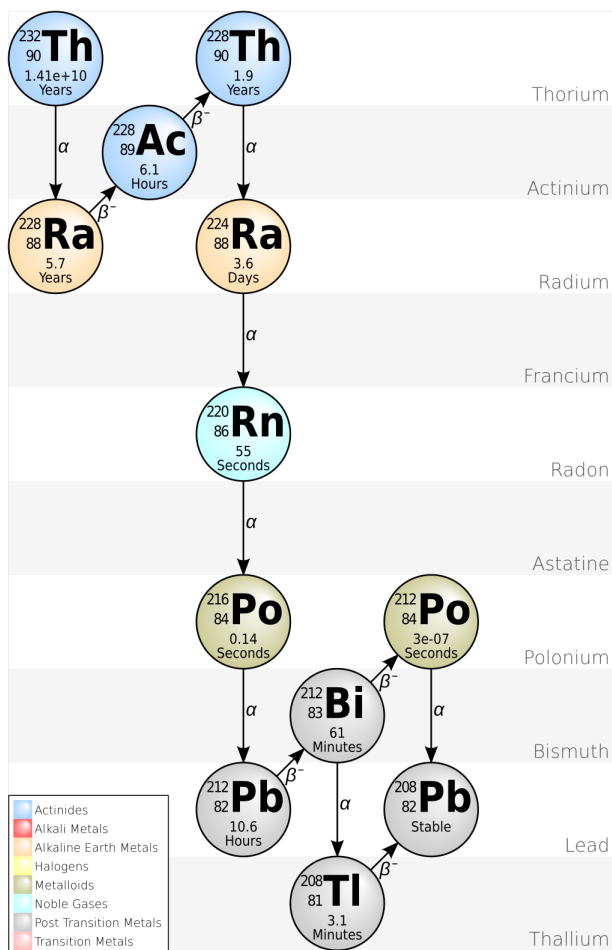


Figure 6.64: Decay chain of  $^{232}\text{Th}$ , included in the concrete-radioactivity source model.

## 6.15 Supplementary Micromegas system diagrams

The gas-system diagram used in the Micromegas chapter is reproduced here at a larger scale to make the line routing, valves, pressure sensors, and operating branches easier to inspect.

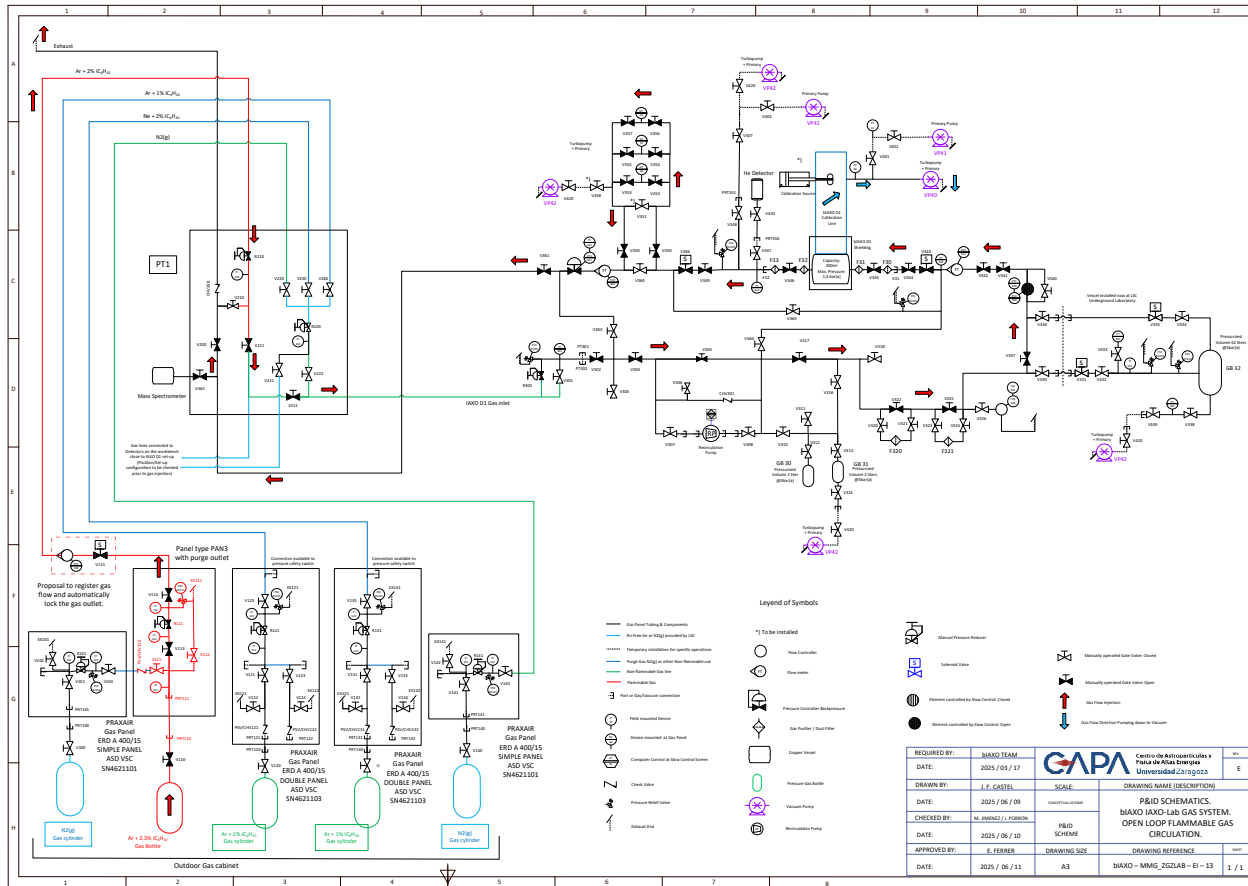


Figure 6.65: Full-page rotated version of the IAXO-D0/D1 gas-system diagram, referenced from Figure 3.21, p. 62.

## 6.16 Supplementary UV-light calibration R&D

During the CEA Saclay internship, a compact Micromegas setup was used to test whether pulsed ultraviolet light could provide a controllable source of photoelectrons for gas-transport and timing studies. The material is kept in the appendix because it is related to detector calibration and Micromegas operation, but it is not part of the baseline BabyIAXO background-model chain. The main text summarizes the methodological relevance of this work in Section 3.6.2.






Figure 6.66: PICOSEC detection concept [77], included as motivation for the use of ultraviolet photons and a photocathode with a Micromegas amplification structure. A charged particle produces Cherenkov photons in a radiator; the photons release photoelectrons at the photocathode, and the electrons are then drifted and amplified in the Micromegas stages. The CEA internship study used the same general idea of UV-induced photoelectrons, but in a simpler calibration-oriented geometry.



(a) Micromegas detector used in the CEA setup.



(b) Open view of the pulsed UV source.

Figure 6.67: Hardware elements used in the CEA UV-light calibration study. The pulsed UV lamp illuminated an aluminized cathode through a UV-transparent window, and the timing of the resulting Micromegas pulse was compared for different drift distances.




Figure 6.68: Simplified geometry of the UV-light calibration concept. A UV-transparent window and photocathode are used to generate primary photoelectrons at a known position, while spacers with different thicknesses change the drift distance. The arrival-time difference between configurations provides an estimate of the electron drift velocity in the selected gas mixture.



---

# Bibliography

- [1] R. J. Crewther et al. “Chiral Estimate of the Electric Dipole Moment of the Neutron in Quantum Chromodynamics”. In: *Phys. Lett.* 88B (1979). [Erratum: *Phys. Lett.* 91B, 487 (1980)], p. 123. DOI: 10.1016/0370-2693(80)91025-4 (cited on p. 19).
- [2] M. Pospelov and A. Ritz. “Electric dipole moments as probes of new physics”. In: *Annals Phys.* 318 (2005), pp. 119–169. DOI: 10.1016/j.aop.2005.04.002. arXiv: hep-ph/0504231 [hep-ph] (cited on p. 19).
- [3] C. Abel et al. “Measurement of the Permanent Electric Dipole Moment of the Neutron”. In: *Phys. Rev. Lett.* 124 (8 Feb. 2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803 (cited on p. 19).
- [4] T. P. Cheng and L. F. Li. *Gauge theory of elementary particle physics*. Oxford Science Publications. Oxford: Clarendon Press, 1984 (cited on p. 19).
- [5] H. Quinn. *The CP Puzzle in the Strong Interactions*. 2001. arXiv: hep-ph/0110050 [hep-ph] (cited on p. 19).
- [6] H. Leutwyler. “The Ratios of the light quark masses”. In: *Phys. Lett.* B378 (1996), pp. 313–318. DOI: 10.1016/0370-2693(96)00386-3. arXiv: hep-ph/9602366 [hep-ph] (cited on p. 19).
- [7] R. Mohapatra and G. Senjanović. “Natural suppression of strong P and T non-invariance”. In: *Physics Letters B* 79.3 (1978), pp. 283–286. DOI: 10.1016/0370-2693(78)90243-5 (cited on p. 19).
- [8] S. M. Barr. “Solving the Strong CP Problem without the Peccei-Quinn Symmetry”. In: *Phys. Rev. Lett.* 53 (4 July 1984), pp. 329–332. DOI: 10.1103/PhysRevLett.53.329 (cited on p. 19).
- [9] R. D. Peccei and H. R. Quinn. “CP Conservation in the Presence of Pseudoparticles”. In: *Phys. Rev. Lett.* 38 (25 June 1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440 (cited on pp. 19–20).

- [10] R. D. Peccei and H. R. Quinn. “Constraints imposed by CP conservation in the presence of pseudoparticles”. In: *Phys. Rev. D* 16 (6 Sept. 1977), pp. 1791–1797. DOI: 10.1103/PhysRevD.16.1791 (cited on p. 19).
- [11] S. Weinberg. “A New Light Boson?” In: *Phys. Rev. Lett.* 40 (4 Jan. 1978), pp. 223–226. DOI: 10.1103/PhysRevLett.40.223 (cited on pp. 19–20).
- [12] F. Wilczek. “Problem of Strong P and T Invariance in the Presence of Instantons”. In: *Phys. Rev. Lett.* 40 (5 Jan. 1978), pp. 279–282. DOI: 10.1103/PhysRevLett.40.279 (cited on pp. 19–20).
- [13] R. Crewther. “Chirality selection rules and the U(1) problem”. In: *Physics Letters B* 70.3 (1977), pp. 349–354. DOI: 10.1016/0370-2693(77)90675-X (cited on p. 20).
- [14] P. Di Vecchia and G. Veneziano. “Chiral Dynamics in the Large n Limit”. In: *Nucl. Phys.* B171 (1980), pp. 253–272. DOI: 10.1016/0550-3213(80)90370-3 (cited on p. 20).
- [15] M. Gorghetto and G. Villadoro. “Topological susceptibility and QCD axion mass: QED and NNLO corrections”. In: *Journal of High Energy Physics* 2019.3 (Mar. 2019), p. 33. DOI: 10.1007/JHEP03(2019)033 (cited on p. 20).
- [16] W. A. Bardeen et al. “Constraints on variant axion models”. In: *Nucl. Phys.* B279 (1987), pp. 401–428. DOI: 10.1016/0550-3213(87)90003-4 (cited on p. 20).
- [17] J. E. Kim. “Weak-Interaction Singlet and Strong CP Invariance”. In: *Phys. Rev. Lett.* 43 (2 July 1979), pp. 103–107. DOI: 10.1103/PhysRevLett.43.103 (cited on p. 20).
- [18] M. Dine et al. “A simple solution to the strong CP problem with a harmless axion”. In: *Physics Letters B* 104.3 (1981), pp. 199–202. DOI: 10.1016/0370-2693(81)90590-6 (cited on p. 20).
- [19] M. Shifman et al. “Can confinement ensure natural CP invariance of strong interactions?” In: *Nuclear Physics B* 166.3 (1980), pp. 493–506. DOI: 10.1016/0550-3213(80)90209-6 (cited on p. 20).
- [20] M. Dine et al. “A Simple Solution to the Strong CP Problem with a Harmless Axion”. In: *Phys. Lett.* 104B (1981), pp. 199–202. DOI: 10.1016/0370-2693(81)90590-6 (cited on p. 20).
- [21] A. R. Zhitnitsky. “On Possible Suppression of the Axion Hadron Interactions. (In Russian)”. In: *Sov. J. Nucl. Phys.* 31 (1980). [*Yad. Fiz.*31,497(1980)], p. 260 (cited on p. 20).

- [22] Y. Chikashige et al. “Are There Real Goldstone Bosons Associated with Broken Lepton Number?” In: *Phys. Lett.* 98B (1981), pp. 265–268. DOI: 10.1016/0370-2693(81)90011-3 (cited on p. 21).
- [23] A. E. Nelson and N. Seiberg. “R-symmetry breaking versus supersymmetry breaking”. In: *Nuclear Physics B* 416.1 (1994), pp. 46–62. DOI: 10.1016/0550-3213(94)90577-0 (cited on p. 21).
- [24] J. Jaeckel and A. Ringwald. “The Low-Energy Frontier of Particle Physics”. In: *Annual Review of Nuclear and Particle Science* 60.1 (2010), pp. 405–437. DOI: 10.1146/annurev.nucl.012809.104433. eprint: <https://doi.org/10.1146/annurev.nucl.012809.104433> (cited on p. 21).
- [25] P. Svrcek and E. Witten. “Axions in string theory”. In: *Journal of High Energy Physics* 2006.06 (June 2006), pp. 051–051. DOI: 10.1088/1126-6708/2006/06/051 (cited on p. 21).
- [26] I. G. Irastorza and J. Redondo. “New experimental approaches in the search for axion-like particles”. In: *Progress in Particle and Nuclear Physics* 102 (2018), pp. 89–159. DOI: 10.1016/j.ppnp.2018.05.003. arXiv: 1801.08127 [hep-ph] (cited on pp. 21–22, 31).
- [27] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (July 2018), A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO] (cited on p. 22).
- [28] P. Sikivie. “Experimental Tests of the ”Invisible” Axion”. In: *Phys. Rev. Lett.* 51 (16 Oct. 1983), pp. 1415–1417. DOI: 10.1103/PhysRevLett.51.1415 (cited on p. 22).
- [29] S. J. Asztalos et al. “SQUID-Based Microwave Cavity Search for Dark-Matter Axions”. In: *Phys. Rev. Lett.* 104 (4 Jan. 2010), p. 041301. DOI: 10.1103/PhysRevLett.104.041301 (cited on p. 22).
- [30] K. Ehret et al. “New ALPS results on hidden-sector lightweights”. In: *Physics Letters B* 689.4 (2010), pp. 149–155. DOI: 10.1016/j.physletb.2010.04.066 (cited on p. 23).
- [31] R. Ballou et al. “New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall”. In: *Phys. Rev. D* 92 (9 Nov. 2015), p. 092002. DOI: 10.1103/PhysRevD.92.092002 (cited on p. 23).
- [32] I. G. Irastorza, F. Avignone, S. Caspi, J. Carmona, T. Dafni, et al. “Towards a new generation axion helioscope”. In: *JCAP* 1106 (2011), p. 013. DOI: 10.1088/1475-7516/2011/06/013. arXiv: 1103.5334 [hep-ex] (cited on pp. 26, 29, 31).

- [33] I. Irastorza et al. *The International Axion Observatory IAXO. Letter of Intent to the CERN SPS committee*. Tech. rep. CERN, 2013 (cited on pp. 26, 31, 34).
- [34] E. Armengaud, F. Avignone, M. Betz, P. Brax, P. Brun, et al. “Conceptual Design of the International Axion Observatory (IAXO)”. In: *JINST* 9 (2014), T05002. DOI: 10.1088/1748-0221/9/05/T05002. arXiv: 1401.3233 [physics.ins-det] (cited on pp. 26, 29–34).
- [35] A. Abeln et al. “Conceptual design of BabyIAXO, the intermediate stage towards the International Axion Observatory”. In: *JHEP* 05 (2021), p. 137. DOI: 10.1007/JHEP05(2021)137. arXiv: 2010.12076 [physics.ins-det] (cited on pp. 26, 30, 33–34, 37–38).
- [36] A. Abeln et al. “Axion search with BabyIAXO in view of IAXO”. In: *PoS ICHEP2020* (2021), p. 631. DOI: 10.22323/1.390.0631. arXiv: 2012.06634 [physics.ins-det] (cited on pp. 26, 34).
- [37] S. Hoof, J. Jaeckel, and L. J. Thormaehlen. “Quantifying uncertainties in the solar axion flux and their impact on determining axion model parameters”. In: *JCAP* 09 (2021), p. 006. DOI: 10.1088/1475-7516/2021/09/006. arXiv: 2101.08789 [hep-ph] (cited on p. 27).
- [38] J. Redondo. “Solar axion flux from the axion-electron coupling”. In: *Journal of Cosmology and Astroparticle Physics* 2013.12 (Dec. 2013), pp. 008–008. DOI: 10.1088/1475-7516/2013/12/008 (cited on p. 27).
- [39] S. Andriamonje et al. “An improved limit on the axion–photon coupling from the CAST experiment”. In: *Journal of Cosmology and Astroparticle Physics* 2007.04 (Apr. 2007), pp. 010–010. DOI: 10.1088/1475-7516/2007/04/010 (cited on pp. 29–30).
- [40] E. Arik et al. “Probing eV-scale axions with CAST”. In: *JCAP* 0902 (2009), p. 008. DOI: 10.1088/1475-7516/2009/02/008. arXiv: 0810.4482 [hep-ex] (cited on pp. 29–30).
- [41] M. Arik, S. Aune, K. Barth, A. Belov, S. Borghi, et al. “CAST solar axion search with  $^3\text{He}$  buffer gas: Closing the hot dark matter gap”. In: *Phys. Rev. Lett.* 112 (2014), p. 091302. DOI: 10.1103/PhysRevLett.112.091302. arXiv: 1307.1985 [hep-ex] (cited on pp. 29–30).
- [42] V. Anastassopoulos, S. Aune, K. Barth, et al. “New CAST limit on the axion–photon interaction”. In: *Nature Physics* 13.6 (June 2017), pp. 584–590. DOI: 10.1038/nphys4109 (cited on pp. 29–30, 54, 63).
- [43] M. Kuster et al. “The X-ray Telescope of CAST”. In: *New J. Phys.* 9 (2007), p. 169. DOI: 10.1088/1367-2630/9/6/169. arXiv: physics/0702188 (cited on p. 30).

- 
- [44] F. Aznar, J. Castel, F. E. Christensen, et al. “A Micromegas-based low-background x-ray detector coupled to a slumped-glass telescope for axion research”. en. In: *Journal of Cosmology and Astroparticle Physics* 2015.12 (2015). Publisher: IOP Publishing, pp. 008–008. DOI: 10.1088/1475-7516/2015/12/008. (Visited on 08/02/2022) (cited on p. 30).
- [45] C. Margalejo Blasco. “Searching for solar axions with ultra-low background Micromegas detectors in the IAXO pathfinder system at CAST”. PhD thesis. Universidad de Zaragoza, Mar. 2026 (cited on pp. 30, 38, 64–66, 75–76, 82).
- [46] L. A. Obis et al. “Background model for the IAXO experiment”. Presentation at VIEWS24, Vienna. Apr. 2024 (cited on pp. 30, 170).
- [47] A. C. Jakobsen, M. J. Pivovarov, and F. E. Christensen. “X-ray optics for axion helioscopes”. In: *Optics for EUV, X-Ray, and Gamma-Ray Astronomy VI*. Vol. 8861. SPIE, 2013, pp. 416–422. DOI: 10.1117/12.2024476. (Visited on 08/02/2022) (cited on p. 32).
- [48] E. Ruiz Chóliz. “Ultra-Low Background Microbulk Micromegas X-Ray Detectors for Axion Searches in IAXO and BabyIAXO”. PhD thesis. University of Zaragoza, 2019 (cited on pp. 33, 57, 139, 188, 253).
- [49] T. Schiffer. “Detector development towards axion searches with BabyIAXO”. PhD thesis. Rheinische Friedrich-Wilhelms-Universität Bonn, 2025. DOI: 10.48565/bonndoc-600 (cited on pp. 33–34, 38).
- [50] U. Schneekloth et al. “Some Thoughts on IAXO Site and SDS”. Presentation at the 22nd IAXO Collaboration Meeting, Dortmund. Sept. 2025 (cited on pp. 35, 221).
- [51] L. Helary et al. “IAXO Technical Coordination Report”. Presentation at the 23rd IAXO Collaboration Meeting, DESY. Mar. 2026 (cited on pp. 36, 221).
- [52] L. Helary, A. Lindner, and I. Manthos. “General Information, IAXO Collaboration Meeting”. Presentation at the 23rd IAXO Collaboration Meeting, DESY. Mar. 2026 (cited on p. 36).
- [53] U. Schneekloth. “Size of (Baby)IAXO Site(s): Stray Fields”. Presentation at the 23rd IAXO Collaboration Meeting, DESY. Mar. 2026 (cited on p. 36).
- [54] C. Cogollos Triviño. “New experimental techniques for axion searches in the RADES and BabyIAXO experiments”. PhD thesis. Universitat de Barcelona, 2025. URL: <https://hdl.handle.net/2445/221031> (cited on pp. 38, 243).
- [55] M. J. Berger, J. H. Hubbell, S. M. Seltzer, et al. *XCOM: Photon Cross Section Database (version 1.5)*. <http://physics.nist.gov/xcom>. National Institute of Standards and Technology, Gaithersburg, MD. 2010. URL: <http://physics.nist.gov/xcom> (visited on 05/16/2024) (cited on pp. 41, 43, 117, 265).

- [56] C. Davisson. “Interaction of gamma-radiation with matter”. In: *Alpha-, Beta- and Gamma-ray Spectroscopy: Volume 1*. Ed. by K. Siegbahn. Chapter in edited volume. Amsterdam: North-Holland Publishing Company, 1965 (cited on p. 43).
- [57] F. Sauli. *Gaseous Radiation Detectors: Fundamentals and Applications*. Cambridge University Press, June 2014. DOI: 10.1017/cbo9781107337701 (cited on pp. 44, 46).
- [58] T. Watanabe, H. W. Schnopper, and F. N. Cirillo. “K X-Ray Fluorescence Yield of Argon”. In: *Physical Review* 127.6 (1962). DOI: 10.1103/physrev.127.2055 (cited on p. 44).
- [59] O. Klein and Y. Nishina. “Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac”. In: *Zeitschrift für Physik* 52.11–12 (Nov. 1929), pp. 853–868. DOI: 10.1007/bf01366453 (cited on p. 45).
- [60] S. Weinberg. *The Quantum Theory of Fields*. Cambridge University Press, June 1995. DOI: 10.1017/cbo9781139644167 (cited on p. 45).
- [61] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097 (cited on p. 47).
- [62] U. Fano. “Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions”. In: *Phys. Rev.* 72 (1 July 1947), pp. 26–29. DOI: 10.1103/PhysRev.72.26 (cited on p. 48).
- [63] H. Schindler and R. Veenhof. *Garfield++ User Guide*. Accessed: 2024-05-20. 2024. URL: <https://garfieldpp.web.cern.ch/garfieldpp/> (cited on pp. 48, 50).
- [64] H. Mirallas Sánchez. “Desarrollo de grandes planos de lectura Micromegas para experimentos de búsqueda de sucesos poco probables”. PhD thesis. Universidad de Zaragoza, 2024 (cited on p. 53).
- [65] G. Charpak, J. Derre, Y. Giomataris, and P. Rebourgeard. “MICROMEGAS, a multipurpose gaseous detector”. In: *Nucl. Instrum. Meth. A* 478 (2002). Ed. by M. Jeitler, M. Krammer, G. Neuhofer, M. Regler, and R. Wedenig, pp. 26–36. DOI: 10.1016/S0168-9002(01)01713-2 (cited on p. 54).
- [66] L. Martinelli. “ATLAS New Small Wheel performance studies after first year of operation”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1063 (2024), p. 169290. DOI: 10.1016/j.nima.2024.169290 (cited on p. 54).

- [67] K. Altenmüller et al. *A new upper limit on the axion-photon coupling with an extended CAST run with a Xe-based Micromegas detector*. arXiv preprint. June 2024. arXiv: 2406.16840 [hep-ex]. URL: <https://arxiv.org/abs/2406.16840> (cited on pp. 54, 63).
- [68] J. Castel, S. Cebrián, T. Dafni, et al. “The TREX-DM experiment at the Canfranc Underground Laboratory”. In: *Journal of Physics: Conference Series* 1468.1 (Feb. 2020), p. 012063. DOI: 10.1088/1742-6596/1468/1/012063 (cited on p. 54).
- [69] W. Zhang, H. Lin, Y. Liu, et al. “Status and prospects of the PandaX-III experiment”. In: *JINST* 18.12 (2023), p. C12001. DOI: 10.1088/1748-0221/18/12/C12001. arXiv: 2311.13396 [physics.ins-det] (cited on p. 54).
- [70] D. Calvet. “A Versatile Readout System for Small to Medium Scale Gaseous and Silicon Detectors”. In: *IEEE Trans. Nucl. Sci.* 61.1 (2014), pp. 675–682. DOI: 10.1109/TNS.2014.2299312 (cited on pp. 56, 68).
- [71] O. Foundation and Contributors. *Node-RED*. Accessed: 2024-05-20. 2024. URL: <https://nodered.org> (cited on p. 65).
- [72] P. Baron, S. Anvar, B. Blank, et al. “AGET, the Get Front-End ASIC, for the readout of the Time Projection Chambers used in Nuclear Physic Experiments”. In: *Nuclear Science Symposium Conference Record* (2011) (cited on pp. 67, 152).
- [73] P. Baron, D. Calvet, E. Delagnes, X. de la Broise, A. Delbart, et al. “AFTER, an ASIC for the readout of the large T2K time projection chambers”. In: *IEEE Trans.Nucl.Sci.* 55 (2008), pp. 1744–1752. DOI: 10.1109/TNS.2008.924067 (cited on p. 67).
- [74] B. Rabenstein and J. Volz. *Prometheus: A Next-Generation Monitoring System (Talk)*. Conference talk. Dublin, May 2015 (cited on p. 69).
- [75] J. Pivarski, H. Schreiner, A. Hollands, et al. *Uproot*. <https://github.com/scikit-hep/uproot5>. Sept. 2017. DOI: 10.5281/zenodo.4340632 (cited on pp. 70, 81).
- [76] IAXO Collaboration. *iaxo-simulations*. <https://github.com/iaxo/iaxo-simulations>. Accessed: 2026-04-14. 2026 (cited on pp. 76, 98).
- [77] J. Bortfeldt, F. Brunbauer, C. David, et al. “PICOSEC: Charged particle timing at sub-25 picosecond precision with a Micromegas based detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 903 (2018), pp. 317–325. DOI: 10.1016/j.nima.2018.04.033 (cited on pp. 78, 283).

- [78] K. Altenmüller, S. Cebrián, T. Dafni, et al. “REST-for-Physics, a ROOT-based framework for event oriented data analysis and combined Monte Carlo response”. In: *Computer Physics Communications* 273 (2022), p. 108281. DOI: 10.1016/j.cpc.2021.108281 (cited on pp. 80, 82, 86, 108, 192).
- [79] S. Agostinelli, J. Allison, K. Amako, et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8 (cited on pp. 80, 87, 108).
- [80] J. Allison, K. Amako, J. Apostolakis, et al. “Geant4 developments and applications”. In: *IEEE Transactions on Nuclear Science* 53.1 (2006), pp. 270–278. DOI: 10.1109/TNS.2006.869826 (cited on pp. 80, 87, 108).
- [81] J. Allison, K. Amako, J. Apostolakis, et al. “Recent developments in Geant4”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225. DOI: 10.1016/j.nima.2016.06.125 (cited on pp. 80, 87, 108).
- [82] R. Brun and F. Rademakers. “ROOT — An object oriented data analysis framework”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 389.1 (1997). *New Computing Techniques in Physics Research V*, pp. 81–86. DOI: 10.1016/S0168-9002(97)00048-X (cited on pp. 80–81).
- [83] C. Haggmann, D. Lange, and D. Wright. “Cosmic-ray shower generator (CRY) for Monte Carlo transport codes”. In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 2. 2007, pp. 1143–1146. DOI: 10.1109/NSSMIC.2007.4437209 (cited on pp. 80, 92, 110).
- [84] D. Díez Ibáñez and L. A. Obis Aparicio. “Efficient cosmic ray generator for particle detector simulations”. In: *Comput. Phys. Commun.* 317 (2025), p. 109805. DOI: 10.1016/j.cpc.2025.109805. arXiv: 2501.07452 [hep-ex] (cited on pp. 80, 90–91, 196).
- [85] R. Chytracek, J. McCormick, W. Pokorski, and G. Santin. “Geometry description markup language for physics simulation and analysis applications.” In: *IEEE Trans. Nucl. Sci.* 53 (2006), p. 2892. DOI: 10.1109/TNS.2006.881062 (cited on p. 89).
- [86] Pivarski, Jim, Elmer, Peter, and Lange, David. “Awkward Arrays in Python, C++, and Numba”. In: *EPJ Web Conf.* 245 (2020), p. 05023. DOI: 10.1051/epjconf/202024505023 (cited on pp. 90, 103).

- 
- [87] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. *CORSIKA: A Monte Carlo code to simulate extensive air showers*. Tech. rep. Karlsruhe: Forschungszentrum Karlsruhe, Feb. 1998 (cited on p. 92).
- [88] L. A. Obis Aparicio. *cosmic-ray-secondaries*. <https://github.com/lobis/cosmic-ray-secondaries>. GitHub repository used to generate cosmic-ray secondary distributions from CRY for the REST-for-Physics simulation workflow. Accessed: 2026-04-14. 2024 (cited on p. 92).
- [89] D. Pfeiffer et al. “Interfacing Geant4, Garfield++ and Degrad for the Simulation of Gaseous Detectors”. In: *Nucl. Instrum. Meth. A* 935 (2019), pp. 121–134. DOI: 10.1016/j.nima.2019.04.110. arXiv: 1806.05880 [physics.ins-det] (cited on p. 93).
- [90] E. Moyse, F. Ali, E. Cortina, R. M. Bianchi, and B. Couturier. “The Phoenix event display framework”. In: *EPJ Web of Conferences* 251 (2021), p. 01007. DOI: 10.1051/epjconf/202125101007 (cited on p. 95).
- [91] L. A. Obis Aparicio. *feminos-daq*. <https://github.com/rest-for-physics/feminos-daq>. Data acquisition software for FEMINOS electronics, including ROOT output, Prometheus monitoring, and the feminos-viewer live data viewer. Accessed: 2026-05-03. 2024 (cited on p. 96).
- [92] D. Thain, T. Tannenbaum, and M. Livny. “Distributed computing in practice: the Condor experience”. In: *Concurrency and Computation: Practice and Experience* 17.2-4 (2005), pp. 323–356. DOI: 10.1002/cpe.938 (cited on p. 98).
- [93] T. Mkrtchyan, K. Chitrapu, V. Garonne, et al. “dCache: Inter-disciplinary storage system”. In: *EPJ Web of Conferences* 251 (2021), p. 02010. DOI: 10.1051/epjconf/202125102010 (cited on p. 101).
- [94] T. Sato. “Analytical Model for Estimating the Zenith Angle Dependence of Terrestrial Cosmic Ray Fluxes”. In: *PLOS ONE* 11.8 (2016), e0160390. DOI: 10.1371/journal.pone.0160390 (cited on p. 110).
- [95] T. Sato. “Analytical Model for Estimating Terrestrial Cosmic Ray Fluxes Nearly Anytime and Anywhere in the World: Extension of PARMA/EXPACS”. In: *PLOS ONE* 10.12 (2015), e0144679. DOI: 10.1371/journal.pone.0144679 (cited on p. 110).
- [96] M. Gordon, P. Goldhagen, K. Rodbell, et al. “Measurement of the flux and energy spectrum of cosmic-ray induced neutrons on the ground”. In: *IEEE Transactions on Nuclear Science* 51.6 (2004), pp. 3427–3434. DOI: 10.1109/TNS.2004.839134 (cited on p. 110).

- [97] L. A. Obis et al. “Status of Cosmic-Ray Induced Background Simulations”. Presentation at the 20th IAXO Collaboration Meeting. Oct. 2023 (cited on pp. 110, 114, 127, 134–135, 141–142).
- [98] M. Guan, M.-C. Chu, J. Cao, K.-B. Luk, and C. Yang. “A parametrization of the cosmic-ray muon flux at sea-level”. In: *arXiv preprint* (Sept. 2015). arXiv: 1509.06176 [hep-ex] (cited on p. 111).
- [99] T. K. Gaisser, R. Engel, and E. Resconi. *Cosmic Rays and Particle Physics*. 2nd ed. Cambridge University Press, 2016. DOI: 10.1017/CB09781139192194 (cited on p. 112).
- [100] S. E. A. Orrigo, J. L. Tain, N. Mont-Geli, et al. “Long-term evolution of the neutron rate at the Canfranc Underground Laboratory”. In: *European Physical Journal C* 82 (2022), p. 814. DOI: 10.1140/epjc/s10052-022-10755-6. arXiv: 2204.14263 (cited on p. 113).
- [101] D. A. Brown et al. “ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data”. In: *Nuclear Data Sheets* 148 (2018), pp. 1–142. DOI: 10.1016/j.nds.2018.02.001 (cited on pp. 118, 121).
- [102] Geant4 Collaboration. *Geant4 Physics Reference Manual*. Release 11.0 documentation. Accessed: 2026-05-13. 2021. URL: <https://geant4.web.cern.ch/docs/> (cited on pp. 118, 122).
- [103] Geant4 Collaboration. *Geant4 Neutron Data Library 4.6*. Neutron data library distributed with Geant4 11.0. Accessed: 2026-05-13. 2021. URL: <https://geant4.web.cern.ch/download/11.0.0.html> (cited on pp. 118, 122).
- [104] D. J. Thomas, P. Kolkowski, N. J. Roberts, B. D. Mellow, and M. J. Joyce. *Investigation of a Possible Replacement for Cadmium as a Thermal Neutron Absorbing Material in Neutron Instruments*. Tech. rep. DQL-RN 017. NPL Report DQL-RN 017. National Physical Laboratory, 2006. URL: <https://eprintpublications.npl.co.uk/3580/> (cited on p. 120).
- [105] R. B. Firestone and Z. Revay, eds. *Database of Prompt Gamma Rays from Slow Neutron Capture for Elemental Analysis*. IAEA STI/PUB/1263. Vienna: International Atomic Energy Agency, 2007. URL: <https://www-nds.iaea.org/pgaa/> (cited on pp. 121, 123).
- [106] G. Rusev, M. Jandel, M. Krtička, et al. “Cascade Gamma Rays Following Capture of Thermal Neutrons on  $^{113}\text{Cd}$ ”. In: *Physical Review C* 88 (2013), p. 057602. DOI: 10.1103/PhysRevC.88.057602 (cited on p. 121).

- 
- [107] L. Weimer, E. Ellingwood, O. Fischer, M. Lai, and S. Westerdale. “G4CASCADE: A Data-Driven Implementation of  $(n, \gamma)$  Cascades in Geant4”. In: *arXiv preprint arXiv:2408.02774* (2024). DOI: 10.48550/arXiv.2408.02774. arXiv: 2408.02774 [physics.comp-ph] (cited on p. 122).
- [108] P. Grabmayr, T. Hehl, A. Stahl, J. Annand, and R. Owens. “A high-resolution, large acceptance scintillation time-of-flight spectrometer”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 402.1 (1998), pp. 85–94. DOI: 10.1016/S0168-9002(97)01068-1 (cited on p. 137).
- [109] Eljen Technology. *EJ-200, EJ-204, EJ-208, EJ-212 Plastic Scintillators Data Sheet*. Accessed: 2024-05-22. Modern equivalent of NE-110. Eljen Technology. 2024. URL: [https://eljentechnology.com/images/products/data\\_sheets/EJ-200\\_EJ-204\\_EJ-208\\_EJ-212.pdf](https://eljentechnology.com/images/products/data_sheets/EJ-200_EJ-204_EJ-208_EJ-212.pdf) (cited on p. 137).
- [110] Luxium Solutions. *BC-400, BC-404, BC-408, BC-412, BC-416 Premium Plastic Scintillators Data Sheet*. Accessed: 2026-05-13. Luxium Solutions. 2021. URL: <https://luxiumsolutions.com/sites/default/files/2021-10/BC400-404-408-412-416-Data-Sheet.pdf> (cited on p. 137).
- [111] K. Altenmüller, J. F. Castel, S. Cebrián, et al. “Background discrimination with a Micromegas detector prototype and veto system for BabyIAXO”. In: *Frontiers in Physics* 12 (2024). DOI: 10.3389/fphy.2024.1384415 (cited on pp. 138, 149, 151–153, 168–170, 192, 219).
- [112] J. B. Birks. *The Theory and Practice of Scintillation Counting*. Oxford: Pergamon Press, 1964 (cited on pp. 139, 205).
- [113] Geant4 Collaboration. *Geant4 Book for Application Developers: Birks Quenching*. Accessed: 2026-05-09. 2026. URL: [https://geant4.web.cern.ch/documentation/dev/bfad\\_html/ForApplicationDevelopers/Detector/birks.html](https://geant4.web.cern.ch/documentation/dev/bfad_html/ForApplicationDevelopers/Detector/birks.html) (cited on pp. 139, 205).
- [114] J. Lindhard, V. Nielsen, M. Scharff, and P. V. Thomsen. “Integral Equations Governing Radiation Effects”. In: *Matematisk-fysiske Meddelelser, Kongelige Danske Videnskabernes Selskab* 33.10 (1963), pp. 1–42 (cited on pp. 139, 205).
- [115] L. A. Obis et al. “Neutron Veto System Simulations”. Presentation at IAXO Collaboration. Aug. 2023 (cited on pp. 141–143, 145, 153–154, 156).
- [116] L. A. Obis et al. “Status of Veto Simulations”. Presentation at the 18th IAXO Collaboration Meeting. Sept. 2023 (cited on pp. 141–142, 153, 156).
- [117] L. A. Obis et al. “Cosmic ray-induced background simulations for IAXO-D1”. Presentation at the 19th IAXO Collaboration Meeting. Mar. 2024 (cited on pp. 164, 166–167).

- [118] S. Cebrian et al. “Radiopurity of Micromegas readout planes”. In: *Astropart. Phys.* 34 (2011), pp. 354–359. DOI: 10.1016/j.astropartphys.2010.09.003. arXiv: 1005.2022 [physics.ins-det] (cited on pp. 190–191).
- [119] F. Aznar et al. “Assessment of material radiopurity for Rare Event experiments using Micromegas”. In: *Journal of Instrumentation* 8.11 (Nov. 2013), pp. C11012–C11012. DOI: 10.1088/1748-0221/8/11/c11012 (cited on pp. 190–191).
- [120] S. Aune et al. “Low background x-ray detection with Micromegas for axion research”. In: *Journal of Instrumentation* 9.01 (Jan. 2014), P01001–P01001. DOI: 10.1088/1748-0221/9/01/p01001 (cited on pp. 190–191).
- [121] M. Jimenez Puyuelo et al. “IAXO-D1 Micromegas Setups at LSC and Zaragoza: Status of Background Model”. Presentation at the 23rd IAXO Collaboration Meeting, DESY. Mar. 2026 (cited on pp. 190–191).
- [122] E. Picatoste et al. “Update the Micromegas Radiopure Electronics”. Presentation at the 21st IAXO Collaboration Meeting. Mar. 2025 (cited on pp. 190–191, 243).
- [123] M. Jimenez Puyuelo et al. “Status of IAXO-D1 Prototypes at Zaragoza and Canfranc”. Presentation at the 21st IAXO Collaboration Meeting. Mar. 2025 (cited on p. 191).
- [124] M. Jimenez Puyuelo et al. “Status of the IAXO-D1 Prototype at Canfranc”. Presentation at the 22nd IAXO Collaboration Meeting, Dortmund. Sept. 2025 (cited on p. 191).
- [125] G. Ros, G. Sáez-Cano, G. A. Medina-Tanco, and A. D. Supanitsky. “On the Design of Experiments Based on Plastic Scintillators Using Geant4 Simulations”. In: *Radiation Physics and Chemistry* 153 (2018), pp. 140–151. DOI: 10.1016/j.radphyschem.2018.09.021. arXiv: 1804.08975 [astro-ph.IM] (cited on p. 205).
- [126] T. Pöschl, D. Greenwald, M. J. Losekamm, and S. Paul. “Measurement of Ionization Quenching in Plastic Scintillators”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 988 (2021), p. 164865. DOI: 10.1016/j.nima.2020.164865. arXiv: 2007.08366 [physics.ins-det] (cited on p. 206).
- [127] F. Pedregosa, G. Varoquaux, A. Gramfort, et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830 (cited on p. 214).
- [128] H. Wulandari, J. Jochum, W. Rau, and F. von Feilitzsch. “Neutron flux underground revisited”. In: *Astropart. Phys.* 22 (2004), pp. 313–322. DOI: 10.1016/j.astropartphys.2004.07.005. arXiv: hep-ex/0312050 (cited on pp. 232, 253).

- [129] P. Benetti, F. Calaprice, E. Calligarich, et al. “Measurement of the specific activity of  $^{39}\text{Ar}$  in natural argon”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 574.1 (2007), pp. 83–88. DOI: 10.1016/j.nima.2007.01.106 (cited on p. 234).
- [130] K. Altenmüller, J. F. Castel, S. Cebrián, et al. “AlphaCMM, a Micromegas-based camera for high-sensitivity screening of alpha surface contamination”. In: *Journal of Instrumentation* 17.08 (Aug. 2022), P08035. DOI: 10.1088/1748-0221/17/08/P08035 (cited on p. 236).