Karlsruhe Institute of Technology

Supernova Neutrino Detection in Dual-Phase Xenon Detectors and the SNEWS Network

Zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften (Dr. rer. nat.)

von der KIT-Fakultät für Physik
des Karlsruher Instituts für Technologie (KIT)

genehmigte

Dissertation

von

M. Sc. Melih Kara

Referentin:	Prof. Dr. Kathrin Valerius Insitut für Astroteilchenphysik Karlsruher Institut für Technologie
Korreferent:	Prof. Dr. Guido Drexlin Institut für Experimentelle Teilchenphysik Karlsruher Institut für Technologie

Tag der mündlichen Prüfung: 31.01.2025

Erklärung:

Ich versichere wahrheitsgemäß, die Arbeit selbstständig verfasst, alle benutzten Hilfsmittel vollständig und genau angegeben und alles kenntlich gemacht zu haben, was aus Arbeiten anderer unverändert oder mit Abänderungen entnommen wurde sowie die Satzung des KIT zur Sicherung guter wissenschaftlicher Praxis in der Fassung vom März 27, 2017 beachtet zu haben.

Ort, Datum Unterschrift

Supernova Neutrino Detection in Dual-Phase Xenon Detectors and the SNEWS Network

Short abstract:
Core-collapse supernovae are among the Universe's most powerful phenomena, releasing a vast flux of neutrinos that carry critical information about the supernova mechanism, stellar evolution, and neutrino properties. Dual-phase liquid xenon time projection chambers, like the XENONnT experiment, originally designed for dark matter searches, are highly sensitive to supernova neutrino signals via coherent elastic neutrino–nucleus scattering. This thesis explores the capabilities of XENONnT in detecting such signals, focusing on data analysis techniques, signal selection methods, and integration into the Supernova Early Warning System. SNEWS.

A framework for simulating and analyzing supernova neutrino interactions in XENONnT is developed, employing machine learning tools and statistical methods to distinguish signals from the background data. The study demonstrates that XENONnT can achieve significant detection probabilities, even under challenging conditions. It also discusses the enhancement of detection confidence through auxiliary detectors of XENONnT, namely its muon and neutron veto systems. The work concludes by exploring the potential of next–generation experiments, such as DARWIN and XLZD, to advance the field of multimessenger astrophysics. This thesis highlights the synergy between dark matter detectors and supernova neutrino studies, preparing for the next galactic supernova event.

Keywords: Supernova, Dark Matter, Supernova Early Warning System, neutrinos

Supernova-Neutrino-Nachweis mit Zweiphasen-Xenon-Detektoren und das SNEWS-Netzwerk

Kurze Zusammenfassung:

Kernkollaps-Supernovae gehören zu den stärksten Phänomenen des Universums und setzen einen kurzfristigen, enorm hohen Neutrinofluss frei, der wichtige Informationen über den Supernova-Mechanismus, die Sternentwicklung und die Neutrino-Eigenschaften enthält. Zweiphasen-Flüssig-Xenon-Zeitprojetionskammern wie das XENONnT-Experiment, das für die Suche nach dunkler Materie entwickelt wurde, sind hochempfindlich für Supernova-Neutrinosignale durch kohärente elastische Neutrino-Kern-Streuung. In dieser Arbeit werden die Möglichkeiten von XENONnT zur Erkennung solcher Signale untersucht, wobei der Schwerpunkt auf Datenanalysetechniken, Methoden zur Signalauswahl und der Integration des Experiments in das Supernova-Frühwarnsystem SNEWS liegt.

Zentrales Element ist die Entwicklung von Simulations- und Analyse–Software zu Supernova-Neutrino-Wechselwirkungen in XENONnT, wobei maschinelles Lernen und statistische Methoden eingesetzt werden, um Signale von Untergrunddaten zu unterscheiden. Diese Arbeit zeigt, dass XENONnT selbst bei niedrigen Signalstärken signifikante Nachweiswahrscheinlichkeiten für Supernovae erreicht. Durch die Einbindung der Myon– und Neutron–Vetosysteme von XENONnT kann die Nachweiswahrscheinlichtkeit noch weiter erhöht werden. Die Arbeit schließt mit einer Untersuchung des Potenzials von Experimenten der nächsten Generation wie DARWIN und XLZD für die Weiterentwicklung der Multimessenger-Astroteilchenphysik. Diese Arbeit verdeutlicht die Synergie zwischen Detektoren der direkten Suche nach dunkler Materie und Untersuchungen von Supernovae und den in diesen Prozessen emittierten Neutrinos. Damit leistet diese Studie einen signifikanten Beitrag zum potenziellen Nachweis von Neutrinos einer zukünftigen Supernova-Explosion.

Schlagwörter: Supernova, Dunkle Materie, Neutrinos, Supernova-Frühwarnsystem

Contents

Preface

Core–collapse supernovae (CCSNe) are among the most energetic and mysterious phenomena in the Universe, driving the synthesis of heavy elements and shaping the evolution of galaxies. The large flux of neutrinos emitted during the collapse carries valuable information about the supernova's mechanism, the properties of neutrinos, and the conditions in the progenitor star. Meanwhile, dark matter, an invisible yet pervasive substance, shapes the cosmos on grand scales, holding the key to understanding cosmic structure formation. The intersection of these phenomena presents a unique opportunity to leverage cutting–edge dark matter detectors to study neutrinos from supernovae, bridging the realms of dark matter research and multimessenger astrophysics.

This thesis investigates the potential of dual–phase liquid xenon time projection chambers (LXe TPCs), such as the XENONnT experiment, for detecting supernova neutrinos through coherent elastic neutrino–nucleus scattering (CEνNS). These detectors, originally designed to probe weakly interacting massive particles (WIMPs), are highly sensitive to low–energy nuclear recoils, making them ideal for observing the burst of neutrinos from a nearby CCSN. This study also explores the integration of such detectors into the Supernova Early Warning System (SNEWS), contributing to multimessenger astronomy efforts and preparing for the next galactic CCSN.

The structure of this thesis is as follows: Chapter 1 introduces the physics of supernovae and dark matter, setting the foundation for the discussions that follow. Chapter 2 focuses on the principles and applications of liquid xenon technologies, detailing the XENONnT experiment and its capabilities. In Chapter 3, the Supernova Early Warning System is presented, highlighting its transition to SNEWS 2.0 and the tools developed in this work for signal publishing and coincidence detection. Chapter 4 discusses the supernova neutrino emission models and their expected interactions in the XENONnT detector, along with the techniques employed to simulate these interactions under realistic detector conditions. The analysis of the simulated supernova signals, integrated with background data, is presented in Chapter 5, where methods for prompt signal identification, background suppression, and significance evaluation are described. Finally, Chapter 6 explores further paths expanding this study, the role of next–generation dark matter detectors like DARWIN and XLZD, and the anticipated impact of these experiments on supernova neutrino detection. This chapter concludes with reflections on the potential of these efforts to provide new insights into the next galactic supernova event.

This work demonstrates the critical role of dark matter detectors in multimessenger astrophysics, using the capabilities of LXe TPCs to decode the physics of supernovae and prepare for the next galactic supernova event.

Dedicated to Turan, Nimet, Hafize and Azim.

Acknowledgments

This thesis would not have been possible without the guidance, support and encouragement of many individuals and institutions.

First and foremost, I express my deepest gratitude to my supervisor, Kathrin Valerius, for her invaluable guidance and unwavering encouragement throughout this journey. I also am very grateful to K. Eitel and G. Drexlin for their mentorship and insightful advice during my thesis work. My heartfelt thanks extend to my collaborators in the XENON collaboration; R. Lang, A. Baxter, R. Peres, S. Ghosh, L.D. Garcia, and A. Molinario, and in the SNEWS collaboration; S. BenZvi, A. Habig, S. Torres–Lara, M. Linvill, J. Vasel, and J. Migenda. Their expertise and collaborative spirit greatly enriched this work. I also wish to acknowledge my colleagues in the Institute of Astroparticle Physics; V.H.S. Wu, S. Dharani, S. Vetter, F. Toschi, and A. Elykov for their camaraderie and for fostering an inspiring and collaborative research environment. Special thanks to Yanina Biondi, whose invaluable assistance transformed my thesis from gibberish to brilliance.

In the preparation of this thesis, I employed AI-based tools, including OpenAI's ChatGPT and Writefull, for cosmetic editing, grammar refinement, and structural suggestions. These tools were used exclusively to enhance the clarity and presentation of the document, with all academic content being my original work.

Finally, I am profoundly grateful to my family; Sevda, Mehmet, Semih, Nil, and Masal Kara, and Burak Bayrak for their unconditional love and support. To my friends, both here in Karlsruhe and back in Turkey, thank you for your encouragement and for keeping me motivated during the challenging moments of this journey.

Melih Kara

Karlsruhe, Germany, December 2024

1

Unveiling the Cosmos: Preparing for the Next Galactic Supernova and Exploring Dark Matter

"The universe is under no obligation to make sense to you."
–Neil deGrasse Tyson

Contents

1.1 Challenges in Particle Physics and Stellar Evolution

The Standard Model of Particle Physics has been highly successful in explaining the fundamental particles (shown in Figure 1.1) and the three fundamental forces of the Universe, accurately predicting numerous experimental outcomes [1, 2, 3]. However, it remains incomplete. It does not account for dark matter (hereafter, DM), an enigmatic form of matter that, while invisible, influences the motion of galaxies and the evolution of large-scale structure in the Universe. Despite compelling astrophysical evidence, no particle in the Standard Model can account for DM, a key limitation of the framework.

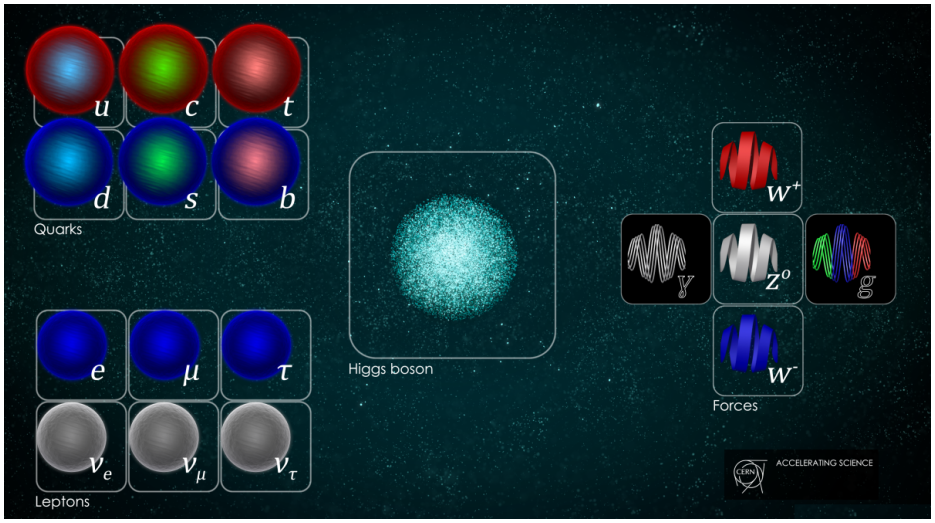

Figure 1.1: The fundamental particles that make up matter and the forces governing their interactions in the Standard Model. Matter consists of quarks and leptons, arranged in three generations of increasing mass and instability, with the lightest particles, like the up and down quarks, making up all stable matter in the Universe. The Higgs boson, responsible for giving particles mass, is centrally depicted. Force carrier particles (bosons) facilitate the interactions between matter: the photon (γ) governs the electromagnetic force, the W and Z bosons (W^+, W^-, Z^0) mediate the weak force, and gluons (g) carry the strong force. Gravity, though fundamental, is not included in the Standard Model. Image credit: Daniel Dominguez/CERN [4].

Neutrinos, another fundamental constituent of the Standard Model, also pose challenges. These nearly massless particles can oscillate between different flavors, a behavior that requires them to have mass, something the Standard Model does not fully explain [1, 2, 3]. Questions surrounding the exact nature of neutrino mass, the hierarchy of neutrino states, and the oscillations mechanisms remain unresolved. Moreover, neutrinos play a critical role in astrophysical phenomena like core-collapse supernovae (CCSNe), where they are produced in vast quantities, offering a unique window into stellar evolution and extreme matter conditions.

CCSNe serve as natural laboratories for investigating both stellar and particle physics. These colossal explosions release an immense amount of neutrinos, carrying vital information about the inner workings of the collapsing star. Detecting these neutrinos can unravel the mysteries of supernova dynamics while probing fundamental properties of neutrinos, such as mass hierarchy and behavior in ultra-dense environments.

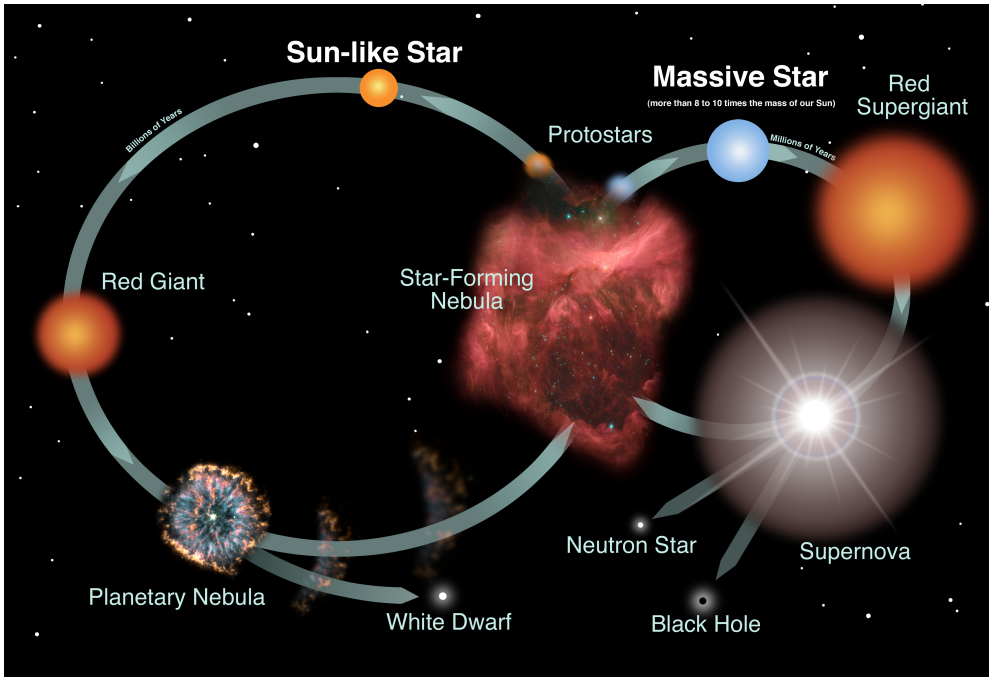

Figure 1.2: Diagram illustrating the lifecycle of Sun-like and massive stars. Stars form in interstellar molecular clouds. Sun-like (low-mass) stars evolve into Red Giants before shedding their outer layers as planetary nebulae, leaving behind a white dwarf core. In contrast, intermediate-mass and massive stars undergo core collapse at the end of their lives, often leading to a supernova explosion. The cores of the most massive stars often collapse further into a black hole whereas intermediate ones typically leave behind a neutron star. Credit: NASA and the Night Sky Network [12].

Neutrino detection is also essential in the emerging field of multimessenger astronomy, which combines data from neutrinos, gravitational waves, and electromagnetic radiation. The simultaneous detection of neutrinos, light, and gravitational waves could yield unprecedented information on the supernova event [5].

1.2 Supernova Physics and Observations

Depending on the progenitor mass, a star can end its life in various ways. Lighter stars, like our Sun, eject their outer shells in a planetary nebula event and become white dwarfs without a supernova explosion [6]. Stars with masses exceeding approximately 20 solar masses are thought to undergo a direct collapse into black holes upon the collapse of their cores, also referred to as a failed supernova [7, 8, 9]. If the progenitor's mass is around 8 to 20 times the mass of the Sun, the core of the star typically collapses under its own gravitational pull. The infalling matter on the outer shell of the star bounces off the core and emits a shock wave leading to a luminous explosion, called a supernova. Consequently, the stellar remnant left after a star collapses is called a neutron star [10, 11]. The different phases of stellar evolution are illustrated in Figure 1.2.

Various types of supernovae can occur, depending on the specific underlying mechanisms involved. The most interesting and well-studied type is the CCSN, which takes place when

the core of a massive star exhausts its nuclear fuel and exceeds the Chandrasekhar limit. The Chandrasekhar limit, named after the astrophysicist Subrahmanyan Chandrasekhar, represents the maximum mass a white dwarf (or stellar core) composed mainly of electron-degenerate matter can have before the electron-degeneracy pressure fails to counteract gravity. This limit is approximately 1.4 solar masses (1.4 $M_{\odot}$) and arises from quantum mechanical principles, specifically the Pauli exclusion principle [13, 14].

In stars below this limit, the electron degeneracy pressure, a quantum effect where no two electrons can occupy the same quantum state, provides enough support to balance gravitational forces, preventing further collapse. However, as the mass of the stellar core approaches the Chandrasekhar limit, the electrons become increasingly compressed, causing their velocities to approach the speed of light. This transition into the relativistic regime alters the relationship between momentum and velocity. In this regime, the momentum of electrons grows faster than linearly with velocity, meaning that the degeneracy pressure no longer increases quickly enough to counteract the growing gravitational force.

When the iron core of a massive star exceeds the Chandrasekhar limit, the electron degeneracy pressure can no longer provide sufficient support to hold back gravity. This results in a collapse of the core. As the core collapses, the densities rapidly increase to extreme levels, and electrons combine with protons to form neutrons, a process known as neutronization. The core rapidly becomes a dense sphere of neutrons, which may stabilize as a neutron star.

However, sudden collapse of the core creates an enormous release of gravitational energy, most of which is emitted in the form of neutrinos. This neutrino flux drives the outer layers of the star outward, triggering a powerful explosion that we observe as a supernova. The intense shock wave generated by the core collapse expels the outer layers of the star into space, enriching the surrounding interstellar medium with heavy elements forged during the star's lifetime and in the explosion itself, also known as Supernova Nucleosynthesis [15]. Elements heavier than iron, like gold and uranium, are synthesized in these extreme conditions, forming building blocks for planets and life. Fusion processes during the explosion create elements beyond those in normal stellar interiors. Expelled material from supernovae significantly contributes to the chemical evolution of galaxies. The origin of the synthesized elements is shown in Figure 1.3.

The study of supernovae offers profound insight into the fundamental processes that govern stellar evolution and the mechanisms of these explosive events. Supernovae serve as natural laboratories for studying extreme physics, such as nuclear reactions at high densities, neutrino interactions, and the behavior of matter under extreme conditions. In their comprehensive reviews, [13] and [14] emphasize the rich physics potential of supernova observations. Supernovae, particularly CCSNe, are prolific sources of neutrinos, which carry approximately 99% of the gravitational binding energy of the collapsing star. These neutrinos provide a unique window into the supernova core, which is otherwise obscured by the surrounding material, revealing perspectives on explosion dynamics and neutrino emissions.

Supernovae are also key targets for multimessenger astronomy. Signals from neutrinos, gravitational waves, and electromagnetic waves are combined to offer a comprehensive understanding of these stellar explosions. Figure 1.4 shows the timing of these signals from

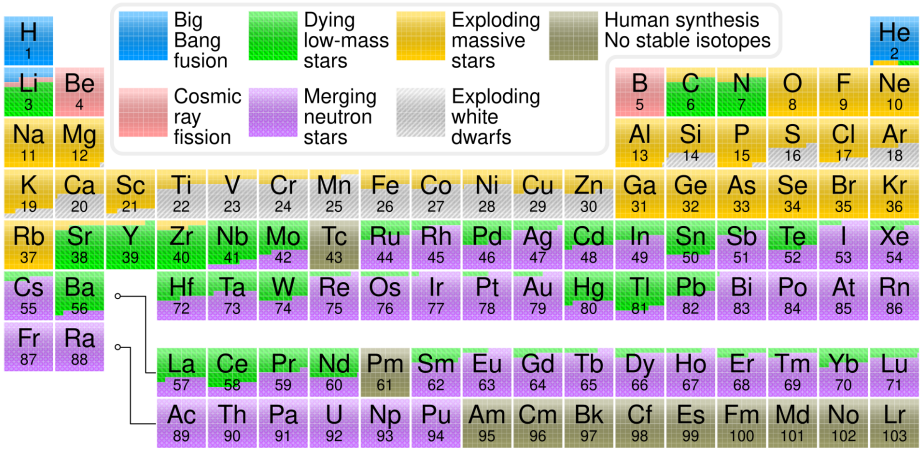

Figure 1.3: The periodic table showing the origin of chemical elements based on the data from the Northern Arizona University Meteorite Laboratory. The nucleosynthesis processes in different stellar environments, including supernovae and neutron star mergers, contribute to the cosmic abundance of elements heavier than iron. Image Credit Wikimedia Commons [16].

a simulated CCSN, emphasizing the complementarity of these messengers in understanding supernova physics [17].

Neutrino observations from supernovae also enable scientists to investigate neutrino physics, studying the masses, oscillations, and interactions with matter under extreme conditions. In [13] the importance of neutrino-neutrino interactions and their role in the dynamics of supernova explosions are discussed, while in [14] the potential for neutrino detectors to reveal new physics beyond the Standard Model is explored.

A pivotal moment in the field of supernova research occurred with the observation of SN 1987A in the Large Magellanic Cloud on February 23, 1987, triggered by the collapse of the blue supergiant star Sanduleak $-69°$ [14]. This supernova was the closest observed supernova since Kepler's Supernova in 1604 [18]. The detection of neutrinos from SN 1987A by several neutrino observatories around the world marked the dawn of neutrino astronomy. These detections provided critical data on the timing, energy spectrum and flux of neutrinos emitted during the explosion, helping to constrain the exotic neutrino properties [14]. A handful of neutrino events were observed by multiple detectors: Kamiokande II (water Cherenkov) detected 11 events [19], the Irvine-Michigan-Brookhaven (IMB) experiment (water Cherenkov) recorded 8 events [20], and the Baksan Neutrino Observatory (crystal scintillator) observed 5 events [21]. Furthermore, the Mont Blanc Underground Neutrino Observatory (Liquid Scintillator Detector, LSD) reported a few events [22], although these detections remain a subject of debate since the reported times were some hours earlier than the other observatories.

The events observed by these detectors, spanning approximately 12 seconds, are illustrated in Figure 1.5. The data from SN 1987A continue to serve as a benchmark for understanding supernova mechanisms and inform current and future observational strategies.

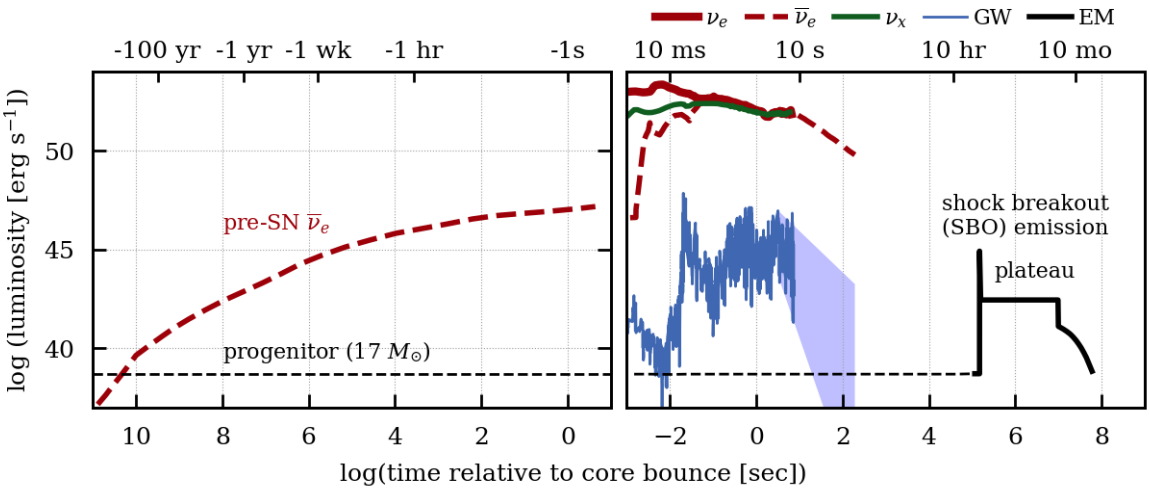

Figure 1.4: Time sequence of signals from neutrinos (ν_e, $\bar{\nu}_e$, ν_x, where ν_x represents heavy lepton neutrinos) shown in solid red, dashed red and solid green lines respectively, gravitational waves (GW) shown in blue, and electromagnetic (EM) radiation shown in black, for a CCSN of a 17 $M_\odot$ progenitor. Solid lines indicate results from CCSN simulations, while dashed lines represent estimates from literature. The x-axis shows the time before and after core bounce, highlighting the pre-CCSN neutrino emissions, core-collapse neutrino burst, and various phases of EM signal emission. The height of the curves does not reflect the total energy output. Figure taken from [17].

1.2.1 A Brief Overview of Galactic Supernovae

Every second, approximately 10 stars end their lives in a supernova explosion within the visible Universe [14]. These cosmic explosions are among the most energetic events in the cosmos, with a single supernova releasing about 10^{46} joules (10^{53} ergs) of energy, equivalent to the Sun's total lifetime energy output compressed into just a few seconds [13, 24]. Despite their frequency on a universal scale, observations within our own Milky Way are far less common due to its relatively small size and our limited vantage point. The first known observation of a supernova in the Milky Way dates back to 185 AD, recorded by Chinese astronomers [25]. This event, now referred to as SN 185, marked one of the first documented instances in which humanity witnessed the death of a massive star.

Throughout history, several significant supernovae have been visually observed in the Milky Way, including SN 1006 [26], SN 1054 [27], which led to the formation of the Crab Nebula, and SN 1604, also known as Kepler's supernova [28]. These ancient recordings, documented by astronomers and scholars of the time, provided early evidence of such cataclysmic events. Although they lacked the scientific instrumentation available today, these historical observations have nonetheless enriched our understanding of stellar evolution. By estimating the frequency of such occurrences and associating the locations of these past events with present-day chemical measurements, we have gained valuable insights into the immense impact supernovae have on the interstellar medium and the chemical enrichment of galaxies.

Despite these early records, the rate of observable supernovae in the Milky Way remains quite low. Recent literature estimates that the rate of CCSNe in a galaxy is approximately 1

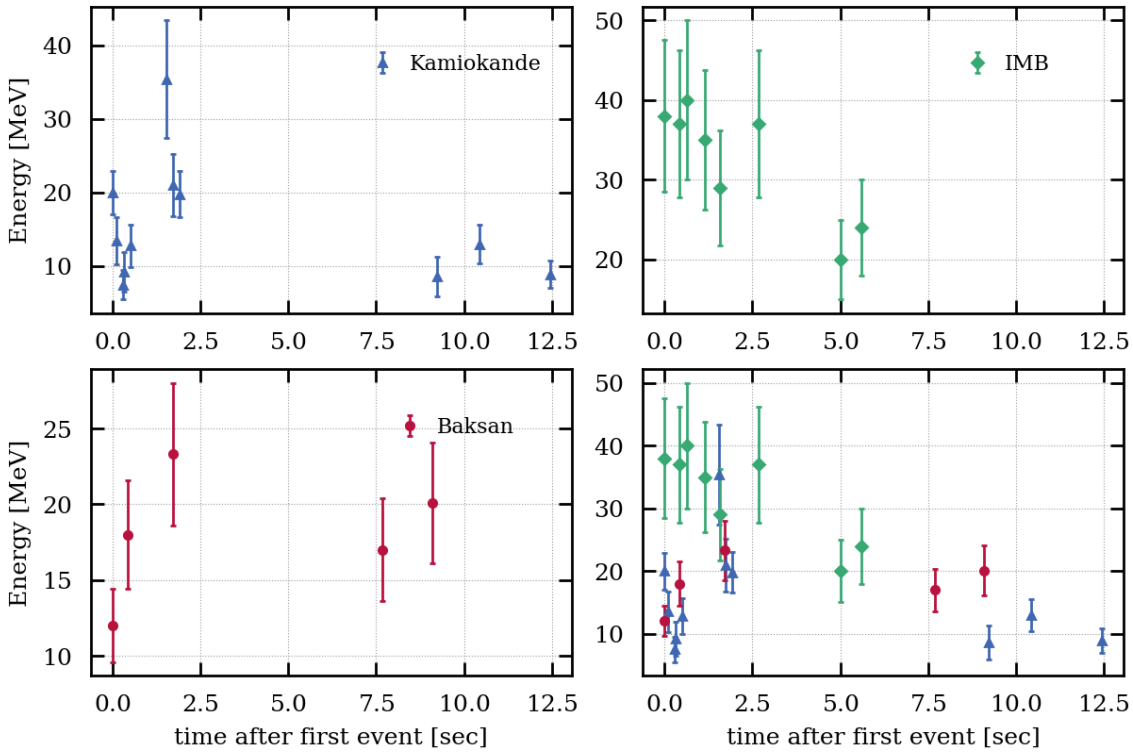

Figure 1.5: Time and energy of neutrino events detected by Kamiokande II detector (blue triangles), IMB detector (green diamonds), and Baksan detector (red circles) during SN 1987A. The bottom right panel shows the combined dataset. Figure remade using the data published in [23, 21]

to 3 per century [29, 30, 31]. In spiral galaxies, the stellar density can serve as a proxy for the probability of supernova progenitors [32]. Using this approach, [31] models the Milky Way as well as the Large and Small Magellanic Clouds to estimate the likelihood of observing supernovae within our vicinity, as illustrated in Figure 1.6. However, due to dust obscuration and the limited observational capabilities within the Milky Way, not all of these events can be directly observed. This scarcity in direct observations highlights the importance of developing alternative methods for detecting and studying supernovae, particularly through neutrinos and other multimessenger signals, such as gravitational waves.

1.2.2 Different Phases of a Supernova Collapse

The explosion of a massive star as a supernova is a complex, multiphase process that occurs over a short timescale, typically transforming the star's core into a neutron star. These phases are crucial for understanding the mechanisms of the explosion and the associated neutrino emission, providing a direct probe of the processes deep within the core. Figure 1.7, taken from [33], provides a visual overview of these phases, which are described in detail below.

During the late stages of a massive star's life, the outer shell undergoes fusion-induced burning of silicon, producing heavier elements like nickel and iron, which sink to the core.

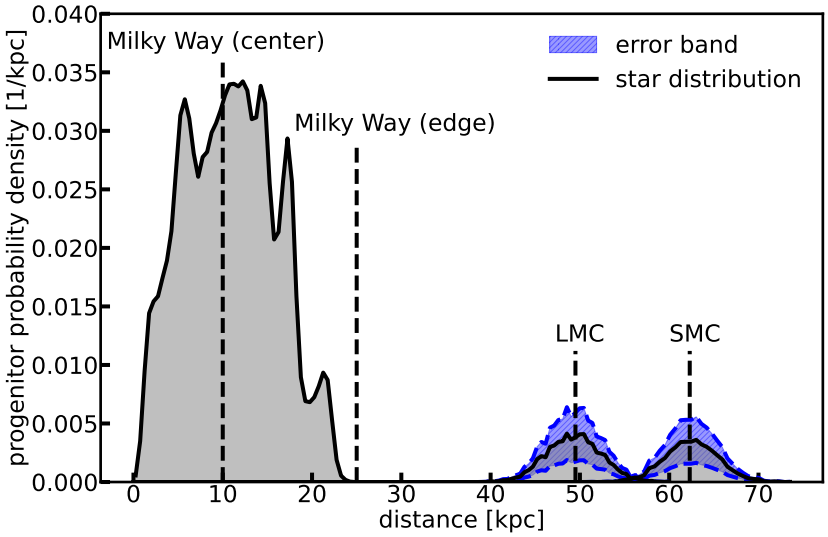

Figure 1.6: Probability distribution of progenitor stars within the Milky Way galaxy. The regions of high density within the satellite galaxies, Large Magellanic Cloud (LMC) and Small Magellanic Cloud (SMC) are also shown with their associated error bands. Figure taken from [31].

At this stage, pre-supernova neutrinos are emitted. Although these neutrinos are inherently fainter compared to those generated during the subsequent supernova explosion, they are detectable by detectors such as KamLAND, Super-Kamiokande, and other detectors currently under construction or proposed. Consequently, they can function as early indicators of the imminent supernova [14, 34, 35]. At this point, the pressure from the degenerate electron plasma and the fusion processes in the outer shell maintain hydrostatic equilibrium against the gravitational pressure. Once the silicon burning phase concludes and the core's mass approaches the Chandrasekhar mass limit, the degenerate electron pressure can no longer sustain equilibrium, and the core collapse begins as the core density, ρ_c exceeds a few times 10^{11} g cm^{-3}.

In the first few milliseconds of the collapse, a sudden increase in density and temperature dissociates the iron-group elements, leaving abundant nucleons for electron capture. As a result, neutronization occurs through electron capture,

$$e^- + p \rightarrow n + \nu_e, \tag{1.1}$$

emitting vast amounts of electron neutrinos, shown in Figure 1.7**(a)**. This is the first clear indication that the collapse has begun [24]. The burst arising from the neutronization process provides the first observable signal of a core collapse.

Approximately 100 milliseconds later, when densities reach $\rho_c \approx 10^{12}$ g cm^{-3}, the mean free path of neutrinos becomes so short that they cannot escape freely from the core. At this point, the neutrinos become "trapped" because their mean free path is extremely short. The dense matter prevents them from escaping freely and they reach thermal equilibrium with their surroundings, exchanging energy through frequent interactions with nucleons and electrons; see Figure 1.7**(b)** [36]. Once neutrino trapping occurs, the core is opaque to neutrinos, and any further neutrino emission that originates from deeper within the core is

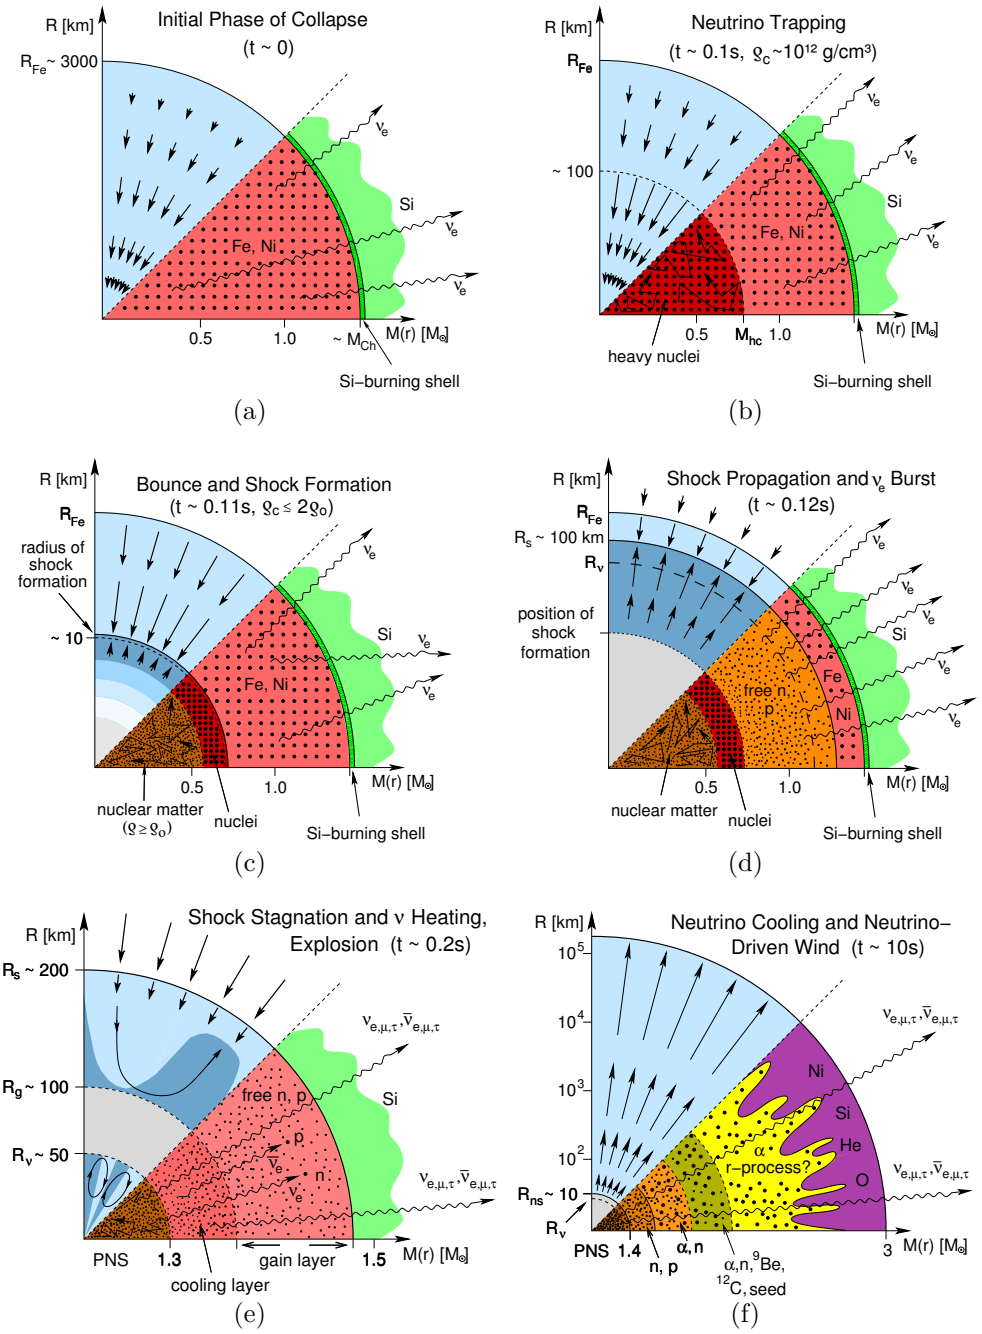

Figure 1.7: Different phases of the supernova explosion and neutrino emission mechanisms. Panels a to f illustrate the initial collapse phase, neutrino trapping, core bounce and shock formation, shock propagation with ν_e burst, shock stagnation, and neutrino heating leading to the explosion, and finally the neutrino cooling and neutrino-driven wind phases. Radial distances $R_{\rm Fe}, R_S, R\nu, R_g$ and $R_{\rm ns}$ correspond to iron–core radius, shock radius, neutrinospheric radius, gain radius (defined as the radius that separates cooling and heating layers) and proto–neutron star raidus. Refer to the text for detailed descriptions. Figure adapted from [33].

delayed until the core bounce and subsequent phases.

As the core continues to collapse, it becomes increasingly dense until it reaches $\rho_c \approx 2{-}3{\times}10^{14}$ g cm^{-3}, the nuclear saturation density ($\rho_0 \approx 2.7{\times}10^{14}$ g cm^{-3}). At nuclear densities, repulsive

forces between nucleons prevent further collapse, as shown in Figure 1.7**(c)**. This sudden halt causes the core to *bounce*, creating a shock wave that moves outward through the surrounding layers [24].

As the shock propagates through the star's layers, it disintegrates the iron nuclei into nucleons—protons and neutrons (Figure 1.7**(d)**). During this process, electron neutrinos are produced in large numbers via electron capture, but they can only escape when the shock reaches regions with lower densities. This results in a neutrino burst being released [13]. However, as the shock propagates, it weakens and the energy carried away by the generated electron neutrinos further diminishes its strength. Eventually, the shock stops and can no longer push through the infalling material.

As the core heats up beyond 100 MeV, positrons and electrons are created through pair production, with positrons capturing neutrons and emitting anti-electron neutrinos. These high temperatures also allow for the production of muon and tau neutrinos and their antineutrinos, collectively known as heavy-lepton neutrinos. The intense neutrino flux generated in the core re-energizes the shock front, reviving its propagation and ultimately leading to the ejection of the outer layers in a supernova explosion about $\sim$ 200 milliseconds after the initial neutronization burst, illustrated in Figure 1.7**(e)** [24, 33].

Following the explosion, the hot, dense proto-neutron star (PNS) begins to cool (see Figure 1.7**(f)**). During this cooling phase, the PNS continues to radiate intense fluxes of neutrinos, which carry away most of the remaining energy of the core. These neutrinos carry approximately 99% of the gravitational binding energy of the PNS, which amounts to around 10^{53} ergs [36, 13]. As the PNS cools, neutrinos stream out from its surface, heating the outer layers and generating a continuous wind-like flow. This wind, composed of protons, neutrons, and alpha particles, is crucial for nucleosynthesis processes. The neutrino-driven wind generates many heavy elements such as gold, platinum, and uranium in the Universe [37, 38]. Over time, the outflow slows and eventually ceases as the PNS transitions to a cold, stable neutron star. The neutrino luminosities as a function of time during these stages are illustrated in Figure 1.8 where each phase is also annotated.

1.2.3 Observing Supernova Neutrinos on Earth

Supernovae are important sources of neutrinos, releasing approximately 99% of their gravitational energy in the form of these particles during the moments surrounding the core collapse. Observing these neutrinos is crucial as they provide a unique window into fundamental particle physics. Neutrino detection experiments explore various interactions to observe these elusive particles, with different types of detectors optimized for different neutrino flavors and energies.

Water Cherenkov detectors, such as the Super-Kamiokande detector, consist of large volumes of ultra–pure water. When antielectron neutrinos, $\overline{\nu_e}$, interact via inverse beta decay (IBD) with protons in the target,

$$\overline{\nu_e} + p \rightarrow e^- + n, \tag{1.2}$$

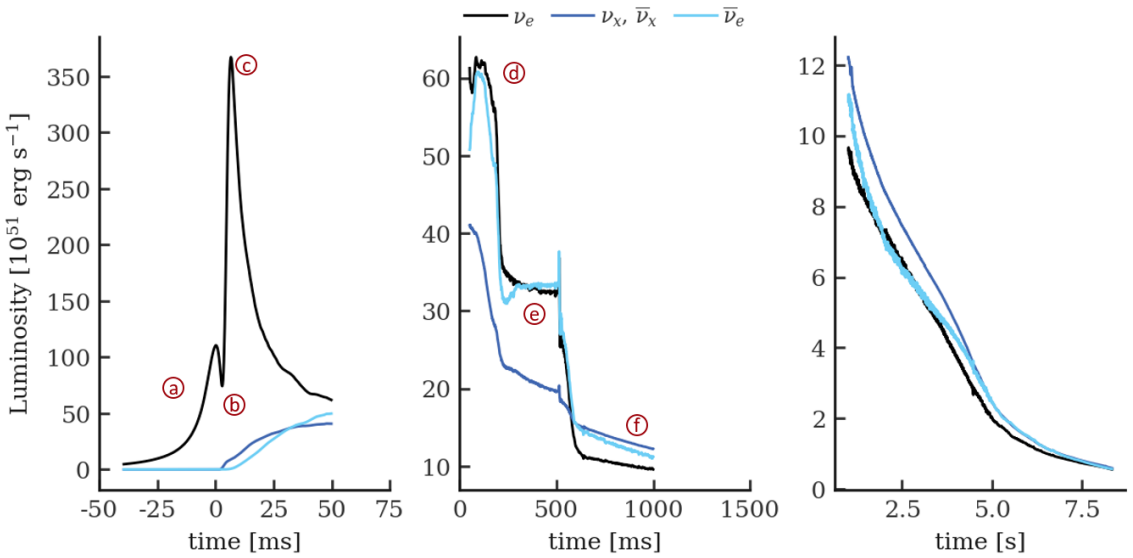

Figure 1.8: Neutrino luminosity distribution of a simulated supernova signal as a function of time after the bounce. The plot highlights the three phases of the neutrino emission process: shock-breakout, accretion phase, and Kelvin-Helmholtz Proto-Neutron Star (PNS) cooling. Note that the y-axis scale differs across the three panels. The circled numbering refers to phases shown in Figure 1.7.

the generated electrons emit Cherenkov radiation, a type of light produced when charged particles move faster than the speed of light in water. This radiation is detected by an array of highly sensitive photosensors. Liquid scintillator detectors like KamLAND [35] and the Sudbury Neutrino Observatory (SNO) [39] are also sensitive to $\overline{\nu_e}$ through IBD but in addition can detect neutrinos of all flavors via elastic scattering (where neutrinos scatter off electrons) and charged-current interactions (where neutrinos interact with nuclei), providing excellent energy resolution [40].

Another important class of detectors are Liquid Argon Time Projection Chambers (LArTPCs), such as the ICARUS detector [41] and the upcoming DUNE experiment [42]. Additionally, DM detectors like XENONnT [43] and LUX-ZEPLIN (LZ) [44], while primarily designed to search for weakly interacting massive particles (WIMPs), are also sensitive to neutrinos from supernovae, particularly through coherent elastic neutrino-nucleus scattering (CEνNS). This interaction dominates at lower energies because the entire nucleus recoils coherently, providing a unique detection method for low-energy neutrinos that complements traditional high-energy neutrino experiments [45, 46, 47]. Each detection method provides complementary information about the supernova neutrino burst, enhancing our overall understanding of both the supernova mechanism and neutrino physics.

1.3 Dark Matter: Defining the Unknown

DM, an enigmatic and nonluminous form of matter, constitutes approximately 25% of the Universe's total energy density; yet, its true nature remains one of the most profound

enigmas in contemporary physics [48, 49, 50]. Unlike ordinary matter, DM does not interact via electromagnetic forces, rendering it invisible to telescopes and undetectable through conventional means. However, its existence is robustly inferred through its gravitational influence on the large-scale structure of the Universe, the motion of galaxies, and its imprint on the cosmic microwave background (CMB) [51]. The pursuit of DM involves multiple disciplines, such as cosmology, astrophysics, and particle physics, as researchers strive to characterize its properties and its role in the formation and evolution of the Universe. A breakthrough in DM detection would not only revolutionize our comprehension of the cosmos but also unveil physics beyond the Standard Model. The following sections will explore the cosmological evidence supporting DM, the observational signatures suggesting its existence, and the experimental methodologies designed to detect this elusive constituent of the Universe.

1.3.1 Cosmological and Observational Evidences

The expansion of the Universe is described by the Hubble parameter $H(t)$ [52], linked to the Universe's energy content via the Friedmann equation for a flat Universe ($k = 0$):

$$H^2(t) = \frac{8\pi G_N}{3}\rho_{\rm tot}(t), \tag{1.3}$$

where G_N is Newton's gravitational constant, and $\rho_{\rm tot}(t)$ is the total energy density at time t, comprising matter, radiation, and dark energy:

$$\rho_{\rm tot}(t) = \rho_{\rm m}(t) + \rho_{\rm r}(t) + \rho_\Lambda(t). \tag{1.4}$$

These components evolve differently over time. The matter density $\rho_{\rm m}(t)$ scales with the expansion factor $a(t)$ as:

$$\rho_{\rm m}(t) = \rho_{\rm m,0}\left(\frac{a_0}{a(t)}\right)^3, \tag{1.5}$$

while radiation, due to redshifting, scales as:

$$\rho_{\rm r}(t) = \rho_{\rm r,0}\left(\frac{a_0}{a(t)}\right)^4. \tag{1.6}$$

The dark energy, represented by the cosmological constant Λ, remains constant over time:

$$\rho_\Lambda(t) = \rho_{\Lambda,0}. \tag{1.7}$$

Combining these components into the Friedmann equation, we obtain the time-dependent Hubble parameter:

$$H^2(t) = H_0^2\left[\Omega_{\rm m,0}\left(\frac{a_0}{a(t)}\right)^3 + \Omega_{\rm r,0}\left(\frac{a_0}{a(t)}\right)^4 + \Omega_{\Lambda,0}\right], \tag{1.8}$$

where H_0 is the present-day Hubble constant[i] and $\Omega_{\mathrm{m},0}$, $\Omega_{\mathrm{r},0}$, and $\Omega_{\Lambda,0}$ are the present-day density parameters for matter, radiation, and dark energy, respectively, defined as the fraction $\Omega_i = \rho_i/\rho_{\mathrm{cr}}$, where ρ_{cr} is

$$\rho_{\mathrm{cr}} = \frac{3H_0^2}{8\pi G_N}. \tag{1.9}$$

Observations suggest that the Universe is flat, meaning that $\Omega_{\mathrm{tot}} = 1$, with approximately 30% of the energy density coming from matter ($\Omega_{\mathrm{m},0} \approx 0.3$), and 70% from dark energy ($\Omega_{\Lambda,0} \approx 0.69$). Of the matter component, only a small fraction is baryonic ($\Omega_{\mathrm{b},0} \approx 0.048$) [51], leaving the bulk of matter as DM ($\Omega_{\mathrm{DM}} \approx 0.25$) showing that DM plays a critical role in shaping the structure and evolution of the Universe. During the expansion of the Universe, particularly around the epoch of matter-radiation equality, DM played a pivotal role in large–scale structure formation. At this time, DM behaved as a cold and nonrelativistic fluid, which means that its particles traveled at speeds much lower than the speed of light ($v \ll c$). This low velocity allowed DM to clump together under the influence of gravity, facilitating the formation of galaxies and galaxy clusters. If DM particles had been moving at relativistic speed, their high velocities would have prevented such large–scale structures from forming, as the particles would have dispersed too quickly. Therefore, the cold and nonrelativistic nature of DM during this phase of cosmic expansion was crucial for the development of the large-scale structure of the Universe [53] as demonstrated by the N–body simulations of the evolution of the Universe [54, 55].

Another important feature of DM is its lack of significant interactions with itself or ordinary matter, other than through gravity making it *collisionless* or *noninteracting* [56, 53]. This is a sharp contrast to baryonic matter, which interacts dominantly through electromagnetic forces, allowing it to cool and form structures like stars and galaxies. DM, in contrast, seems to be unable to cool through radiation, consequently exhibiting a behavior different from baryonic matter throughout the evolution of the Universe.

DM, if it is a particle, must be stable or have a decay time much longer than the age of the Universe, which is approximately 13.8 billion years or 4.35×10^{17} seconds, to persist through history [50]. Along with its stable and noninteracting properties, DM is thought to exhibit adiabatic inhomogeneities. This implies that initial density fluctuations in DM mirrored those in ordinary matter and photons, originating from quantum effects during early Universe inflation. These fluctuations were pivotal in the formation of today's large-scale structures, as the interaction between DM and ordinary matter in these early phases was essential for the formation of galaxies, clusters, and the cosmic web [53].

The gravitational effects of DM manifest across a wide range of astronomical scales, from the rotation curves of galaxies to the distribution of galaxy clusters and the formation of the cosmic web. These gravitational signatures provide some of the strongest support for the existence of DM. Furthermore, cosmological probes such as the cosmic microwave background (CMB), baryon acoustic oscillations (BAO), and gravitational lensing all indicate a dominant

[i]The Hubble constant, H_0, is commonly expressed as $H_0 = 100\,h\,\mathrm{km/s/Mpc}$, where h is a dimensionless scaling factor. As a result, density parameters Ω_i are often quoted as $\Omega_i h^2$, which provides a value independent of the precise value of H_0, making it more convenient to compare cosmological models.

DM component that has shaped the evolution of the Universe. The observational evidence for the existence of DM can be summarized as follows:

- **Galactic Rotation Curves**: One of the earliest and most compelling pieces of evidence for DM comes from the study of galactic rotation curves. Observations of spiral galaxies, such as those by Vera Rubin in the 1970s [57], revealed that the rotational velocities of stars remain constant or even increase with distance from the galactic center, contrary to expectations based on visible matter alone. According to Newtonian mechanics, the outer stars should exhibit decreasing rotational speeds if only the visible mass was present. This discrepancy suggests that galaxies are embedded in massive halos of DM, extending far beyond the visible disk [58].
- **Galaxy Cluster Collisions**: Galaxy cluster collisions provide a unique laboratory for studying DM. One of the most famous examples is the Bullet Cluster, where observations of the hot gas (detected via X-rays) and gravitational lensing maps reveal a clear separation between visible matter and the bulk of the mass [59]. During the collision, the intergalactic gas is slowed down by ram pressure, while the DM, which interacts only gravitationally, passes through largely unaffected. This separation provides direct evidence for the existence of DM as a component distinct from ordinary matter. The illustration of the interactions within the Bullet Cluster is shown in Figure 1.9.

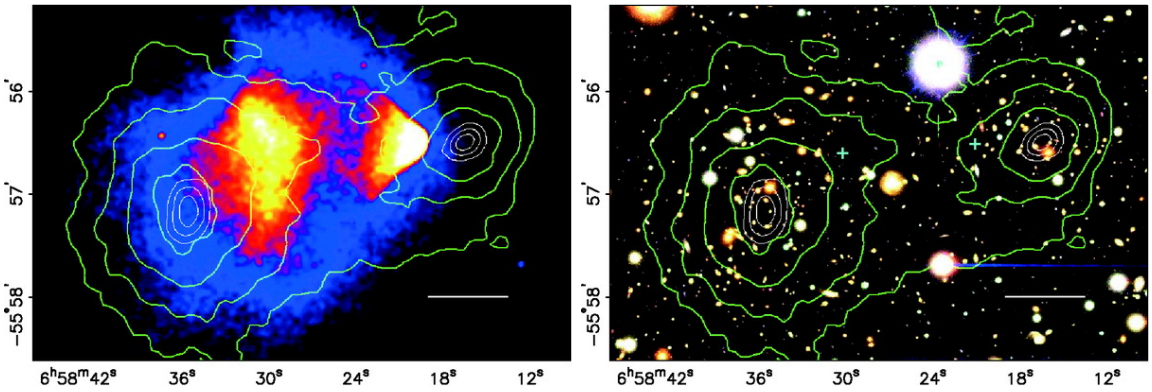

Figure 1.9: Bullet Cluster (1E0657-558) [56], illustrating the merging of two galaxy clusters. The white bar at the bottom right of each panel represents 200 kpc at the distance of the cluster. The green contours highlight the distribution of the total mass, reconstructed through weak gravitational lensing, indicating the presence of DM. In the left panel, the color image beneath the contours displays X-ray data from the Chandra telescope, revealing the position of the hot gas, which constitutes approximately 80% of the total baryonic mass. This figure clearly demonstrates that while the DM halos, traced by the green contours, pass through each other without interaction, the collisional gas (in red and blue) lags behind due to its interaction during the merger. The collisionless stars and galaxies follow the gravitational attraction of DM. Figure adapted from [59].

- **Gravitational Lensing**: Gravitational lensing offers another compelling line of evidence for DM. This phenomenon occurs when a massive object, such as a galaxy cluster, bends the light from background objects, creating multiple images or distortions of the background source. The degree of light bending allows to map the mass distribution of the lensing object. Studies consistently show that the visible mass of galaxies and clusters is insufficient to account for the observed lensing effects, implying a substantial presence of DM [60].

- **Cosmic Microwave Background (CMB)**: The CMB is the remnant radiation from the Big Bang, offering a snapshot of the Universe approximately 380,000 years after its inception, when the first neutral atoms formed and photons decoupled from matter. First detected in the 1960s, CMB carries imprints of the early conditions of the Universe and provides invaluable clues about its composition and evolution [61].

  Measurements of the anisotropies (temperature fluctuations) in the CMB, particularly by the Planck satellite [51], reveal how structures in the Universe formed and evolved. DM plays a crucial role in this process by affecting the gravitational potential wells in which baryonic matter accumulates. The patterns of these temperature fluctuations, when plotted as a function of angular scale (or multipole moment, ℓ), form what is known as the CMB power spectrum. The power spectrum quantifies how much the temperature fluctuations vary with angular scale and is characterized by a series of peaks that correspond to the sound waves in the early Universe's photon-baryon fluid.

  The gravitational effects of DM are crucial in determining the height and position of these peaks, providing an accurate calculation of the abundance of DM relative to ordinary matter. The CMB power spectrum reveals that DM constitutes about five times more mass than baryonic matter, with the precise ratio determined by fitting cosmological parameters such as the baryon density (Ω_b), DM density (Ω_{DM}), and the Hubble constant (H_0) to the observed data.

  Changes in cosmological parameters affect the shape and height of the CMB power spectrum, with the positions of the peaks corresponding to sound horizon scales and the overall shape determined by the interaction between radiation and matter at the time of decoupling. The gravitational influence of DM, in particular, increases the amplitude of certain peaks, thereby providing strong evidence for its existence. These temperature anisotropies, combined with other cosmological probes, indicate that DM constitutes about 25% of the total energy density of the Universe. A visual representation of how changes in density parameters affect the CMB temperature spectrum is shown in Figure 1.10.

- **Large-Scale Structure**: On cosmological scales, the spatial distribution of galaxies and galaxy clusters throughout the Universe provides compelling evidence for the existence of DM. The large-scale structure observable today is thought to have originated from primordial density fluctuations. In this context, the gravitational influence of DM served as a framework or scaffolding around which ordinary baryonic matter could aggregate, subsequently leading to the formation of stars, galaxies, and other cosmic structures. Detailed numerical simulations of the evolution of the Universe that include DM components have shown a remarkable concordance with the empirically observed distribution of galaxies and the intricate structure of the cosmic web. This strong agreement underscores the critical role of DM in contemporary cosmological models and theories [63].

The cumulative evidence from galactic rotation curves, galaxy cluster collisions, gravitational lensing, the cosmic microwave background, and large-scale structure consistently understood when introducing DM. These observations converge on the conclusion that DM is a nonluminous and noninteracting substance, distinct from ordinary matter, yet crucial for

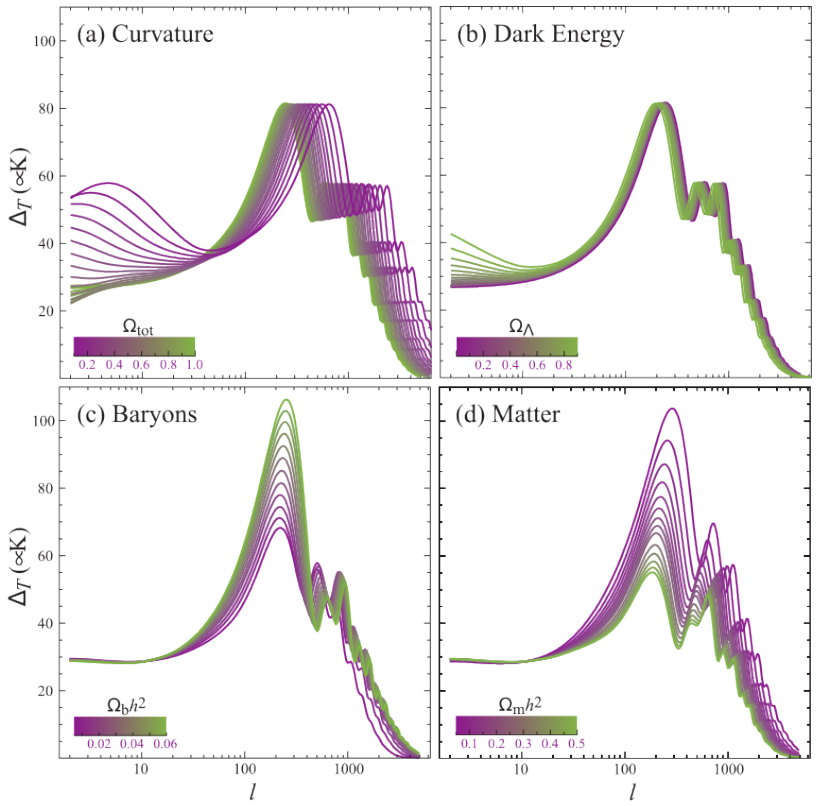

Figure 1.10: Acoustic temperature spectrum of the CMB with varying cosmological parameters. The figure illustrates the sensitivity of the acoustic peaks in the CMB temperature spectrum to different cosmological parameters. Each panel shows the temperature power spectrum as a function of the multipole moment ℓ, highlighting the effects of changes in the parameters. The upper left panel illustrates the effect of the curvature parameter, where higher total density values shift the peaks to higher ℓ, while the upper right panel shows how the dark energy density Ω_Λ changes the spectrum, with lower values shifting the spectrum towards higher multipole moments and causing a larger power on smaller angular scales. The bottom left panel demonstrates the impact of varying the baryon density $\Omega_b h^2$, with higher baryon densities increasing the amplitude of the odd peaks relative to the even ones. The bottom right panel shows how the DM density (displayed as the total matter density $\Omega_m h^2$) affects the peak spacing, with higher DM densities shifting the peaks to smaller angular scales (higher ℓ). This figure highlights the key role that the observation of CMB anisotropies plays in constraining fundamental cosmological parameters. Figure taken from [62].

understanding the composition and evolution of the Universe. Direct detection experiments, utilizing advanced technologies and sensitive detectors, aim to unveil the particle nature of DM, potentially revolutionizing our understanding of the cosmos. Different approaches to detect DM are discussed below.

1.3.2 Detecting DM: Direct, Indirect and Collider Searches

The particle nature of DM remains elusive, with several theoretical candidates proposed to explain its nonluminous and noninteracting behavior. One of the most compelling candidates is the WIMP, which emerges naturally in many extensions of the Standard Model of particle physics, such as supersymmetry [64]. WIMPs are appealing because they would have been

thermally produced in the early Universe, with a relic abundance that coincides with the observed DM density today, a concept often referred to as the "WIMP miracle" [65].

Beyond WIMPs, other candidates have been proposed, including axions, sterile neutrinos, and primordial black holes. Axions, originally introduced to solve the strong charge–parity problem in quantum chromodynamics (QCD), offer another plausible explanation for DM, particularly in the case of ultralight particles [66]. Sterile neutrinos, on the other hand, extend the neutrino sector by adding a right-handed component, providing a mechanism for DM if these particles are heavy and weakly interacting [67]. Primordial black holes, hypothesized to have formed in the early Universe, present a more exotic candidate, but recent constraints limit their contribution to the DM density [68, 69].

Among these candidates, WIMPs remain the focus of a majority of experimental searches because of their strong theoretical motivation and the potential for detection. These particles are predicted to interact with regular matter via weak-scale interactions, making them detectable through a variety of methods, including direct detection, indirect detection, and searches at particle colliders.

These three primary strategies for detecting DM as a particle are illustrated in Figure 1.11. Direct detection experiments search for the interaction of DM with atomic nuclei or bound electrons, while indirect detection focuses on identifying the Standard Model particles produced from DM annihilation or decay in high-density regions of the Universe, such as the galactic center. Collider searches, on the other hand, attempt to produce DM particles in high-energy collisions of Standard Model particles. The figure highlights these approaches and the central, yet unknown, interaction mechanism that connects DM with Standard Model particles.

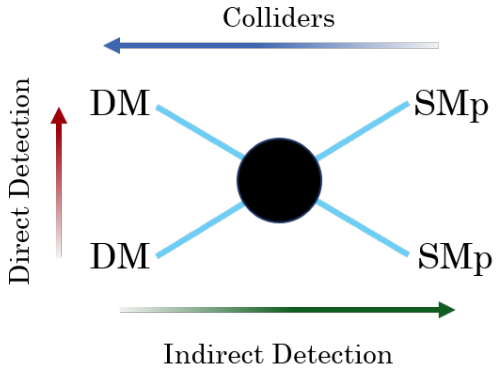

Figure 1.11: Detection strategies for DM as a particle. The green arrow represents indirect detection, which searches for Standard Model particles (SMp), such as gamma rays, that could be produced by the annihilation or decay of DM particles in regions of high DM density, such as the center of the Milky Way. The blue arrow shows collider searches, where DM particles might be produced through the interactions of Standard Model particles, a process that is the reverse of indirect detection. The red arrow illustrates direct detection, where experiments aim to observe the recoil or other signals resulting from DM particles interacting with Standard Model particles. The central black circle symbolizes the unknown interaction mechanism responsible for these processes, which remains to be discovered.

Direct detection experiments aim to measure the rare interactions of WIMPs with atomic nuclei, where DM particles scatter elastically off a target nucleus (or bound electrons),

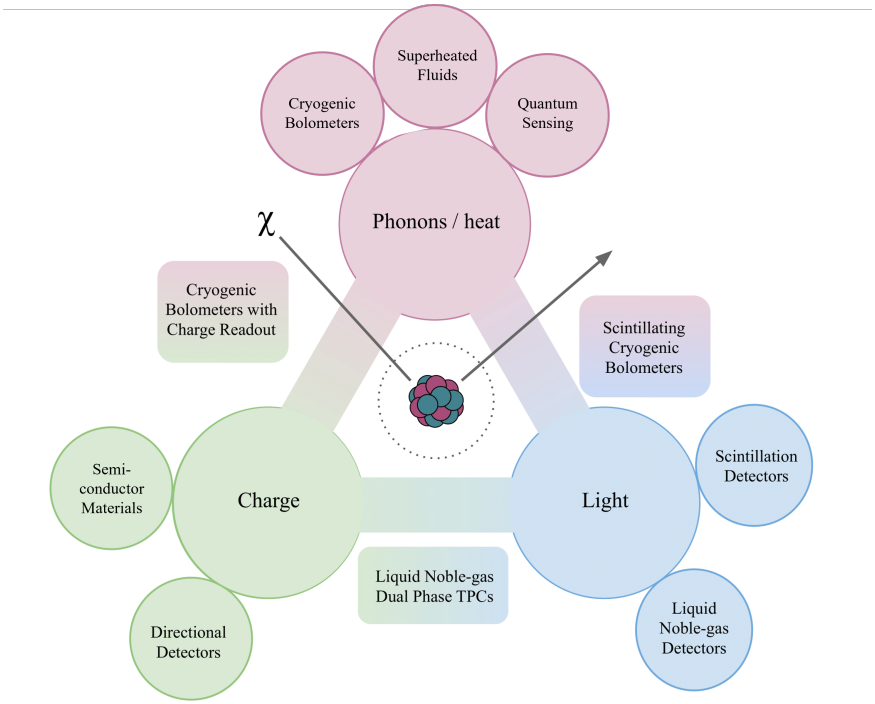

Figure 1.12: Overview of the primary detection strategies for WIMP DM interactions. WIMPs (χ) can transfer energy to a nucleus, resulting in the production of phonons/heat, charge, or light, which are measured by different detection technologies. Phonons/heat are detected by cryogenic bolometers, superheated fluids, and quantum sensing methods. Charge is measured using semiconductor materials, directional detectors, and liquid noble-gas dual-phase Time Projection Chambers (TPCs). Light (scintillation) signals are captured by scintillation detectors and liquid noble-gas detectors. Some detectors, such as scintillating cryogenic bolometers and cryogenic bolometers with charge readout, combine multiple signals to enhance sensitivity and background discrimination. Figure taken from [76].

transferring a small amount of energy in the process. Recent leading experiments include XENONnT [70], LUX–ZEPLIN [44], and PandaX [71], and future planned experiments such as PandaX-4T [72], DARWIN [73], and XLZD [74, 75], which employ large volumes of liquid xenon in dual-phase TPCs to detect WIMP interactions. These detectors are pushing the limits of sensitivity, particularly in the low-mass WIMP region. As illustrated in Figure 1.12, these experiments typically focus on two of the three possible signals from a WIMP interaction: phonons (heat), charge, and light. Each detection method is designed to observe one or more of these signals, taking advantage of different physical principles to maximize sensitivity to potential DM interactions.

In the phonon/heat detection approach, cryogenic bolometers and superheated fluids are used to measure the tiny amount of heat or vibrational energy (phonons) produced when a WIMP scatters off a nucleus. The Cryogenic Dark Matter Search (SuperCDMS) experiment [77] and the CRESST-III experiment [78] are examples of leading efforts in this category. SuperCDMS uses cryogenic germanium and silicon detectors, while CRESST-III uses scintillating calcium tungstate and sapphire crystals. This signal is often extremely small, requiring highly sensitive detectors cooled to mK–scale temperatures to minimize thermal noise. Advances in quantum sensing techniques offer promising improvements in the detection of these minute energy deposits, thus enhancing the sensitivity of this approach [79].

Charge detection focuses on measuring the ionization produced by a WIMP interaction. When a WIMP scatters off a nucleus, it can liberate electrons from the atoms of the target material, generating a measurable electric charge. Technologies such as semiconductor detectors and liquid noble–gas dual–phase TPCs are specifically designed to capture this ionization signal. TPCs, in particular, have gained widespread use because of their ability to measure both the charge and light signals from interactions, providing a powerful way to distinguish DM signals from background noise and scalability.

Light detection, or scintillation, occurs when the energy transferred in a WIMP interaction causes the target material to emit light. Scintillation detectors and liquid noble-gas detectors rely on this principle, where materials like xenon or argon are used to detect scintillation light. These detectors are highly sensitive to nuclear recoils and can provide additional discrimination against background events by measuring the light output alongside the charge or heat.

Certain technologies, including scintillating cryogenic bolometers (measuring light and heat), cryogenic bolometers with charge readout (measuring charge and heat), and liquid noble-gas dual-phase TPCs (measuring light and charge), integrate these detection mechanisms to facilitate concurrent measurements of two out of three possible signals. This approach of using multiple signals significantly enhances the ability to differentiate between possible WIMP interactions and background events primarily caused by α, β and γ radiation, which typically results in a recoiling electron rather than a recoiling nucleus.

Looking further ahead, next–generation dark matter experiments such as DARWIN and XLZD, with their large target masses and advanced detection technologies, are anticipated to detect supernova neutrinos at rates comparable to dedicated neutrino observatories [47]. These detectors will enhance our ability to study supernovae, offering valuable insight into the early stages of core collapse and the mechanisms behind these extraordinary astrophysical events.

The next chapter focuses on the principles of liquid noble–gas dual–phase TPCs, with a particular emphasis on xenon technology and the XENONnT experiment. This lays the groundwork for subsequent discussions on the advancements in dark matter detection and the integration of XENONnT into the Supernova Early Warning System (SNEWS). In the following chapters, the mechanisms of supernova neutrino emission, their expected signals in dual–phase TPCs, and the methodologies for their detection are explored. The analysis framework developed for identifying supernova neutrino interactions within XENONnT is also presented, highlighting the experiment's sensitivity under various scenarios. These investigations demonstrate the potential of XENONnT and similar detectors to contribute significantly to multimessenger astrophysics.

2

Liquid Xenon Time Projection Chambers and the XENONnT Experiment

"Equipped with his five senses, man explores the universe around him and calls the adventure Science."
–Edwin Hubble

Contents

2.1 Xenon as a target

Noble gases such as argon and xenon are frequently used in particle detection due to their high purity and their ability to ionize and scintillate when interacting with particles [80]. Xenon is ideal for DM searches because of its high atomic mass, dense liquid state, and self-shielding properties, making it sensitive to rare particle interactions. The XENON project, with iterations like XENON100 [81], XENON1T [82], and XENONnT [70], has evolved to increase the sensitivity to DM through improvements in detector size and design. This chapter covers unique properties of the xenon element for DM detection, dual-phase xenon technologies, and how XENONnT fits into the broader search for DM.

2.1.1 Xenon Technology

With a high atomic number and a density of 2.862 g/cm^3 in its liquid phase, xenon is an ideal target for DM searches. Its atomic mass of approximately 122.3 GeV/c^2 is well-matched to WIMP-nucleon cross-sections, covering predicted DM mass ranges from a few GeV/c^2 to several TeV/c^2 [83]. Additionally, xenon's range of isotopes, both even and odd numbered, enables DM searches via spin-dependent and spin-independent interaction mechanisms [83].

When hit by a particle, xenon emits both scintillation light and electrons. These two signals provide critical information about the interaction, such as the energy deposited and the interaction type. With high light and charge yields, xenon detectors achieve an energy resolution of around 1% at the MeV scales [84]. Its liquid state at $T \approx 175\,\mathrm{K}$ allows scalable detector designs under cryogenic conditions that are relatively straightforward to maintain.

Different xenon isotopes (see Table 2.1) also allow investigation beyond DM. For example, ^{136}Xe and ^{124}Xe are used in studies of neutrino properties, including two-neutrino double electron capture [85] and neutrinoless double beta decay [86, 87].

Table 2.1: Properties of naturally occurring xenon isotopes relevant to DM detection and neutrino studies. The table includes the mass number, nuclear spin, and natural abundance of each isotope. Values are taken from [88].

Isotope	Mass Number	Spin	Abundance [%]
^{124}Xe	124	0	0.09
^{126}Xe	126	0	0.09
^{128}Xe	128	0	1.92
^{129}Xe	129	1/2	26.44
^{130}Xe	130	0	4.08
^{131}Xe	131	3/2	21.89
^{132}Xe	132	0	26.89
^{134}Xe	134	0	10.44
^{136}Xe	136	0	8.87

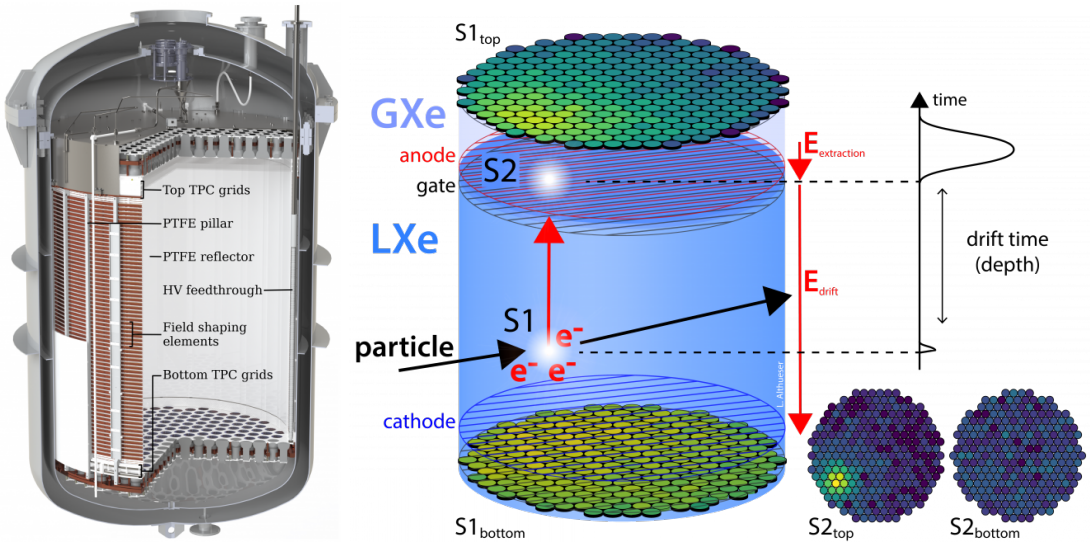

Figure 2.1: Schematic of the XENONnT TPC and illustration of the detection principle. The left panel displays the TPC structure, highlighting key components such as PMT arrays, electrodes, and reflectors. The right panel illustrates a particle interaction in the LXe target, where a prompt scintillation signal is generated, followed by the drift and extraction of free electrons producing a delayed scintillation signal in the gaseous phase. The timing difference between S1 and S2 (drift time) and light patterns in the PMTs are also depicted.

Furthermore, xenon's high density provides significant self-shielding against external radiation, allowing for a "fiducialization" approach. This method isolates a low-radioactivity inner volume within the detector, enhancing sensitivity by reducing background noise. Together, these features make xenon a desired target for DM detection.

2.1.2 Xenon Dual-Phase Time Projection Chambers

Dual-phase detection technologies play a crucial role in particle identification within DM detectors. Liquid xenon (LXe) dual-phase detectors, particularly TPCs, are widely used in DM searches due to their high sensitivity and efficient background rejection [44, 89, 90] and in neutrinoless double beta experiments [91]. This design utilizes LXe as the primary detection medium, complemented by a gaseous xenon layer to facilitate secondary signal generation after charge extraction from the liquid phase to the gaseous phase.

A typical TPC setup, such as the XENONnT TPC shown in Figure 2.1, includes an inner volume of LXe surrounded by vacuum ultra-violet (VUV) light reflectors for optimal capture of scintillation light. The detector is equipped with two arrays of photomultiplier tubes (PMTs) positioned at the top and bottom of the vessel to collect light signals produced during particle interactions. There are three electrodes: the cathode, the gate, and the anode. The cathode is positioned at the bottom, the gate is located just beneath the liquid-gas boundary, and the anode is located at the top. These electrodes generate electric fields that drift ionized electrons in the TPC to the liquid–gas surface and extract them from the LXe to the gaseous xenon layer, where a secondary signal is generated.

The detection process in a TPC begins with a particle that interacts with xenon atoms, which can lead to ionization and excitation of xenon. This interaction produces an immediate, or *prompt*, light signal called the prompt scintillation signal, **S1**. The freed electrons, driven by the applied drift field, move towards the liquid–gas surface, where they are extracted by the extraction field. Upon entering the gaseous phase, these electrons interact with xenon atoms, resulting in a secondary (*proportional*) scintillation signal that is proportional to the number of extracted electrons known as **S2**. The S1 and S2 time delay corresponds to the interaction depth, while the x–y position can be inferred based on the intensity pattern of the S2 light observed on the top PMTs, enabling full 3D position reconstruction, hence the term Time Projection Chamber.

2.1.3 Nuclear and Electronic Interactions

A particle that interacts with a xenon atom in the detector can produce either a *nuclear recoil* (NR) by scattering off atomic nucleus or an *electronic recoil* (ER) by interacting with electrons. In ER events, the deposited energy is mainly transferred to the light and charge quanta, detected via prompt and proportional scintillation light. For NR events, a larger fraction of the deposited energy is converted to heat, which is undetectable by TPCs. In both cases, the energy deposit results in ionized (Xe^+, or ion) and excited (Xe^*, or exciton) xenon atoms. These react to form diatomic excimers (Xe_2^*, or dimers), and de-excitation of the dimers emits scintillation light, $h\nu$ [92]. This interaction process follows two primary reactions; for ionization,

$$\begin{aligned} \mathrm{Xe} + e^- &\rightarrow \mathrm{Xe}^+ + 2e^- \\ \mathrm{Xe}^+ + \mathrm{Xe} + \mathrm{Xe} &\rightarrow \mathrm{Xe}_2^+ + \mathrm{Xe} \\ \mathrm{Xe}_2^+ + e^- &\rightarrow \mathrm{Xe}^{**} + \mathrm{Xe} \\ \mathrm{Xe}^{**} + \mathrm{Xe} + \mathrm{Xe} &\rightarrow \mathrm{Xe}_2^* + \mathrm{Xe} + \mathrm{heat} \\ \mathrm{Xe}_2^* &\rightarrow 2\mathrm{Xe} + h\nu. \end{aligned}$$

Similarly, for the excitation process:

$$\begin{aligned} \mathrm{Xe} + e^- &\rightarrow \mathrm{Xe}^* + e^- \\ \mathrm{Xe}^* + 2\mathrm{Xe} &\rightarrow \mathrm{Xe}_2^* + \mathrm{Xe} \\ \mathrm{Xe}_2^* &\rightarrow 2\mathrm{Xe} + h\nu. \end{aligned}$$

The dimers in both cases eventually dissociate, emitting a VUV photon with wavelength $\lambda = 178$ nm [92]. The energy of this photon, which is about 7 eV, is lower than the excitation energy of the xenon atom. Thus, the emitted light can travel without being reabsorbed by another atom.

Typical ER interactions in the TPC include gamma rays from radioactive backgrounds and neutrino–electron scattering from solar neutrinos. Other ER sources are beta particles, often from intrinsic radioactivity in materials, and alpha particles, primarily from radon decay

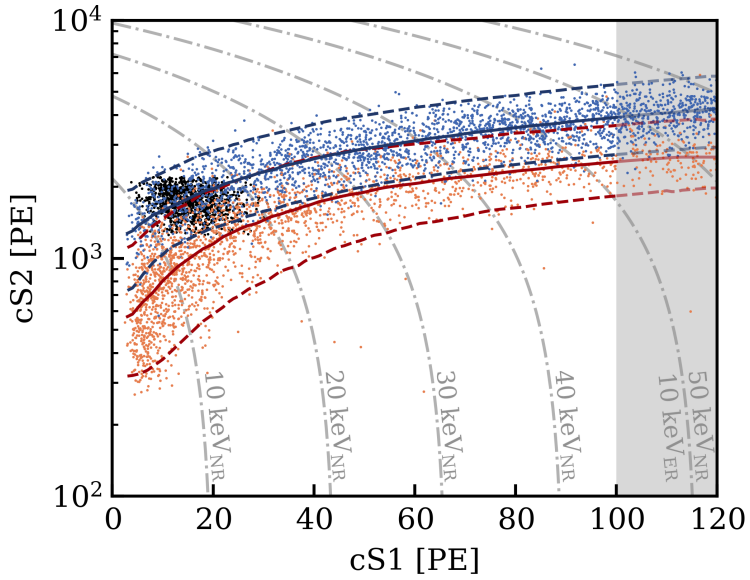

Figure 2.2: NR and ER response model of XENONnT in the first science run. The bands are derived from the calibration sources ^{241}AmBe shown in orange dots, ^{222}Rn shown in blue, and ^{37}Ar shown in black to model the response for NR and ER interactions. The reconstructed NR energies (keV$_{NR}$) are also shown with dotted-dashed grey lines. Figure from [93].

products within the detector. NR, on the other hand, is mainly caused by fast neutrons, hypotetical DM particles (e.g. WIMPs) and neutrinos through CEνNS.

The TPC distinguishes between NR and ER events primarily based on the relative light and charge yields they produce. Under the applied electric field, electron-ion recombination is suppressed, allowing separated electrons to drift toward the gas phase. Since the total number of generated quanta (light and charge) remains nearly constant, an anticorrelation exists between the two. In Figure 2.2 the first science run NR and ER response models of XENONnT are presented. The models are derived using calibration data ^{241}AmBe (orange dots) for NR and ^{222}Rn (blue dots) and ^{37}Ar (black dots) for ER.

2.1.4 Neutrino Interactions with xenon: CC, NC and CEνNS

The probability of a collision between two particles, determined by their atomic properties and relative energies, is quantified by the interaction cross–section. Neutrinos can interact with xenon nuclei through three main channels with different cross–sections: charged current (CC) interactions, neutral current (NC) interactions, and CEνNS [94, 95, 96].

In CC interactions, a neutrino exchanges a W boson with a xenon nucleus, converting into a corresponding charged lepton, often producing a distinctive signal. At neutrino energies near 1 GeV, the approximate cross–section for CC interactions is $\sigma_{CC} \sim 10^{-38}\,\mathrm{cm}^2$ [94] . However, this energy scale is an order of magnitude larger than the typical supernova neutrino energies of several tens of MeV. Although this cross–section is relatively small, neutrinos from a sufficiently close supernova could still produce detectable CC events in xenon [97].

In NC interactions, a neutrino exchanges a Z boson with a xenon nucleus, potentially leaving it in an excited state or leading to secondary particle emission. The cross–section for NC interactions are typically in the range $\sigma_{NC} \sim 10^{-39} - 10^{-38}\,\mathrm{cm}^2$ at 1 GeV [94]. However, at supernova neutrino energies, NC interactions are not prominent because of their low cross–sections.

In contrast, CEνNS is highly relevant at MeV energies typical of supernova neutrinos. In CEνNS, a neutrino scatters off a xenon nucleus coherently, causing the entire nucleus to recoil with minimal change in neutrino energy. This interaction has an approximate cross–section of $\sigma_{\mathrm{CE}\nu\mathrm{NS}} \sim 10^{-39}\,\mathrm{cm}^2$ for neutrinos in the MeV range [95, 96], making it the dominant interaction channel for supernova neutrino detection in XENONnT. Although CC and NC interactions play an important role in other detectors that employ lighter nuclei or focus on specific neutrino flavors, CEνNS is the most relevant channel in xenon-based detectors such as the XENONnT experiment. At energies associated with supernova neutrinos, CEνNS provides a highly sensitive means of detecting the broad neutrino flux [94].

For different recoil energies, the differential interaction cross–section can be calculated by

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_R}(E_R, E_\nu) = \frac{G_F^2}{4\pi} Q_W^2 \left(1 - \frac{m_N E_R}{2E_\nu^2}\right) F^2(E_R), \tag{2.1}$$

where m_N is the mass of the target nucleus, E_R and E_ν denote the recoil and neutrino energies, respectively, and G_F is the Fermi coupling constant [96]. The term $Q_W = N - (1 - 4\sin^2\theta_W)Z$ represents the weak nuclear hypercharge, with N and Z as neutron and proton counts. The weak mixing angle, $\sin^2\theta_W \approx 0.2386$, applies to small momentum transfers, while F^2 is the squared Helm form factor, which describes the nuclear charge distribution.

The Helm form factor [98] used in eq. (2.1) accounts for the finite size of the atomic nuclei when calculating the cross–section for interactions such as CEνNS. The form factor, $F(q)$, is a function of momentum transfer q, which is related to the energy and angle of scattering of the incoming particle. $F(q)$ accounts for the spatial distribution of the charge within the nucleus and it is expressed as

$$F(q) = 3j_1(qr_n)\frac{1}{qr_n}\exp\left(-\frac{(qs)^2}{2}\right), \tag{2.2}$$

where $q^2 = 2m_N E_R$ represents the squared momentum transfer in terms of the recoil energy, $j_1(qr_n)$ is the spherical Bessel function of the first kind, r_n is the nuclear radius parameter and $s = 0.9$ fm is the nuclear skin thickness. The nuclear radius is defined as

$$r_n^2 = c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2, \tag{2.3}$$

where $c = 1.23A^{1/3} - 0.60$ fm and $a = 0.52$ fm, with A denoting the atomic mass number [99]. Taken ^{131}Xe as an example, the squared Helm form factor is shown in the left panel of Figure 2.3, with the corresponding cross–section in the right panel, calculated using eq. (2.1). The color map indicates the cross–section values in $\mathrm{m}^2/\mathrm{keV}$, with the neutrino energy, E_ν on the x-axis and the recoil energy, E_R on the y-axis. The Helm form factor suppresses the cross–section for higher momentum transfers, which corresponds to higher recoil energies. A notable dip occurs around $E_R = 90$ keV in the form factor for ^{131}Xe, where the destructive

interference of the nuclear structure causes the probability of coherent scattering to approach zero. This implies that the contribution of CEνNS is strongly reduced for recoil energies near this region. At higher neutrino energies, the probability of interaction increases, but it is more concentrated at lower recoil energies due to the nature of the scattering kinematics, which favors low–energy recoils. Here we discuss the form factor for ^{131}Xe, but the form factors and cross–sections are also computed for all isotopes listed in Table 2.1. The left panel of Figure 2.4 displays other xenon isotope form factors, while the right panel shows those for the much lighter argon isotopes for comparison.

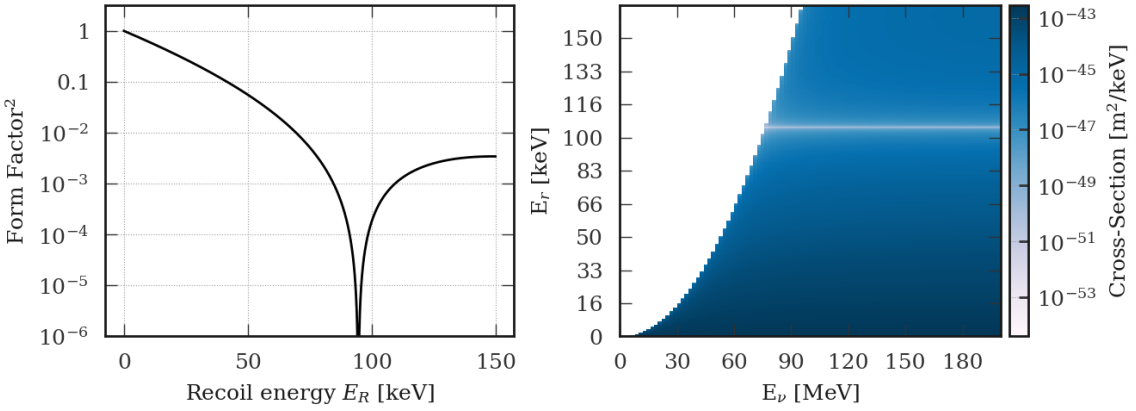

Figure 2.3: Squared Helm form factor, $F^2(q)$ for ^{131}Xe (*left*), and the cross-section probabilities of ^{131}Xe based on the neutrino energies, E_ν and the expected recoil energies, E_R (*right*).

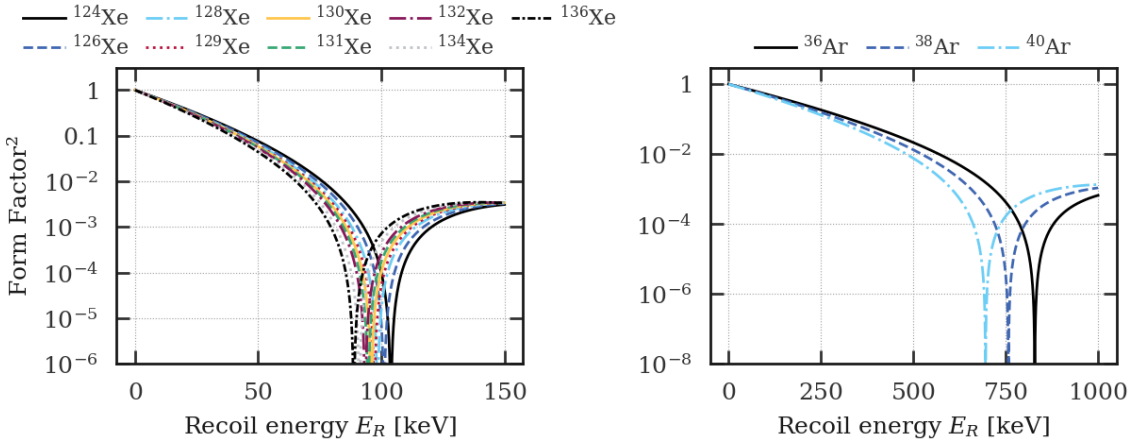

Figure 2.4: Squared Helm form factor, $F^2(q)$ for different xenon isotopes (*left*) and argon isotopes (*right*) discussed in this work. The cross–section of neutrino interactions changes based on the interacting isotope as different isotopes have different form factors.

2.2 XENONnT experiment

The current phase of the XENON project, XENONnT, which began to take science data in 2021, extends the search for DM. The project was first initiated by XENON10 [100] in 2005-2007. Over time, the project progressed through XENON100 [81] from 2008 to 2016, and then XENON1T [101] from 2010 to 2019, each phase improving experimental sensitivity and setting stringent limits on DM properties.

The XENONnT setup comprises three nested detectors at the INFN Laboratori Nazionali del Gran Sasso (LNGS) in Italy: a central cryogenic TPC with 5.9 tonnes of LXe in its active volume[i], surrounded by neutron and muon veto detectors.

2.2.1 XENONnT Data Processing and Data Structure

The XENONnT experiment uses a triggerless data acquisition (DAQ) system [102] to digitize photoelectric signals from PMTs. All data is stored and can be reprocessed with varying settings and corrections. The custom software `strax` [103] in Python reconstructs the recorded signals, and an XENONnT-specific extension called `straxen` [104] manages the detector-specific conditions for signal processing.

Signal reconstruction in `straxen` involves transforming individual PMT *pulses* into high-level *events*, as shown in Figure 2.5. Reconstruction algorithms first identify *hits*, defined as waveform intervals that exceed a threshold, with a 30-ns pre–threshold and 200-ns post-threshold window. Hits in different PMTs are clustered if they are within 700 ns, whereas isolated hits (from afterpulses or dark counts) are stored separately [105].

An illustration showing the waveforms corresponding to the S1 and S2 peaks in an event is shown in Figure 2.6 along with the PMT hit pattern for S2 on the PMTs of the top array. The classification as S1 or S2 depends on the total PMT charge, the rise time (10% to 50% charge quantiles), central 50% width, tight coincidence within 50 ns of the peaklet maximum, and area fraction in the top (AFT) PMT array [105]. Events are identified in the processing as time regions around S2 peaks larger than 100 PE called “triggering peak”. The event is defined as the time window 2.45 ms before and 0.25 ms after the triggering peak, and if multiple event windows overlap, they are merged. To avoid a high number of triggers due to photoionization signals that follow a large S2, only the large S2s that have fewer than 8 other S2s with areas less than 50% of the triggering S2 within ± 10 are accepted as triggering peaks. The largest S1 in the event time window is identified as the primary S1, and the largest S2 signal that follows the primary S1 is called the primary S2.

The experiment stores both low–level data and the software tools for efficiently constructing high–level data. The study in this thesis primarily uses high–level signals; peaks and events. Data types are generally organized in tabular format, with parameters extracted from

[i] The active volume is the volume within the cryostat affected by the electric field

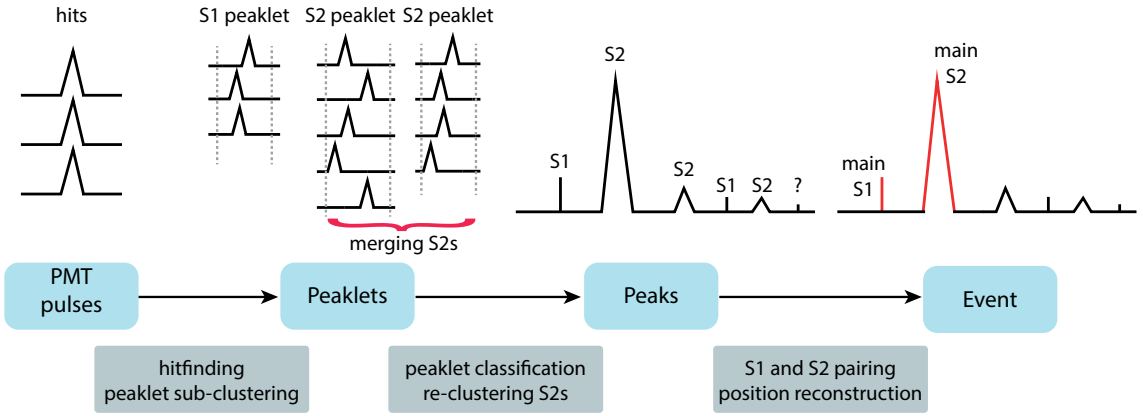

Figure 2.5: Signal reconstruction in XENONnT DAQ. Individual PMT pulses are classified as hits, sub-clustered, and labeled as S1 or S2 peaklets based on waveform shape. Some S2 peaklets are merged if they originate from the same electron cloud. Signals integrated over all PMTs are labeled as S1 or S2 peaks, while unclassified peaks are marked with a question mark. Figure adapted from [105].

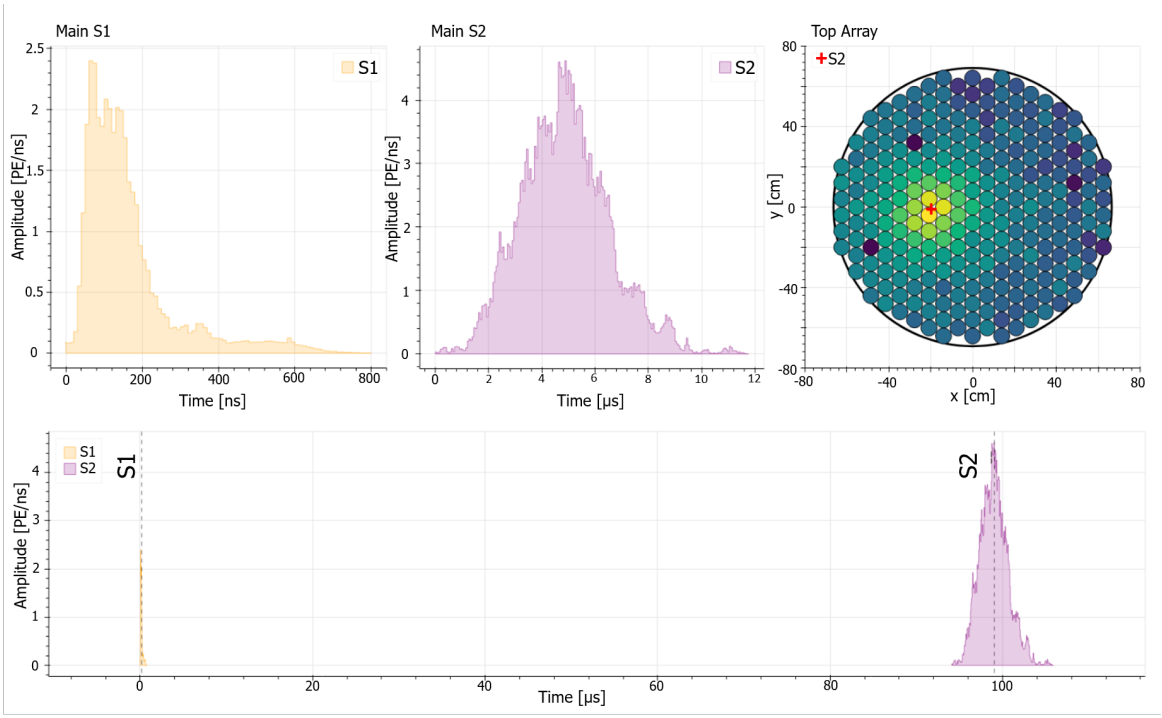

Figure 2.6: Plot of summed waveform and S2 hit pattern for a krypton calibration event. The bottom panel shows the time span between the S1 and S2 signals, with zoomed-in views of S1 and S2 signals in the two upper panels. Notice the different time scales for the two different peak types. The S1 signal exhibits a sharp rise and exponential decay, while the S2 signal has a broader, Gaussian–like shape with a higher amplitude. The hit pattern displays light intensity registered by all PMTs, with yellow indicating the highest intensity.

waveforms such as timing, positions, rise times, summed waveform area, and waveform widths within 50% and 90% area containment. Additional features, such as proximity information (e.g., the number of nearby peaks occurring before and after a given peak), are also sometimes computed to further characterize peaks. A new approach called *subtyping* has been introduced to subdivide detected peaks into more refined classes based on peak features and proximity information, which helps to identify the source of observed signals with greater precision.

The scientific data acquisitions, referred to as science run (SR) began in July 2021 with the SR0 and included 107 days of live-time loss-corrected exposure. This was followed by SR1, which extended the search with an additional live-time exposure of 208 days. During each SR timeline there were additional calibration runs and other special conditions that caused extra down-times. For this study, we focus specifically on a subset of SR1 data collected in background monitoring mode, excluding periods with localized high-rate events (hot spots) and broader rate increases (warm spots). The considered SR1 dataset contains 592 individual runs of different lengths spanning over 83 days of live–time. In the following chapters, a comparison of the treatment for SR0 and SR1 is presented. However, the main analysis framework is built for the SR1 conditions to be able to understand and use the most recent detector conditions.

In LXe, large S1 and S2 signals can generate *delayed electrons* by photoionizing impurities [106]. Between SR0 and SR1, an increase in photoionization (PI) strength led to differences in S2 signal rates. After maintenance between SR0 and SR1, the PI strength increased tenfold, likely due to new impurities from improved radon removal with a high–flow extraction of radon from the LXe target. These impurities do not affect the electron lifetime but cause background rate variations, especially in low-energy events from ^{210}Pb plate-out on the TPC walls [46]. Delayed electron signals, or "single-electron backgrounds", are a major background source in low-energy searches, mimicking signals from real nuclear interactions. These occur when electrons trapped on impurities or surfaces are gradually released after large S1 or S2 signals, forming lingering "electron trains" [107]. XENONnT and previous experiments show that these backgrounds significantly affect low-energy event detection, crucial for DM and CEνNS searches. A common strategy to mitigate these signals incorporates time–based criteria to eliminate signals following large S2s, as well as improved purification and detector materials to decrease impurities [108, 109]. Also, a fidicualization is applied to exclude near-wall volumes of the detector. These strategies enhance the sensitivity of XENONnT to low-energy events, thereby mitigating the influence of delayed electron backgrounds.

2.2.2 Neutrino Interactions in the Veto Detectors

The XENONnT experiment includes two outer veto detectors, each filled with ultra-pure, demineralized water. These detectors are primarily designed to identify external backgrounds, such as muons, but they also offer an additional channel for detecting neutrino interactions via inverse beta decay (IBD). In IBD, an electron antineutrino ($\bar{\nu}_e$) interacts with a proton

in water, producing a positron and a neutron

$$\bar{\nu}_e + p \rightarrow e^+ + n. \tag{2.4}$$

The positron from IBD emits Cherenkov light, detectable by photomultiplier tubes in the water, and the delayed capture of the neutron provides a second signal. This double coincidence, positron followed by neutron capture, significantly enhances the probability of neutrino identification [110].

Although XENONnT's primary focus, with its highest sensitivity, is to detect supernova neutrinos via CEνNS in the TPC, it is also helpful to monitor IBD interactions in water-based veto detectors to enhance the overall detection strategy. The dual detection capability of XENONnT, utilizing both the CEνNS signals in the TPC and the IBD signals in the veto detectors, allows for improved sensitivity to a supernova neutrino burst. Together, the TPC and veto detectors form a multi–component detection system that can increase the significance of a supernova signal when analyzed jointly. This complementary approach will be further discussed in subsequent chapters, highlighting how integrating signals from all three detectors can provide a powerful, comprehensive response to a nearby supernova.

2.2.3 Recent Results

As of late 2024, XENONnT continues to collect data and analyzes the results of its initial science runs. Recent works [93, 105, 111] discuss the signal reconstruction methods, calibrations, and WIMP search modeling of the experiment. These papers detail the signal processing and the blind analysis approaches used for unbiased searches, with a dual-phase TPC for precise event detection, the results of the WIMP searches over 1.1 tonne-years, providing stringent exclusion limits despite no significant findings.

The exclusion curves for the spin-independent WIMP-nucleon cross–section for WIMP masses between 6 GeV/c^2 and the TeV/c^2 scale, derived from XENONnT, together with those of other experiments and the projected limits of future experiments are presented in Figure 2.7.

Also shown in the figure is the *neutrino floor* in light blue, the ultimate sensitivity limit for direct DM detection experiments due to an irreducible background from solar, atmospheric and diffuse supernova neutrinos [112]. At these ultra-low interaction rates, CEνNS from solar and atmospheric neutrinos produce signals indistinguishable from those of WIMPs, creating what is often referred to as the 'neutrino fog'. This background forms a floor below which it becomes increasingly challenging to differentiate between WIMP interactions and neutrino-induced events, as WIMPs with cross–sections below this threshold would be effectively obscured by the neutrino background. However, this sensitivity level opens a unique avenue for studying neutrino interactions, particularly through CEνNS, which XENONnT has already demonstrated by its recent measurement of CEνNS signals from solar ^{8}B neutrinos marking a milestone for DM detectors in the field of neutrino physics. In a blind analysis with an exposure of 3.51 tonne-years, XENONnT observed 37 events above 0.5 keV, with an expected background of $(26.4^{+1.4}_{-1.3})$ events [46]. The background-only

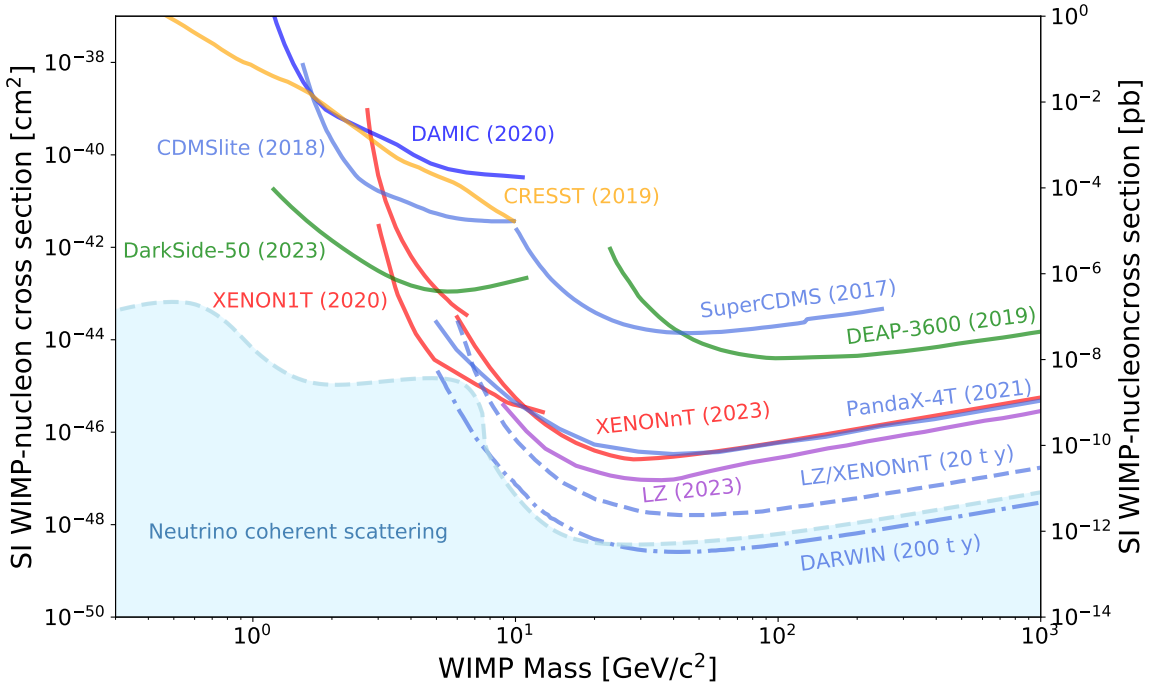

Figure 2.7: Exclusion limits (solid) and projections of limits (dashed) of spin-independent WIMP-nucleon cross–section for existing and future detectors. Figure adapted from [74, 112]

hypothesis was rejected with a statistical significance of 2.73 σ, leading to a measured ^{8}B neutrino flux of $(4.7^{+3.6}_{-2.3}) \times 10^6$ cm^{-2}s^{-1}. This measurement is consistent with results from dedicated solar neutrino experiments, confirming the sensitivity of liquid xenon detectors to low-energy neutrino interactions and validating the Standard Model prediction for the CEνNS cross–section on xenon.

This capability also underpins the potential for similar observations from supernova neutrinos, as both CEνNS channels would allow for detailed studies of low-energy neutrino physics and enhance multimessenger astrophysics efforts.

2.2.4 Successor of XENONnT: XLZD

The development of large-scale LXe TPCs for dark matter research continues with the next-generation XLZD project, a collaborative effort that unites the XENON, LZ, and DARWIN collaborations [73]. Often referred to as "the ultimate DM experiment," XLZD is expected to feature a target mass of approximately 50 tons of LXe, significantly expanding the detection capabilities beyond those of XENONnT [74, 75]. Notably, the XLZD project builds upon the groundwork laid by the DARWIN concept, which initially proposed this scale of detector but has since evolved into the larger, unified XLZD collaboration.

In Figure 2.7, the sensitivity projections for this next-generation experiment are shown under the name DARWIN, as the XLZD collaboration had not yet been formally established at the time of the publication [74]. Despite the name difference, the projections are applicable to

the capabilities envisioned for XLZD. It can be seen that with an estimated 200 ton-year exposure, XLZD is expected to explore the entire parameter space accessible for WIMP DM and reach the neutrino floor, the fundamental limit below which the WIMP signals are effectively indistinguishable from the background neutrino interactions.

Although XLZD is primarily optimized to search for WIMPs above masses of about 5 GeV, its ultra–low background levels also make it well suited to observe other rare interactions, including CEνNS. This interaction provides a unique window into low-energy neutrino physics, which has implications for supernova neutrino detection and other astrophysical neutrino sources.

In the event of a galactic supernova, a detector of the XLZD scale is expected to detect on the order of 1000 neutrino interactions, as we will demonstrate in Chapter 6. This number approaches the statistics collected by dedicated neutrino observatories, positioning XLZD as an important player in multimessenger astrophysics. Its ability to detect supernova neutrinos through CEνNS could complement data from other neutrino observatories, allowing for a richer understanding of the explosion dynamics of supernovae, neutrino properties and the role of neutrinos in such astrophysical events [47].

3

The Supernova Early Warning System: SNEWS

"The cosmos is within us. We are made of star-stuff. We are a way for the universe to know itself."
–Carl Sagan

Contents

The physics of supernova neutrino emission is discussed in Section 1.2.1, and the XENONnT DM experiment is introduced in Section 2.2. As it will be discussed in Chapters 4 and 5, XENONnT is well positioned to actively contribute to the global network of neutrino observatories and to be a significant player in multimessenger astronomy, aligning with the collective objective of promptly detecting the next galactic supernova.

This chapter describes the Supernova Early Warning System (SNEWS) and its major upgrade to SNEWS 2.0. The two main software essential for this upgrade, called SNEWS Publishing Tools and SNEWS Coincidence System, detailed in Section 3.2 and in Section 3.3 respectively, are developed within this thesis. The work described in this chapter is also published in the Journal of Instrumentation (JINST) [113].

3.1 Overview of SNEWS

SNEWS [114] is a global network of neutrino observatories. SNEWS coordinates among various experiments to provide an early warning system for the next CCSN by detecting the burst of neutrinos that precedes the optical signal. This chapter introduces SNEWS and outlines the specifics of its publishing tools and coincidence servers in their respective sections. The following chapters will elaborate on how XENONnT interacts with SNEWS, explaining the role and incorporation of XENONnT within the detection network.

The occurrence of a future CCSN in our galaxy is expected to yield signals detectable through neutrino emissions. In preparation for such an event, we foresee three key signals: detectable neutrino emissions, gravitational waves, and electromagnetic waves, all observable from Earth. This combination of signals forms a rich dataset in Multi-Messenger Astronomy (MMA), significantly enhancing our understanding of these remarkable astrophysical events.

Such cosmic events are predicted to happen just once or twice every hundred years [29, 115], making proactive preparation essential to optimize scientific outcomes. Coordination among global observatories is crucial to ensure optimal data collection. Comprehensive coverage of the entire spectrum, including neutrino emissions, provides a deeper understanding of stellar evolution and neutrino physics [40]. Conversely, without adequate preparation, we risk missing this unique opportunity to capture and study a rare and extraordinary phenomenon, thereby limiting our ability to gain valuable insights into the complexities of these cosmic events and their implications for astrophysics and particle physics [14].

SNEWS traces its roots to a workshop in 1998, which unified neutrino experimentalists, theorists, and astronomers in the exploration of establishing a global network [116, 117]. The core functions of the system emphasize early detection for the next galactic CCSN, inferred from coincident neutrino signals worldwide, allowing coordinated efforts among observatories to study this astronomical phenomenon comprehensively. Today, SNEWS comprises both neutrino observatories and DM detectors capable of detecting neutrino signals, such as XENONnT, LZ, and the future XLZD [118, 44, 74]. The current collaborative effort involves various aspects of preparation for a galactic supernova, including modeling, signal prediction, alert formation, triangulation, and coordination for follow-up observations across different

wavelengths. These observations span the electromagnetic spectrum and involve ground- and space-based telescopes, as well as gravitational wave observatories [119, 120, 121]. Traditional procedures involve extensive data analysis, testing, and cautious data sharing, leading to a slower process. In contrast, SNEWS enables the registration of even tentative observations, initiating an automated coincidence search to expedite the detection process. Signals that simultaneously occur in multiple independent experiments are highly unlikely to be mere noise, providing a robust basis for automated systems within the network to issue prompt alerts to the global astronomical community.

Using the combined data from various experiments, SNEWS optimizes the lead time for global observers to prepare for the imminent electromagnetic phenomena accompanying a supernova event. This automated and collaborative approach ensures rapid and reliable information dissemination, enhancing the overall efficiency of early warning systems in studying these cosmic occurrences.

3.1.1 From SNEWS 1.0 to SNEWS 2.0

The initial iteration of SNEWS, referred to as SNEWS 1.0 entered a testing phase in 2001 and reached full operational status on 1 July 2005 [116]. Since then, it has possibly been the first and longest-running MMA project with seven neutrino experiments involved[i].

SNEWS 1.0 functions by collecting simple datagrams through the User Datagram Protocol (UDP) on its central server at Brookhaven National Laboratory (BNL). A backup server was strategically established at INFN Bologna. These datagrams convey the time of observation and several optional fields, although the additional fields are rarely utilized in practice. The primary role of the server was to monitor the experimental alerts. If two or more alerts are received within a 10-second window, the server promptly notifies an email list comprising interested parties. Since its inception, this mailing list has attracted almost 7000 subscribers. Notably, no core-collapse supernova events have been observed within our galaxy during this period, rendering this mailing list possibly the least traffic-intensive on the Internet.

To balance sensitivity and specificity, participating experiments within SNEWS 1.0 adhered to a self-imposed limitation of one false alarm per week. If an experiment generates alarms too frequently, it was flagged as noisy, resulting in its two-fold coincidences being demoted to nonautomated status. This constraint keeps the Poisson probability of accidental coincidence lower than the actual supernova rate, estimated at 1.63 ± 0.46 per century [115]. Although this strict criterion ensures a very low false-positive rate, it may introduce the possibility of false negatives.

To address this, SNEWS is being upgraded to SNEWS 2.0. Among other notable upgrades such as introduction of new data tiers beyond the traditional coincidence alert tier, this iteration relaxes the strict limit on requesting only high-confidence detections and includes the incorporation of new observatories and enhanced versions of existing ones. SNEWS 2.0 aims to balance reducing false positives and avoiding potential false negatives. It adopts

[i] Super-K (Japan), LVD (Italy), Ice Cube (South Pole), KamLAND (Japan), Borexino (Italy) KM3NeT (Mediterranean), HALO (Canada)

a more inclusive approach, anticipating any signal that might originate from a galactic supernova. It employs a similar coincidence algorithm to assess false probabilities, triggering when two or more observatories report times within a predetermined coincidence window described in Section 3.3.

3.1.2 SNEWS Today

Today, SNEWS 2.0 features enhanced functionalities by integrating more neutrino detectors and observatories, increasing sensitivity and range. This includes water Cherenkov detectors like Super-Kamiokande and the upcoming Hyper-Kamiokande [122, 123], which target low-energy neutrinos via inverse beta decay and elastic scattering in large water volumes. Another type uses natural settings such as lakes, Antarctic ice, or deep-sea sites such as Baikal, IceCube, and KM3NET [124, 125, 126] for high-energy atmospheric and astrophysical neutrinos. LArTPCs such as DUNE [127] and ICARUS [128] detect neutrinos through interactions with argon nuclei. DM detectors such as XENONnT and LUX-ZEPLIN contribute via CEνNS. Scintillator detectors, including the Borexino [129] and future JUNO [130], also detect via inverse beta decay and elastic scattering. Future experiments like KM3NeT and Hyper-Kamiokande will significantly increase the sensitivity and observational capabilities of SNEWS 2.0 [118].

Although the primary objective of SNEWS 2.0 is to provide a prompt and reliable alert for a galactic CCSN by detecting the high flux of neutrinos that precede the event, the system has expanded its scope to include additional scientific goals across different operational tiers. SNEWS 2.0 operates with three key tiers: coincidence, significance, and timing, as described below and illustrated in Figure 3.1.

- **Coincidence Tier**: Focuses on identifying simultaneous neutrino signals from multiple detectors, a crucial indicator of a potential galactic CCSN event. Using a sophisticated algorithm, an alert is triggered when two or more detectors report observations within a specified time window. This tier ensures a rapid response to potential supernova signals, maximizing the lead time for subsequent observations across different wavelengths (described in Section 3.3).

- **Significance Tier**: Enhances supernova detection by assessing the statistical significance of each signal. Users submit p-values and corresponding time bin widths, allowing for a refined analysis of observed signals. This tier enables the combination of marginal signals from different detectors to achieve a global significance that is large enough to be of interest. This is particularly useful for pre-supernova neutrinos, where even the best detectors have a limited range, thereby improving the overall scientific value of the data.

- **Timing Tier**: Focuses on the temporal aspects of neutrino detections. Users submit lists of individual neutrino signal times, offering a detailed temporal profile of the signals. This information aids in studying the dynamics of supernova processes and refining neutrino emission models. Additionally, timing data from multiple detectors is

used to triangulate the signal's location, providing valuable information for the broader astronomical community.

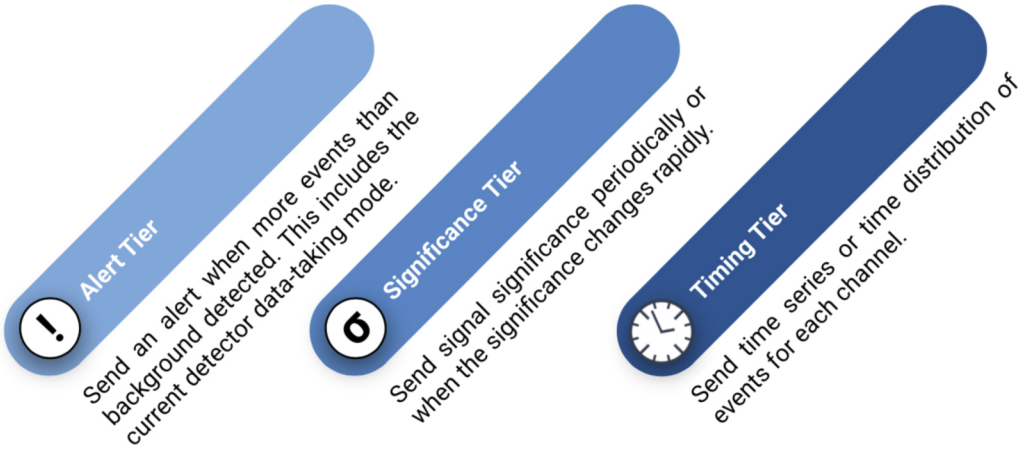

Figure 3.1: Representation of the three core tiers of SNEWS. Figure is from [118].

Beyond these main tiers, the network also addresses broader aspects, including modeling, coordination for follow-up observations in other wavelengths, multilateration, global significance inference, and presupernova neutrino detection.

Collaborating with AAVSO [131] and REFITT [132], SNEWS mobilizes astronomers for a rapid optical follow-up of a supernova neutrino alert. Due to their geographic diversity and flexibility, amateur astronomers offer rapid responses and global coverage, unlike professional observatories that require scheduling. Although professional telescopes can be overwhelmed by the brightness of nearby supernovae, amateur telescopes capture initial data, guiding professional follow-ups and optimizing equipment use, thereby enhancing early data capture.

SNEWS is conducting long-term photometric monitoring of Red Super Giant (RSG) stars, which are potential candidates for CCSN. By observing 640 identified RSG stars, SNEWS aims to detect photometric variations that might indicate pre-supernova activity [133]. Although this may not precisely predict an explosion, it could offer insights into the final stages of stellar evolution, enhancing the understanding of massive stars' behavior before a supernova. These observations can also provide valuable information about supernova progenitors following a core-collapse supernova event.

In addition to astronomical follow-up, SNEWS plays a crucial role in assisting gravitational wave observatories, such as LIGO and Virgo [134], by narrowing down the time window for their searches based on the neutrino signal. Studies show that when the neutrino signal is detected, gravitational wave observatories can significantly increase their sensitivity within that time frame, improving the chances of gravitational wave detection associated with the supernova event [40].

SNEWS is also exploring multilateration (or, informally known as triangulation), utilizing data from global detectors to identify the direction of incoming neutrinos. This method can estimate the supernova's sky location, potentially enhancing follow-up observations by localizing within a few degrees [135].

Currently, collaborations between experiments are being solidified through Memoranda of Understanding (MoU). These agreements outline the integration of experiments into the SNEWS network, ensuring that they can provide and receive timely alerts. XENONnT is actively participating in this process and the formal inclusion in SNEWS 2.0 is underway, with future MoUs expected to solidify this partnership.

SNEWS 2.0 harnesses advances in computational environments and resources to optimize its performance. The introduction of new software tools, developed in Python, brings improved flexibility, performance, and additional features. In addition, this new development environment allows for long-term maintainability. At the heart of SNEWS 2.0, two primary software packages, *SNEWS Publishing Tools* and *SNEWS Coincidence System*, have been developed to streamline its functionality.

The first package serves as the backend, operating on the server to execute the coincidence algorithm, detect trigger coincidences, and distribute alerts. Meanwhile, the second package functions as a front-end, providing users with an interface to interact with the SNEWS server. The subsequent sections delve into a detailed explanation of these tools, explaining their roles and functionalities within the upgraded SNEWS 2.0 framework.

3.2 SNEWS Publishing Tools

Serving as a user interface, the SNEWS Publishing Tools (`snews_pt`) is a robust Python package equipped with a variety of features. The essence of online communication is delegated to Scalable Cyberinfrastructure to support Multi-Messenger Astrophysics (SCiMMA) [136], using Kafka protocols (see Section 3.2.1), thus enabling `snews_pt` to concentrate on additional functionalities. The primary tools offered to users include the ability to subscribe to incoming alerts issued by the server, publish their observations to the extent they prefer, and transmit periodic heartbeats to ensure communication integrity. In the subsequent subsections, the discussion will be structured as follows: first, an examination of the authentication framework governing communication channels; second, an exploration of subscription and publishing functionalities in their respective subsections; and finally, an introduction to remote commands and their operational mechanisms.

3.2.1 Credentials and Kafka Topics

To establish secure connectivity between servers and authentic users, credentials play an important role in identifying users. The SCiMMA community offers a hopskotch package [137] designed to handle authentication and credential management. This package utilizes Kafka [138] protocols to establish secure connections between authenticated users and authorized servers. Kafka is a distributed streaming platform that is used to build real-time data pipelines and streaming applications. It can publish and subscribe to streams of records, store them, and process them. Users of member experiments can create accounts

on the website using institutional email addresses for authentication and request access to specific topics. A *topic* in this context refers to a specific channel or subject within the messaging system, typically used to organize and manage different categories of messages.

Administrators within the SNEWS Collaboration then grant *read-only* access to coincidence alert topics, *write-only* access to observation topics, and both *read and write* access to test topics. In contrast, the server is endowed with *write* access to the coincidence alert topic and *read-only* access to the observation topic. This configuration ensures that observation messages, containing only suspected supernova observation timing, can be read only by the server. They remain undisclosed to the network unless they trigger the coincidence logic, at which point all members can access the alert message issued by the server on the coincidence alert topic. This approach guarantees secure and confidential communication between users and the server.

The test topics serve various purposes, including general purpose testing and connection testing for users, as explained in the following.

3.2.2 Subscription

The subscription feature, designed for broad accessibility, extends beyond the confines of SNEWS membership to include a wider audience. This inclusive approach allows anyone to subscribe and monitor the incoming alerts, which are inherently public. Facilitated by the SNEWS Publishing Tools (`snews_pt`), users have two distinct avenues for connection and alert monitoring: a command-line interface (CLI) and an interactive Python application programming interface (API), such as Jupyter notebooks.

Upon executing the subscription command, an uninterrupted connection is established to the relevant coincidence alert topic. When an alert is received from the server, its contents are saved in a local JSON file while the alert-listening process continues. Alerts are triggered as soon as two detectors form a coincidence and generate the initial alert labeled as `"New Coincidence"`. Subsequent updates with tags `"Alert Update"` are received iteratively as more detectors contribute their data, all of which are appended to the existing JSON file. Different alert types are discussed in the following sections.

In addition to the core listening functionality, the CLI introduces a notable feature: plug-in scripts that can optionally be used to perform some additional tasks using the alert data. These scripts are executed automatically upon receiving an alert. Users are guided to incorporate two lines of code at the beginning of their custom follow-up script.

Upon receiving an alert, the subscribe function creates a time-stamped JSON file and triggers the execution of the custom script, passing the filename as a parameter. This design allows users to automate and customize follow-up procedures seamlessly, leveraging the information in the JSON file.

This dual functionality and automated follow-up scripting offer a flexible and responsive framework for users to interact with the alert system, accommodating a range of preferences and requirements within the scientific community. For instance, a pointing script can

be integrated to use the initial neutrino time information from each coincidence-forming experiment to infer the potential position of the neutrino source. This script will be executed automatically upon receiving each alert.

3.2.3 Publishing

The `snews_pt` communication tool offers a versatile mechanism for users to publish their observations and heartbeats. Adopting an object-oriented paradigm, this package enables users to pass relevant information and interact with the message object before final submission. In particular, messages undergo validation upon creation, ensuring integrity and adherence to specified criteria.

As mentioned in the previous Section 3.1, `snews_pt` accommodates the distinct tiers within its design, allowing users to build messages tailored for various purposes. To streamline user interaction, a default configuration file is used as input for the software and it predefines topics for both production and development environments. In addition, it incorporates essential information about the detector.

Upon initial execution, users are recommended to set their names using the tools provided. Once set, this information is recorded in the configuration file and becomes the default argument for future interactions, eliminating the need to enter the detector name each time. A similar approach applies to specifying Kafka topics, and the default topics are already predefined in the configuration file. Users only need to indicate whether the created message is intended for testing, and the correct topic is automatically selected.

Furthermore, `snews_pt` offers various message templates for different tiers and other purposes. These messages are dynamically generated based on the input provided and their content is validated using predefined message models. For example, a message containing the initial neutrino time would be classified as a coincidence-tier message, whereas a message with detector status information would be classified as a heartbeat message. Users can easily input the desired information and `snews_pt` will automatically produce the corresponding messages, validating their content in the background. The range of possible messages and their required arguments are explained in the following and also summarized in Table 3.1. This approach improves user efficiency by reducing tasks and offering an intuitive interface to create messages specific to various communication tiers.

In addition, users can view and interact with the generated messages. In the event of an invalid message, a warning message is displayed to inform the user about the issue, ensuring transparency and effective error handling.

The message constructor function accepts various arguments, each with a specified format. To qualify as a message for a particular tier, there are typically two essential inputs. The first is the detector name, which can be retrieved from the configuration file once the user has set it. The other input(s) determine the tier. Here is a breakdown of different types of messages and their respective input requirements:

	Format	Coincidence	Significance	Timing	Heartbeats	Retraction
detector name	str	**Required***	**Required***	**Required***	**Required***	**Required***
initial neutrino time	str	**Required**	-	**Required**	-	-
machine time	str	Optional	Optional	Optional	Optional	Optional
observation p-value	float [0,1]	Optional	Optional	Optional	-	-
p-values for time bins	list [float]	-	**Required**	-	-	-
width of time bins	float	-	**Required**	-	-	-
neutrino time series (histo-grams)	list [str (int)]	-	-	**Required**	-	-
detector status	str ON|OFF	-	-	-	**Required**	-
retract latest	int	-	-	-	-	**Required**

Table 3.1: The required and optional arguments for different intended uses. (*) The detector name can be set once by the experiment, and in the following uses, the tools always fetch this name.

1. **Coincidence Tier Message:**
   - Generated when the user provides the "initial neutrino time" as an argument in International Organization for Standardization (ISO) format or a string that can be formatted to this standard.
   - Optional arguments may include machine time indicating the time when the data was read and the overall detection p-value.
2. **Significance-Tier Message:**
   - Users can provide p-values corresponding to each time bin alongside their specified widths.
   - P-values format: a list of floats; bin width format: a float number.
   - Optional arguments, such as machine time and observation p-values, can also be included as needed.
3. **Timing-Tier Message:**
   - Users are required to pass a list of strings containing individual neutrino times.
   - Optional arguments, like machine time and observation p-values, can be included if available.
4. **Heartbeat Messages:**

- Crucial for monitoring experiment connectivity, calculating false alarm probabilities during triggered alerts, and potentially coordinating downtimes within the network.

- Creating a heartbeat message is straightforward; users only need to pass their detector status as either "ON" or "OFF".

- All messages are automatically timestamped with the current date-time upon submission to SNEWS to enable latency-tracking.

5. **Retraction Messages:**

   - Users employ retraction messages to retract previously submitted messages from the cache.

   - Play a crucial role in mitigating false triggers and minimizing the risk of high false alarm rates.

   - To initiate a retraction, the user must provide a "retract latest" argument with an integer value indicating the number of messages to retract from the last submitted messages.

   - This feature proves particularly useful in scenarios where there may be programmatic mistakes leading to the unintentional submission of several observation messages.

Predefined arguments serve multiple purposes and cater to various intentions using the same message builder class. To allow flexibility, users may include specific information as key-value pairs. This additional information is stored in a separate field called a meta-field, which is not processed immediately by the server nor included in distributed alerts but is stored for later access. When receiving arguments, the builder parses them to determine their intended use and constructs the appropriate messages. Each message is stored in a list that becomes an attribute of the object. During message creation, meta-fields are generated only for keys that do not belong to any other generated message, preventing duplication of information and maintaining clarity in messages.

To transmit the generated messages to the server, users can invoke a method on the message object. This method offers optional arguments, such as indicating whether the message is designated for the testing topic (also known as the firedrill topic) or the main production topic. Once executed, the generated messages are automatically delivered to the relevant Kafka topic, making them available to the server. The timing tier and significance tier messages are correctly labeled and inserted into the stream. However, it is important to note that the coincidence system only selects and handles the messages of the coincidence tier, along with heartbeat and retraction messages; see Section 3.3.

3.2.4 Remote Commands

In addition to message submission, users have the capability to initiate remote commands that prompt the server to take specific actions. These commands serve different purposes, including cleaning the cache, testing the connection to the server, and asking for heartbeat feedback. Users can execute a designated function to trigger commands. These commands essentially represent a unique category of messages recognized by the server. Upon detecting these special messages, the server invokes different functions to execute the corresponding actions.

The cache cleaning command proves particularly useful during the development phases. When implementing changes, developers often seek to refresh the cache to assess the effects of recent modifications. Specifically, the `"clear_cache"` command purges the existing entries in the cached database on the server and creates a new one. This ensures a clean slate for testing and evaluating the system's response to recent updates. The two additional commands, namely `"request_heartbeat_feedback"` and `"test_connection"`, cater to the broader user base.

There are instances where users may send messages and desire confirmation that these messages have been successfully received by the server. To address this concern, users can initiate a connection test command, which involves creating a message containing the detector name and execution time, tagged as "sending" before. Subsequently, upon submission, the system opens the third topic mentioned in Section 3.2.1, called the test topic, and awaits a response. Upon receiving the test connection command, the server duplicates the message contents, alters the status tag to "received", and sends it back to the testing topic. The user's system monitors this test topic until a specified timeout is reached (an adjustable argument). If the user identifies their message with the correct detector name and time, modified as expected, they can confidently conclude that the message has successfully bounced back, confirming the connection. This process provides users with tangible verification that their message has reached the server.

Lastly, users have the option to request feedback on their heartbeats by submitting a specific command. This command includes an argument for inputting email addresses to which the user wants feedback. Upon receiving this command, the server undergoes an authentication process by comparing the provided email addresses with the known authorized member email addresses for the given observatories. Once authenticated, the server compiles data containing the heartbeat information of the experiment, including computed latencies, and sends an email to the specified addresses. The email includes relevant data and a graphical representation, as shown in Figure 3.2. More details on this process are provided in Section 3.3.4.

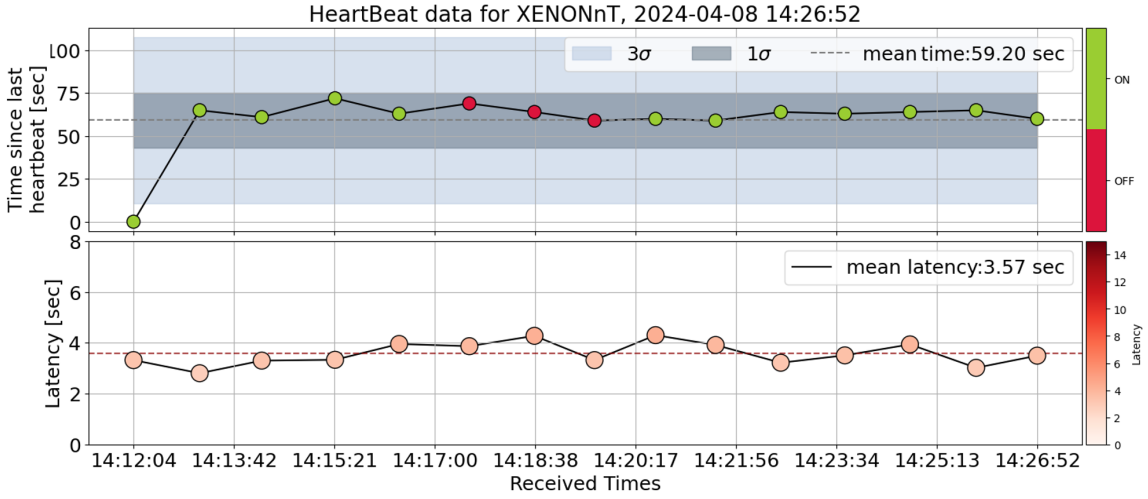

Figure 3.2: The figure illustrates a simulated example from the XENONnT experiment spanning a 15-minute interval, showcasing key metrics related to the server's reception of heartbeats. The top plot displays the time intervals representing the duration between successive heartbeats reaching the server. Marker colors differentiate between 'ON' and 'OFF' states of the received signals. Below, the latency is depicted, delineating the time span from when the signal is dispatched by the member to when it is processed by the server. Darker shades of red indicate higher latency levels.

3.3 SNEWS Coincidence System

The SNEWS Coincidence System (`snews_cs`) serves as the back-end algorithm that runs on the servers, functioning as the main coincidence logic within the Supernova Early Warning System. This central component processes the messages of the coincidence tier as described in Section 3.1.2. Primarily utilized and managed by SNEWS developers and administrators, this script plays a pivotal role in determining coincidences and initiating alerts, subsequently disseminated through various channels explained in the ensuing subsections.

Given its centrality in the information flow, the script considers various scenarios to ensure robustness. The subsequent subsections detail the capabilities of this module, commencing with an explanation of the data flow (Section 3.3.1) and handling mechanisms (Section 3.3.2).

Continually operating on the server, the coincidence system is designed to withstand continuous operation, with measures in place to minimize downtimes and facilitate maintenance, as explained in Section 3.3.5.

3.3.1 Coincidence Logic

The SNEWS Coincidence System employs a designed algorithm to ensure accurate handling of coincidences. Its primary objective is to cover all potential scenarios while avoiding any oversight of coincident events. The system defines a default coincidence window of 10 seconds, indicating that a coincidence alert is triggered when at least two detectors report

initial neutrino times within this temporal proximity. The sequence of these messages is not important, and other experiments can join the existing alert by reporting a time within the same window. A flow chart explaining the process of coincidence handling is shown in Figure 3.3.

In cases where a detector reports an initial neutrino time that aligns with the second detector but not the first, the system adeptly identifies these situations. It introduces 'sub-groups', simultaneously enabling multiple coincidence windows to coexist. Managing cached data becomes crucial as detectors may belong to several subgroups simultaneously.

Although the system promptly triggers a coincidence alert upon the formation of a coincidence between a pair of detectors, it continually receives additional information and updates its cached data. When a new detector joins an existing coincidence alert, the system recognizes

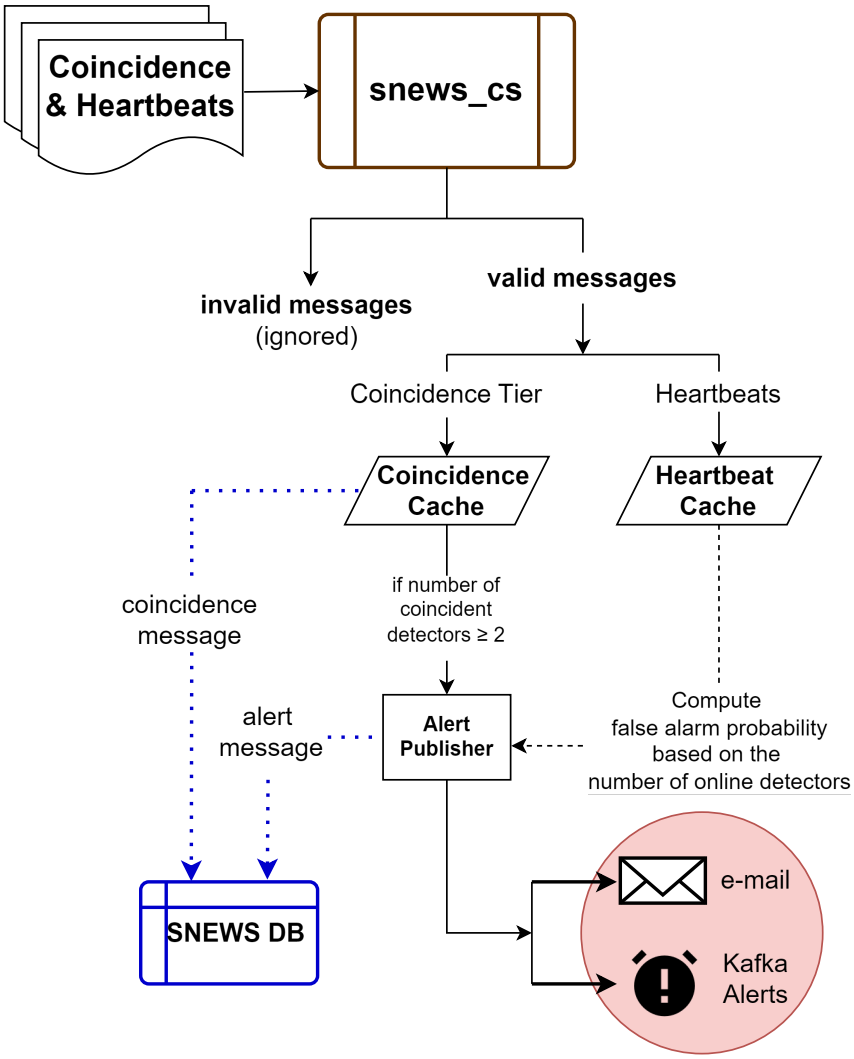

Figure 3.3: Flowchart illustrating the logic of the SNEWS coincidence system. Coincidence and heartbeat messages are initially registered in the `snews_cs`, where they undergo validation for their input. The *CoincidenceCache* and *HeartbeatCache* classes handle coincidence and heartbeat messages, respectively. Upon detecting a coincidence between two or more detectors, an alert is generated and published. Both coincidence and alert messages are stored for record-keeping purposes. The *HeartbeatCache* class computes false alarm rates for alerts, which are then appended to the alert message. Finally, alerts are distributed via Kafka topics, emails, and other channels.

this and issues an update alert, incorporating information from the additional detector. If a subsequent coincidence involves the new and existing detectors, distinct alerts with new subgroup IDs are generated and transmitted.

The system also demonstrates sophistication in recognizing detector updates. In cases where a detector sends an update and a coincidence has already been established, a new alert marked with an `"Alert Update"` flag circulates, emphasizing the modified time of the relevant detector. If no coincidence had occurred before the update but formed afterward, the system considers the latest value and circulates a new initial alert.

In instances where a detector retracts its message, particularly if it had already formed a coincidence with another existing detector observation leading to an alert, the system issues an update alert. This alert exclusively features the remaining detector, signaling the retraction of the observation by the initial detector. This approach ensures the accuracy and reliability of the coincidence logic within the SNEWS Coincidence System.

The observation messages containing suspected initial neutrino observation times undergo a two-tier storage approach. Initially, they are cached in a Pandas DataFrame[ii] for a 24-hour duration, allowing continuous comparison with other messages in the cache. Simultaneously, these observation messages are archived in the SNEWS database, providing a more permanent storage solution. Similarly, coincident alerts formed between two or more detectors are also stored in the database. Although the precise memoranda of understanding regarding long-term data storage between the SNEWS network and individual experiments are not finalized, the current implementation offers flexibility to accommodate future agreements.

For registered heartbeats, the default storage policy suggests a 48-hour retention period in a comma-separated file. This file is continuously updated, and older entries are systematically deleted to make room for newer additions. This storage approach ensures efficient management of heartbeat data over time.

The core working principle of coincidence logic is visually depicted in Figure 3.3, providing an illustrative overview of the functionality of the system. Specifics regarding command handling on the server, alert distribution, heartbeat handling, and the feedback mechanism are outlined in their respective subsections, offering a more detailed understanding of the overall system architecture and operation.

3.3.2 Command Handling at the Server

The functionality of the SNEWS Coincidence System (`snews_cs`) is structured around various class objects and a modular design. At its core, a persistent Kafka subscriber runs within the coincidence script, consuming messages from the observation topic, as detailed in Section 3.2.1. Each incoming message undergoes a series of steps. Initially, it is directed to a *CommandHandler*, where the message content and intended tier are examined. The

[ii] Pandas is a widely used Python library for data manipulation and analysis. A *DataFrame* is a labeled data structure provided by Pandas, designed to store and manage tabular data efficiently, supporting various data types.

CommandHandler verifies the format and validity of the message. Specifically, `snews_cs` manages messages properly formatted and dispatched using `snews_pt` for coincidence tier, heartbeat, or command purposes. However, messages directed at the significance or time tier, as explained in Section 3.1.2, are recognized but ignored. The flow chart shown in Figure 3.4 explains the process.

Properly formatted messages, those sent through `snews_pt`, include a field indicating the message ID. This ID plays a pivotal role within the *CommandHandler*, categorizing the message content and guiding subsequent actions. For remote commands introduced in Section 3.2.4, the `Commands` class is invoked, where further checks are executed before the actual command is processed. The commands and their respective actions on the coincidence server can be listed as follows.

- The cache–reset command is exclusively reserved for software maintainers for testing purposes. When this remote command is received, it is expected to contain a field with a simple password. The server authenticates this password, granting access to reset the cache. Once validated, the command handler instructs the coincidence script to reset its cache and initiates a new Pandas DataFrame.

- For the test connection command, as detailed in Section 3.2.4, user input is manipulated and reintroduced into the test topic stream. Validation is performed on the user's end with `snews_pt`.

- When the received message pertains to a heartbeat, a distinct class object managing heartbeats is called. For the heartbeat request feedback command, the requesting email addresses are compared against the addresses registered for the requesting experiment. If the addresses belong to a known authorized user, feedback is generated and sent to the user.

- For retraction messages, the *CommandHandler* directs the coincidence script to remove the last N numbers of messages, as specified in the retraction request. Subsequently, it triggers a rerun of the coincidence logic to identify any consequential changes. If the redacted messages had initially formed a coincidence, this process might lead to a retraction update, ensuring the timely and accurate adjustment of alerts in response to retractions.

- If the message involves a neutrino time report, i.e. a Coincidence Tier message, the *CommandHandler* validates the message content before authorizing the main coincidence script to search for coincidences. These messages represent reported initial neutrino observation times from different detectors. Each message undergoes validation, and if it meets the specified criteria, it actively contributes to the ongoing coincidence search. Furthermore, the *CommandHandler* identifies whether the message is intended for production or testing by examining the `"is_test"` field within the message content. The test messages, designed for testing purposes, bypass the usual time checks. On the other hand, all other neutrino times undergo sanity checks, ensuring their validity, and only messages from the last 48 hours are accepted for further processing. This validation process ensures accuracy and effectiveness in the detection of potential galactic core-collapse supernovae events.

The process is illustrated in Figure 3.4. After completion of the validation process outlined above, the *CommandHandler* initiates the next step in the process. Once the *CommandHandler* gives a green light, the coincidence search scripts are activated to begin searching for coincidences among the different messages received. This phase in the workflow ensures that potential coincidences are identified and properly evaluated for further action.

3.3.3 Coincidence Forming and Alert Distribution

The primary functions of the system are to identify Coincidence Tier messages that occur within a specified time window, generate an alert message, and distribute this alert to the appropriate Kafka topics. This system is designed to accurately detect coincidences and promptly provide critical alerts for potential galactic core-collapse supernova events.

Once the *CommandHandler* approves the processing of these messages, the coincidence script, utilizing the *CoincidenceDistributor* module, incorporates each new message into its cache. With every new message in the cache, a sequence of coincidence checks unfolds. The messages are initially inserted into the existing cache, with the first message designated as subgroup=0. Subsequent messages are then compared to this initial subgroup for coincidences. If no coincidence is found, the new message is labeled subgroup= 1, and this pattern continues. This process is handled by *CacheManager*. In the coincidence search process, each new message is scrutinized against the neutrino times in each subgroup. If the message originates from an experiment already present in the subgroup, checks are conducted to determine whether this update triggers a new coincidence. Likewise, retraction messages trigger subgroup inspections. When a coincidence is formed, additional checks are performed to determine if it is a newly established coincidence, in which case the alert message is labeled `"New Coincidence"`, or if it involves an existing coincidence being modified by a new experiment joining or an existing one retracting. If it is the latter, the messages are labeled `"Alert Update"`. At each step, the cache undergoes appropriate adjustments, and the neutrino times are consistently sorted based on the earliest time in the subgroup, and the neutrino time differences relative to the first neutrino time in the subgroup are recomputed.

The information flow is visually represented in Figure 3.4, where each module is represented by a rounded colored box. Upon identification of a coincidence, the *AlertDistributor* module is activated, initiating the creation of a comprehensive alert message. This message encompasses crucial details, including the names of the experiments involved in the coincidence, their respective initial neutrino times, time differences between experiments, subgroup numbers, information indicating whether it is a new alert or an update, and a computed false alarm probability. To calculate the probability of a false alarm, the module collaborates with the *HeartbeatHandler*, employing a Poisson probability computation based on the number of active detectors at that particular moment.

Following its creation, the alert is distributed through multiple channels. Initially, it is broadcasted to the Kafka alert topic, ensuring rapid distribution across the network. Simultaneously, the alert can be sent via email to designated alert mailing lists and transmitted to internal Slack channels through a dedicated Slack bot. The flexibility of the system

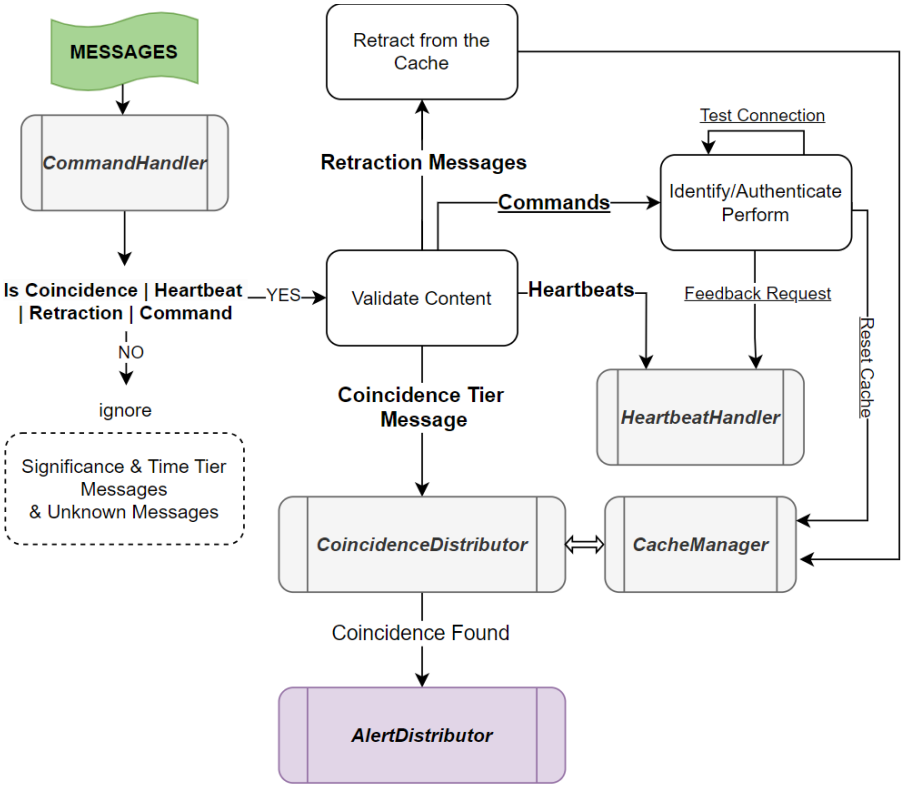

Figure 3.4: The process flow of `snews_cs`. Incoming messages undergo initial processing by the *CommandHandler*, which identifies and filters out time and significance tier messages, along with unknown message types. Validated message types are then routed to their respective functions. Retraction messages are handled by the *CacheManager* and removed from the cache. Commands are identified, authorized, and executed accordingly. Heartbeats are directed to the *HeartbeatHandler* for registration, while coincidence messages are forwarded to the *CoincidenceDistributor* for coincidence search.

allows for the construction of plugins that consume Kafka alerts and send notifications through various platforms such as GCN notices, VOEvent packets, and other communication channels[iii]. Efforts are also underway to establish a connection between the Kafka alert topic and smartphone notifications through a dedicated app, further enhancing the accessibility and promptness of alerts.

3.3.4 Heartbeat Handling and Feedback Mechanism

The processing of heartbeats occurs through a distinct script on the server, separate from the primary coincidence script. When the *CommandHandler* receives a heartbeat, it transmits the message to the *HeartbeatHandler* module. Similar to the *CoincidenceDistributor*, this module maintains a cache and logs the incoming heartbeat messages for each detector. Heartbeat messages carry `"detector_status"` information and an optional `"machine_time"` (refer to

[iii] GCN (Gamma-ray Coordinates Network) notices and VOEvents (Virtual Observatory Events) are two forms of astronomical alert messages used to communicate information about transient astronomical events.

Table 3.1), coupled with a `"sent_time"` field upon execution by `snews_pt`. Upon reception, these messages are time-stamped with a `"received_time"` and latencies are calculated by assessing the time delay between message injection into the Kafka stream and consumption by the *HeartbeatHandler* module on the server.

The *HeartbeatHandler* organizes these messages within a dynamic Pandas DataFrame, retaining entries from the last 48 hours for each detector. This dataframe undergoes continuous updates, ensuring that the latest entries are incorporated while removing the older ones. To mitigate potential crashes or serve as a backup, this dataframe is consistently written to a file. This approach also enables the storage of heartbeat records per day in a separate file, accommodating longer storage needs. The flexibility of this storage method aligns with the specific agreements made with individual experiments. Importantly, all this information remains confidential and only resides on the server. Access to such data is only granted through user feedback requests or in scenarios where a coincidence alert is triggered, utilizing the data to compute false alarm probabilities.

Individual experiments may occasionally experience false triggers, as dictated by Poisson statistics. The probability of observing k events or more in a counting experiment is given by

$$P(X \geq k) = 1 - \sum_{i=0}^{k-1} \frac{e^{-\lambda}\lambda^i}{i!}, \tag{3.1}$$

when expecting an average of λ events.

According to the probability mass function of such a model, detectors may encounter signals that mimic (here, the frequency of such events is referred to as *imitation frequency*) the expected counts they would observe in the event of a supernova. To mitigate this, SNEWS requires each individual experiment to adjust their thresholds such that reported observations are deemed false at most once per week. This practice helps minimizing the occurrence of accidental coincidences between different detectors.

Although individual false alarm rates may not strictly follow Poisson statistics and can fluctuate over time, we can estimate the expected frequency of false alarms by assuming a uniform false alarm rate of once per week for each experiment. The imitation frequencies of the coincidences can then be computed as a function of the number of detectors available and for a required n-fold coincidence, as illustrated in Figure 3.5.

The imitation frequencies are calculated using the combined rate R_{comb},

$$R_{\text{comb}} = (C(n,r) + 1) \times F_{d1} \times F_{d2} \times \ldots \times F_{d_r} \times \delta t^{(r-1)} \tag{3.2}$$

which corresponds to the inverse of the imitation frequency, and F_{d_i} denotes the false alarm rate of the ith detector. The term $C(n,r)$ represents the combination of n options with r possibilities, and δt represents the time window.

In Figure 3.5, it is evident that even in the extreme scenario where 15 experiments are actively collecting data and only a 2-fold coincidence is observed, the imitation frequency is approximately once every 10 years.

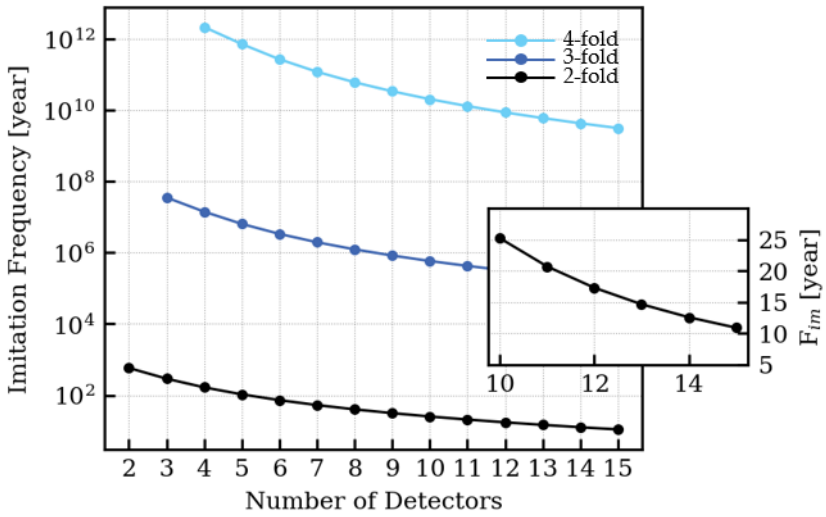

Figure 3.5: The imitation frequency, F_{im}, of n–fold coincidences among N experiments, considering a 10–second coincidence window and individual alarm rate of one per week. The lowest imitation frequencies occur for 2-fold coincidences involving a large number of operational experiments. For example, a false coincidence between 2 detectors out of 15 operational ones happens once every ≈ 10 years (see inlet with linear scale).

In the event of a triggered alert within `snews_cs`, heartbeats play a central role in assessing the number of detectors available at the time of the event. Using the number of detectors that contribute to the formed coincidence, the system estimates how often such occurrences could be attributed to individual false alerts imitating a coincidence.

3.3.5 Logic of Distributed Server

Ensuring stable operation of the server and the comprehensive search for coincidence occurrences in received observation messages are essential. To maintain server stability and guarantee continuous operation, requires the implementation of backup servers. In addition, to prevent the simultaneous triggering of multiple alerts by different servers in the event of a coincidence, distributed server logic is being developed.

The primary server is currently operational at Purdue University, Indiana, USA, serving as the central hub for processing observation messages and detecting coincidences. In addition, plans are underway to establish two backup servers, strategically located in Oxford, UK and Karlsruhe, Germany. These backup servers will serve as fail-safes, ready to assume the role of the primary server should it become unavailable due to unforeseen circumstances. The backup server is expected to be lightweight and can be hosted on any machine with minimalistic requirements [113].

The distributed server logic involves establishing communication channels between the primary server and the backup servers, enabling seamless data transfer and task delegation. In the event of a server outage or failure, the backup servers will automatically assume the responsibilities of the primary server, ensuring uninterrupted operation and timely detection of coincidences.

By implementing a distributed server architecture, redundancy was introduced into the system, enhancing reliability and resilience. This approach not only safeguards against potential server failures, but also mitigates the risk of alert duplication, thereby minimizing confusion and streamlining the alert dissemination process.

The SNEWS system is constantly undergoing significant upgrades, incorporating several enhanced systems and logics, as well as welcoming new experiments. One of the notable changes in SNEWS 2.0 is the adoption of more relaxed, flexible, and interpretable Python software developed within this thesis. This new software provides a user-friendly interface, allowing users to easily publish their observations, while a robust coincidence system continuously monitors for coincidences between experiments [113]. These new functionalities in the software significantly improve the efficiency of SNEWS communication. In addition, new server logics, such as distributed backup servers, are being implemented to enhance system reliability. With these upgrades, SNEWS employ the software developed within this thesis for its future operations, ensuring a more efficient and responsive network for detecting potential galactic core-collapse supernovae events.

4

Supernova Neutrino Signatures in Dual–Phase TPCs

"What we know is a drop, what we do not know is an ocean."
–Isaac Newton

Contents

In the previous chapters, we introduced key components of the Standard Model, delved into direct DM detection with a particular focus on LXe TPC detectors, and provided a detailed overview of the XENONnT experiment. We further explored supernova physics, covering mechanisms of neutrino emission, and introduced recent upgrades for the SNEWS 2.0, including the development of specialized tools within this work.

This chapter begins with an in-depth explanation of supernova neutrino emission models, covering various simulations for different progenitor stars in Section 4.1. Following this, we calculate the expected neutrino interactions in LXe TPCs using the XENONnT experiment. We present an analysis of the expected interaction rates in Section 4.2. In addition, we describe methods for simulating these interactions in XENONnT under realistic detector conditions. In Section 4.3, we introduce custom software developed as part of this work to enhance the analysis capabilities of supernova neutrino detection in XENONnT.

4.1 Understanding Supernova Models and Neutrino Fluxes

In Chapter 1 the physics of supernova explosions was presented, with a focus on modeling efforts using sophisticated computational techniques and other approaches, including observations such as SN1987A. Despite extensive research, different modeling techniques have not yet fully agreed on one model [139, 140, 141, 142], highlighting the complexities involved in accurately modeling supernova events. In this section, established and emerging supernova models are explored, emphasizing model-dependent luminosity and energy distributions and their dependence on the properties of the progenitor star.

The different models are incorporated into the Python software package called `snewpy` developed under SNEWS 2.0. Key properties such as luminosities, neutrino energy spectra, and neutrino fluxes for all flavors are accessible through `snewpy`. The package provides a unified and accessible platform for exploring these models. Using selected established models, the time and neutrino energy distributions of supernova events were analyzed. To understand the expected supernova neutrino signal on Earth, the flux from model-dependent luminosity and energy profiles was investigated, and different characteristics are discussed in the following subsections.

4.1.1 Characterizing Neutrino Fluxes

Neutrino bursts originating from the core of a supernova travel through space, exhibiting varying number densities of neutrinos at different energies depending on the predicted luminosity. Consequently, the flux depends on both the luminosity and the neutrino energy spectrum as a function of time. Neutrino detection on Earth possess two primary features: the spectrum of their energy and the emitted flux over time in relation to the post-bounce phase (t_{pb}) of the core collapse.

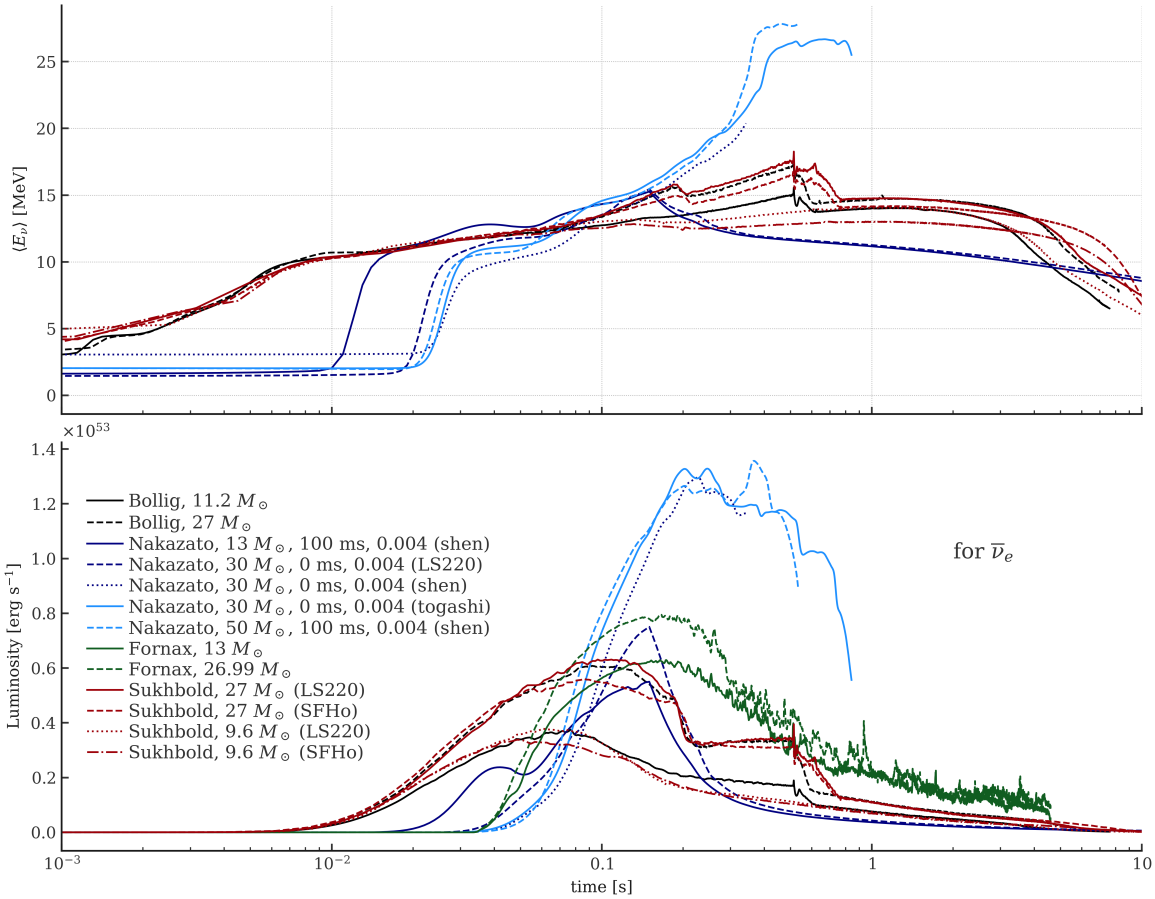

Figure 4.1: The mean neutrino energy evolution with time for electron antineutrinos for all models considered in this study (*top*), and the luminosity distribution (*bottom*).

As the energy distribution of these neutrinos varies over time, it leads to different mean energies across the various phases of a supernova event. The top panel of Figure 4.1 shows the energy distribution at a given time, while the bottom panel depicts the mean energy distribution over time for all models in this work.

Models typically provide time-dependent luminosities $L_{\nu_\beta}(t_{pb})$, for each neutrino flavor ν_β, including electron, muon, and tau neutrinos, as well as their antineutrinos. These luminosities represent the total energy radiated by the supernova per unit time as a function of post-bounce time (t_{pb}). In addition, models often supply the energy distribution of neutrinos at a given time, denoted by $\Phi_{\nu_\beta}(E_\nu, t_{pb})$, for each neutrino flavor. When the exact neutrino energy distribution is not available, tools such as `snewpy` employ an approximate parameterization based on the pinched thermal spectrum introduced by [143]. This parameterization describes how the neutrino energies are distributed for a given flavor and post-bounce time and is often modeled using a gamma distribution:

$$\Phi_{\nu_\beta}(E, t_{pb}) = A(t_{pb}) \left(\frac{E}{\langle E_{\nu_\beta}(t_{pb}) \rangle} \right)^{\alpha_\beta(t_{pb})} \exp\left[-\frac{(\alpha_\beta + 1)E}{\langle E_{\nu_\beta}(t_{pb}) \rangle} \right], \tag{4.1}$$

where the normalization constant $A(t_{pb})$ ensures that the total probability distribution integrates to 1, that is, $\int dE\, \Phi_{\nu_\beta}(E, t_{pb}) = 1$.

The parameter α_β, known as the pinching parameter, quantifies the deviation of the energy distribution from a pure thermal spectrum. It is determined by the ratio of the second moment to the square of the first moment of the energy distribution:

$$\alpha_\beta = \frac{2\langle E_{\nu_\beta}\rangle^2 - \langle E^2_{\nu_\beta}\rangle}{\langle E^2_{\nu_\beta}\rangle - \langle E_{\nu_\beta}\rangle^2}. \tag{4.2}$$

Using the time-dependent luminosity and neutrino energy distribution, the neutrino flux, f_{ν_β}, (defined as the number of neutrinos arriving at the detector per unit area and per unit time) can be calculated for a supernova at a distance d [144]:

$$f_{\nu_\beta}(E_\nu, t_{pb}) = \frac{L_{\nu_\beta}(t_{pb})}{4\pi d^2} \frac{\Phi_{\nu_\beta}(E_\nu, t_{pb})}{\langle E_{\nu_\beta}(t_{pb})\rangle}. \tag{4.3}$$

Understanding the fluxes in different neutrino flavors can also provide helpful information, although detection in the CEνNS channel is flavor–blind. Detectors may be more effective with models that predict specific neutrino energy and flavor distributions. Liquid scintillators and water Cherenkov detectors are sensitive to specific neutrino flavors primarily through inverse beta-decay interactions. For example, if a model suggests a higher flux of electron antineutrinos, detectors optimized for inverse beta decay are more likely to observe the event. In the next galactic supernova, the total energy emitted in all flavors will be crucial in understanding the mechanisms driving the supernova explosion. In particular, CEνNS interactions remain unaffected by neutrino oscillations, which can alter the flavor contributions that reach the detectors.

4.1.2 Modeling Supernova Mechanisms

The physics of CCSN is extensively studied and different models are developed by different groups. Simulations of different CCSN models often include multiple variations of the same stellar evolution model, such as different progenitor masses or various equations of states (EOS). The diversity of models and assumptions results in a broad spectrum of potential outcomes.

This study compares widely used models and progenitors, selected for their frequent use in recent publications. Models with similar progenitor masses were chosen for easy comparison. In addition, variations within the same model were examined to assess differences from mass changes. The selected models with assumed properties are summarized in Table 4.1. For easier reference, the `snax` index, which is an identifier for the custom software developed in the context of this work, is also listed in this table.[i]

The Bollig 2016 [145] and Nakazato 2013 [146] models are widely used in supernova neutrino studies due to their detailed, multidimensional simulations and wide range of progenitor masses. Bollig's model incorporates state-of-the-art treatment of neutrino transport and

[i] The `snax` software is introduced in Section 4.3.

Table 4.1: List of selected supernova models and their progenitors. Also listed is the `snax` id which is an identifier for each model progenitor, mass of the progenitor in units of solar masses, $M_\odot$, the supernova shock revival time, metallicity (Z) of the progenitor star, and the assumed EOS of the matter. Some properties, such as metallicity, revival time, or EOS, are not specified (denoted as *NS) for certain simulations.

Model name	snax id	mass	$t_{\mathbf{revival}}$	Z	EOS
Bollig 2016	1	11.2 $M_\odot$	*NS	*NS	LS220
Bollig 2016	2	27.0 $M_\odot$	*NS	*NS	LS220
Nakazato 2013	2	13.0 $M_\odot$	100.0 ms	0.004	shen
Nakazato 2013	13	30.0 $M_\odot$	0.0 ms	0.004	LS220
Nakazato 2013	14	30.0 $M_\odot$	0.0 ms	0.004	shen
Nakazato 2013	15	30.0 $M_\odot$	0.0 ms	0.004	togashi
Nakazato 2013	20	50.0 $M_\odot$	100.0 ms	0.004	shen
Fornax 2021	2	13.0 $M_\odot$	*NS	*NS	*NS
Fornax 2021	15	26.99 $M_\odot$	*NS	*NS	*NS
Sukhbold 2015	1	27.0 $M_\odot$	*NS	*NS	LS220
Sukhbold 2015	2	27.0 $M_\odot$	*NS	*NS	SFHo
Sukhbold 2015	3	9.6 $M_\odot$	*NS	*NS	LS220
Sukhbold 2015	4	9.6 $M_\odot$	*NS	*NS	SFHo

stellar evolution physics. Nakazato models, in addition to offering a wide range of progenitor masses and different equations of state, provide the feature of simulating failed supernovae, where the core collapses into a black hole, resulting in a sharp cut-off in the neutrino signal. These failed supernovae represent unique astrophysical scenarios that are worth investigating. Both models have been extensively referred to in the recent literature [47, 147], underscoring their reliability and relevance to neutrino detection experiments.

The Sukhbold 2015 [148] model is notable for studying progenitors with different EOS, probing their impact on collapse and neutrino emissions. The Fornax 2021 [142] model uses 3D simulations with realistic physics, including turbulence and rotation, offering insight into supernova neutrino signals.

Together, these models provide a comprehensive framework for studying supernova neutrino detection, covering a wide range of astrophysical phenomena, allowing comparisons and understanding of supernova physics. To ensure consistency and provide a standard reference throughout this thesis, the Bollig 2016 model with the 27 $M_\odot$ progenitor is used as a benchmark. This approach allows for a coherent comparison across different scenarios and provides a deeper understanding of how specific parameters influence outcomes.

Models usually include four neutrino flavors: electron neutrinos (ν_e), electron antineutrinos ($\bar{\nu}_e$), and a combined category for muon (μ) and tau (τ) neutrinos, and their antineutrinos, represented as ν_x and $\bar{\nu}_x$, respectively. The luminosities and energy spectra of each flavor are characterized by different models. Another important aspect is that the time evolution of these fluxes is also unique in each simulation. The mean energy distribution for each neutrino flavor over time is shown in the left panel of Figure 4.2, while the luminosity distributions for different flavors are shown on the right. For a comparison of $\bar{\nu}_e$ luminosities see Figure 4.1.

The apparent variations in the peaks observed in Figure 4.1 may suggest significant differences

between the models. However, this can be deceptive since the initial timing of the models is arbitrary. In addition, amplitude variations are dependent on distances, which can lead to different models producing analogous signal patterns. Thus, distinguishing between models is complex, and methods for doing so will be addressed in Chapter 6.

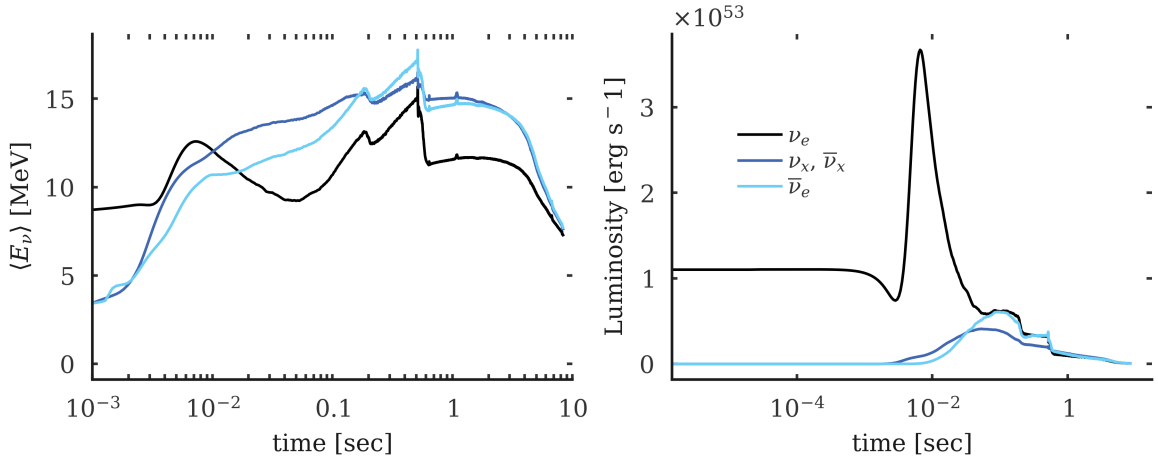

Figure 4.2: The mean energy distributions (*left*) and the bolometric luminosity distributions (*right*) of different neutrino flavors as a function of time for the Bollig 2016 model with a 27 $M_\odot$ progenitor.

4.2 Expected Supernova Signals in XENONnT

The investigation of neutrino signals from supernovae relies on precise modeling and simulation techniques to predict observables such as prompt scintillation (S1) and proportional scintillation (S2) signals in dual-phase LXe TPCs such as XENONnT. This section outlines the methods and tools employed to simulate and project the observables associated with supernova neutrinos.

The expected interaction rates at different recoil energies and times for various supernova models, which form the basis of the signal simulation, are discussed in Section 4.2.1. First, a simplified approach to modeling the expected signals is presented, followed by a more advanced technique.

The basic modeling approach, described in Section 4.2.2, uses neutrino fluxes from the luminosity curves of a selected supernova model. For each recoil energy, the resulting photon and electron yields are calculated using semi-analytical models in the Noble Element Simulation Technique (NEST) [149]. Detector effects and the final expected signal are then estimated statistically to provide an initial understanding of neutrino interactions in xenon detectors.

In Section 4.2.3, more detailed simulation techniques are introduced using the `fuse` (Framework for Unified Simulation of Events) toolkit [150], which integrates detailed microphysics and detector models. This advanced framework enables deeper statistical analysis, including survival probabilities, detection efficiencies, and signal generation probabilities, allowing refined predictions and insights into the detector response to supernova neutrinos.

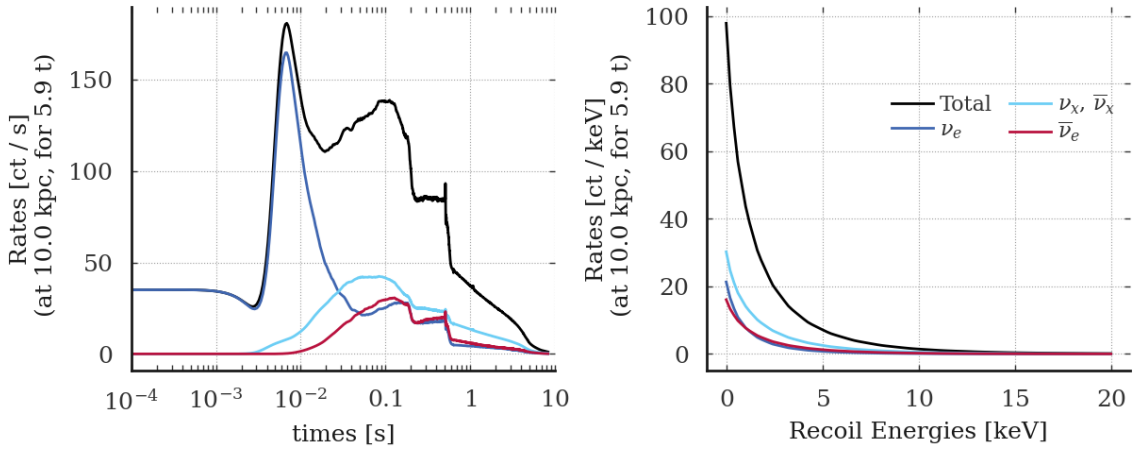

Figure 4.3: The expected differential interaction rates in XENONnT for Bollig 2016 model with 27 $M_\odot$ at 10 kpc are shown for individual neutrino flavor fluxes as well as the sum over all fluxes. On the left (right) panel, rates are integrated over all possible recoil energies (times) and shown as a function of time (recoil energy).

4.2.1 Differential Interaction Rates

Neutrinos from CCSNe are expected to interact with the xenon atoms within the XENONnT TPC, causing the xenon nuclei to recoil with varying energies according to their interaction cross–sections, leading to different observable signals. The number of neutrinos arriving on Earth decreases as a function of the distance from the source. For example, the neutrino flux of a supernova at 10 kpc will be reduced by a factor 10^{48} arriving on Earth compared to the flux at the source. The differential rate of these neutrino interactions can be expressed as

$$\frac{\mathrm{d}^2 R}{\mathrm{d}E_R \mathrm{d}t_{pb}} = \sum_{\nu_\beta} N_{\mathrm{Xe}} \int_{E_\nu^{\min}} \mathrm{d}E_\nu \, f^0_{\nu_\beta}(E_\nu, t_{pb}) \, \frac{\mathrm{d}\sigma}{\mathrm{d}E_R}(E_\nu, E_R), \tag{4.4}$$

as a function of post-bounce time t_{pb} and recoil energies E_R, where $E_\nu^{\min}$ is the minimum neutrino energy required to produce a recoil of energy E_R, $f^0_{\nu_\beta}$ is the initial neutrino flux for flavor β, $\mathrm{d}\sigma/\mathrm{d}E_R$ is the differential cross–section for neutrino scattering that produces recoil energy E_R, and N_{Xe} is the number of xenon targets.

To accurately simulate and analyze the supernova interactions within the XENONnT TPC, the time-integrated differential rates were computed first. This approach allows for the investigation of the cumulative effect of neutrino interactions over the duration of the supernova event and the extraction of the recoil spectrum shown in the right panel of Figure 4.3. The differential rate of recoil events at a given recoil energy, integrated over time from t_0 to t_f, is given by:

$$\frac{\mathrm{d}R}{\mathrm{d}E_R} = \sum_{\nu_\beta} N_{\mathrm{Xe}} \int_{t_0}^{t_f} \mathrm{d}t_{pb} \int_{E_\nu^{\min}} \mathrm{d}E_\nu \, f^0_{\nu_\beta}(E_\nu, t_{pb}) \frac{\mathrm{d}\sigma}{\mathrm{d}E_R}(E_\nu, E_R). \tag{4.5}$$

Additionally, the differential rate of recoil events per unit time, integrated over all possible recoil energies, was derived, providing the total interaction rate over time, which is crucial for assessing the signal timing characteristics. This is expressed as:

$$\frac{\mathrm{d}R}{\mathrm{d}t_{pb}} = \sum_{\nu_\beta} N_{\mathrm{Xe}} \int_{E_R^{\min}}^{E_R^{\max}} \mathrm{d}E_R \int_{E^\nu_{\min}} dE_\nu \; f^0_{\nu_\beta}(E_\nu, t) \; \frac{\mathrm{d}\sigma}{\mathrm{d}E_R}(E_\nu, E_R) \tag{4.6}$$

where $E_R^{\min}$ and $E_R^{\max}$ are the minimum and maximum recoil energies detectable by the detector. This integral provides the total rate of interactions per unit time, shown in the left panel of Figure 4.3 for the benchmark model, accounting for all neutrino energies and all detectable recoil energies.

Both the recoil spectra and the time distributions are needed to understand the detector's response to supernova interactions. They allow for the calculation of expected signal rates and their distribution over time and energy. The total integrated rates shown in Figure 4.3 were used to determine the number of expected interactions. It is important to note that these expected interactions vary depending on the supernova's distance and the detector's energy threshold. The comparison of the total expected interactions among the various models examined is shown in Figure 4.4. Although in reality the experiments have an energy threshold, the rates were calculated down to 0 keV, implying a perfect detection sensitivity, represented by blue bars. The same figure also uses a more realistic 0.7 keV detection threshold to recalculate the expected counts, marked with black bars.

In practice, the energy threshold for detection is not a sharp cut–off. Instead, detection efficiency gradually increases with recoil energy, starting lower at low energies and approaching full efficiency at higher levels. It is noticable that, while some models show higher expected interaction counts when all interactions are considered, the numbers change when an energy threshold is applied. This indicates that a significant fraction of interactions in these models occur below the assumed energy threshold. The distribution of the number of expected interactions as a function of the energy threshold is also shown in Figure 4.5.

Although the overall anticipated interaction rates are informative, examining the temporal and energy distribution of these interactions can lead to a better understanding of the signal characteristics. The next section describes simulations that employ various supernova model profiles as input to explore each model in greater depth.

4.2.2 Approximate Modeling of Neutrino-Induced Signals

A simplified approach can be used to approximate the expected signal after an interaction. Starting from the recoil energy E_R, the average number of quanta (photons N_{ph} and electrons N_{el}) produced can be estimated as

$$\langle N_{ph} \rangle = E_R \times L_y(E_R), \quad \langle N_{el} \rangle = E_R \times Q_y(E_R) \tag{4.7}$$

where $L_y(E_R)$ is the light yield, and $Q_y(E_R)$ is the charge yield at recoil energy E_R.

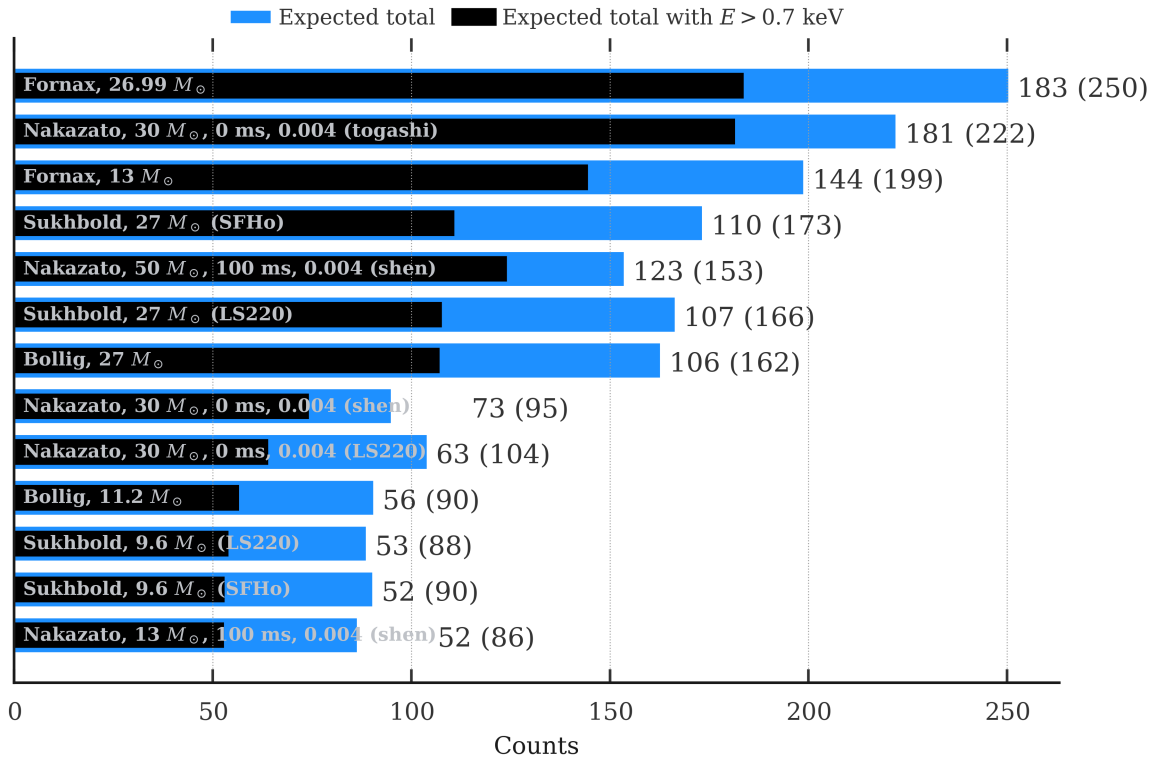

Figure 4.4: Expected number of interactions in XENONnT for different supernova models at 10 kpc, computed by integrating eq. (4.5) (or eq. (4.6)). The blue bars assume perfect detection efficiency down to 0 keV interactions. The largest number of interactions is expected from the Fornax 2021 model with 26.99 $M_\odot$ progenitor, while the lowest numbers are obtained from the Nakazato 2013 model with 30 $M_\odot$, shen EOS, 100 ms revival time and 0.004 metallicity progenitor. A more conservative step function energy threshold is applied for the same calculation, and the expected interactions above 0.7 keV are shown with the black bars. The exact numbers in both cases are also shown next to each bar.

The number of photons detected by the PMTs is then approximated using a binomial distribution with the photodetection efficiency g_1:

$$N_{PE} = \text{Binomial}(N_{ph}, g_1), \tag{4.8}$$

and the S1 signal is modeled by a Gaussian distribution centered on N_{PE} with variance proportional to $\sqrt{N_{PE}}$:

$$S1 = \text{Gauss}(N_{PE}, \alpha\sqrt{N_{PE}}). \tag{4.9}$$

For the S2 signal, the survival probability of electrons depends on the electron lifetime τ, drift velocity v_d, and the distance Δz they travel:

$$p_{\text{sur}} = \exp\left(-\frac{\Delta z}{v_d \tau}\right), \tag{4.10}$$

with the number of surviving electrons $\tilde{N_{el}}$ sampled as:

$$\tilde{N_{el}} = \text{Binomial}(N_{el}, p_{\text{sur}}). \tag{4.11}$$

Finally, the S2 signal is modeled with a Gaussian distribution around the surviving elec-

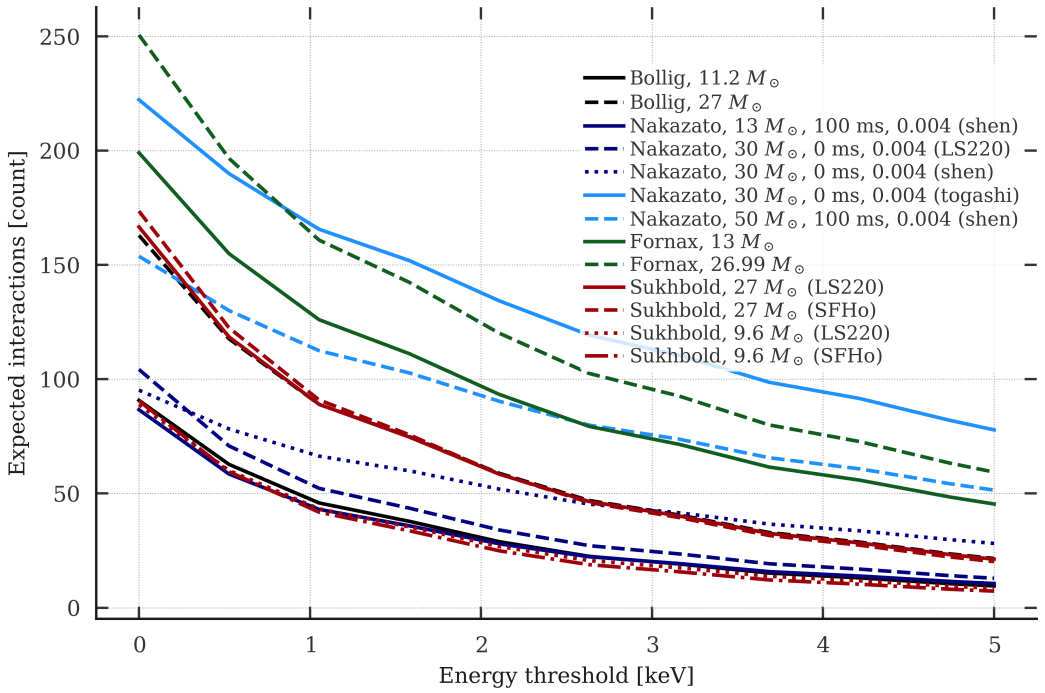

Figure 4.5: Number of expected interactions in XENONnT above a specified detection threshold (step function). The number of expected counts decreases more rapidly for some models due to a significant fraction of neutrinos having lower energies, resulting in lower expected recoil energies.

trons:

$$S2 = \text{Gauss}(\tilde{N_{el}}, \beta\sqrt{\tilde{N_{el}}}). \tag{4.12}$$

By applying these processes over the recoil energy distribution and accounting for interaction times, a simplified approximation of the detector's response to neutrino interactions can be obtained.

The coefficients α and β used in this basic modeling approach to estimate the signal distributions of S1 and S2 are generally derived from experimental data, providing a simplified view of signal generation in dual-phase LXe TPCs. This approximate modeling serves as an introduction to the essential dynamics of how neutrino interactions generate observable signals within the detector. By offering a foundational understanding, this method sets the stage for the more sophisticated simulations discussed in the next section, which incorporate non-linear effects such as electric field distortions and light collection efficiencies. These advanced simulations give a more accurate and comprehensive picture of the detector response, improving our ability to interpret supernova neutrino signals.

4.2.3 Advanced Simulation Techniques

Although simplified models offer a basic understanding of neutrino interactions and how the observed signals such as S1 and S2 are generated, more advanced simulations are essential to accurately capture the complexities of the real detectors. These simulations account for key

detector properties, such as light collection efficiency (LCE), photoelectron gain, and spatial variations in response, all of which influence the observed signals.

For this purpose, the `fuse` framework was developed for the XENONnT experiment. This tool generates realistic detector signals by incorporating data-driven characterization and fine-tuning detector parameters based on real calibration measurements, allowing for detailed simulations of nonlinear detector effects like electric field distortions and proper S1 and S2 signal propagation.

The simulation chain, processed by `straxen`, ensures consistency by processing both simulated and real data in the same manner. This approach provides a detailed output, including accurate S1 and S2 signal areas, spatial distributions, and time profiles.

`fuse` can initiate simulations at two distinct levels, each focusing on different aspects of the interaction and detector response:

- **Microphysics Simulation**: This step simulates the initial physical processes at the interaction site using parameters such as position, time, energy deposition, and type of interaction. The simulation chain integrates NEST to generate the appropriate number of quanta based on energy deposition and local electric field conditions. These quanta are then passed to the next level for detector physics simulations.
- **Detector Physics Simulation**: In this stage, the detector response to the generated quanta is determined. This includes simulating the drift of electrons in the TPC under the influence of the electric field, light scattering and reflection within the detector, and other physical processes. The quanta, along with information about the interaction time and position, are combined with auxiliary files representing detector-specific conditions to produce realistic detector signals.

Simulations can be initiated from either level using an *instructions* file containing interaction parameters. Figure 4.6 illustrates the steps involved in microphysics simulation, while Figure 4.7 shows the process for detector physics simulation.

The figures illustrate the structure and workflow of the microphysics and detector physics simulations used in the XENONnT experiment. The microphysics simulation includes light- and charge-yield computations followed by the integration of various simulation volumes. The detector physics simulation involves stages from electron drift and extraction to timing, scintillation light production, and photon propagation, including handling delayed electrons. These simulations ensure detailed and accurate modeling of the physical processes and responses in the XENONnT experiment.

To simulate supernova neutrino CEνNS interactions, the expected interaction numbers at a given distance along with the time and energy profiles were obtained for each SN model as described in Section 4.2.1. An instruction file containing the required information for each interaction was created and used as input for `fuse` simulations. The simulated data were processed with the standard processing software, `straxen`. The final data products are then have the same format as the background data. The area histograms of the S1 and S2 signals using sample background data were compared with the simulated signals of the Bollig 2016 model with a 27 $M_{\odot}$ progenitor in Figure 4.8.

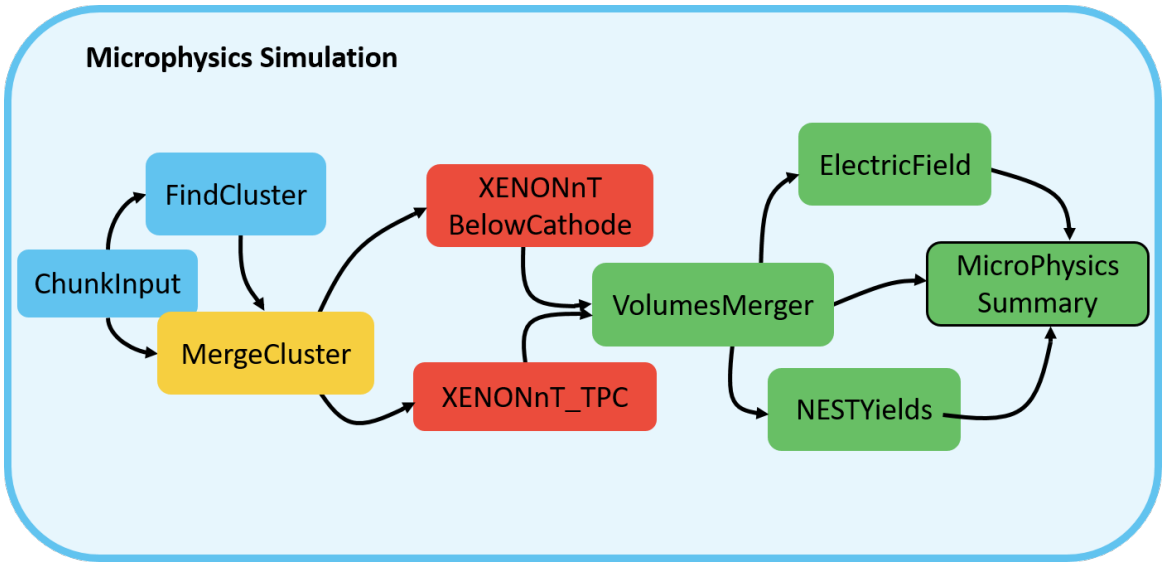

Figure 4.6: Steps involved in the `fuse` microphysics simulation. The simulation generates quanta using NEST based on energy deposition and the local electric field at the interaction site. This process forms the foundation for realistic signal generation in the detector. Figure is adapted from [150].

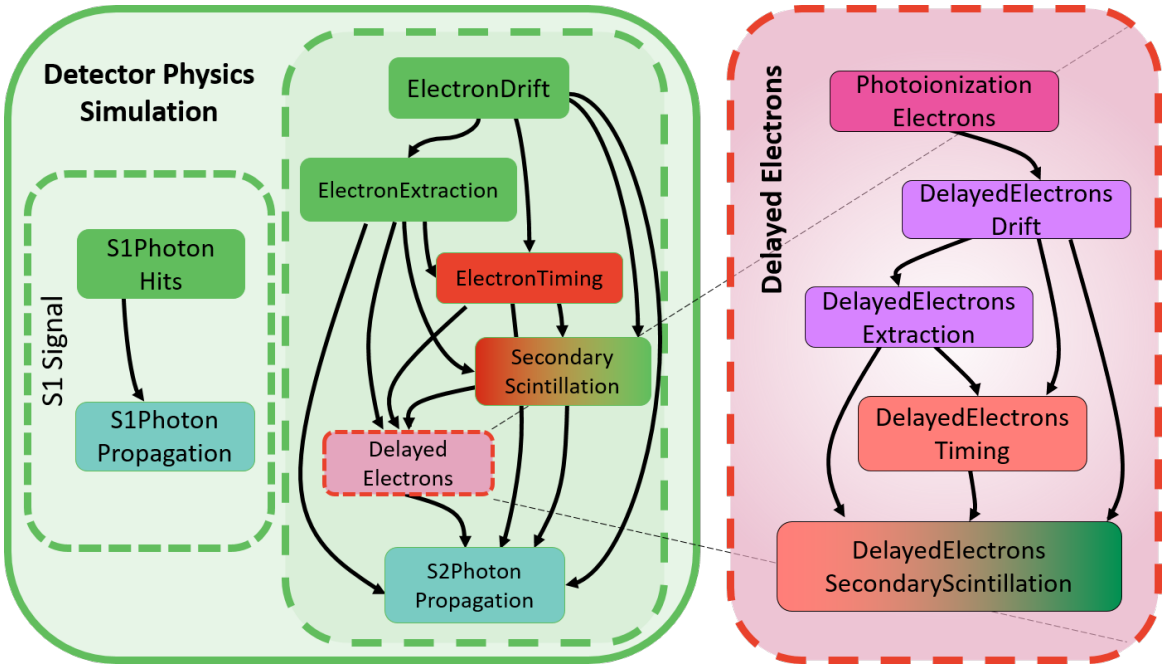

Figure 4.7: Steps involved in the `fuse` detector physics simulation. The two main observable signals in the XENONnT TPC, S1 and S2, are simulated by propagating the results from the microphysics simulations. The S1 signal is simulated by propagating the light within the detector geometry until it is captured by one of the PMTs. For the S2 signal, the drift of electrons to the liquid-gas interface and their extraction to the gaseous phase under specific electric field conditions are simulated. Additionally, the simulation can separately consider delayed electrons. Figure is adapted [150].

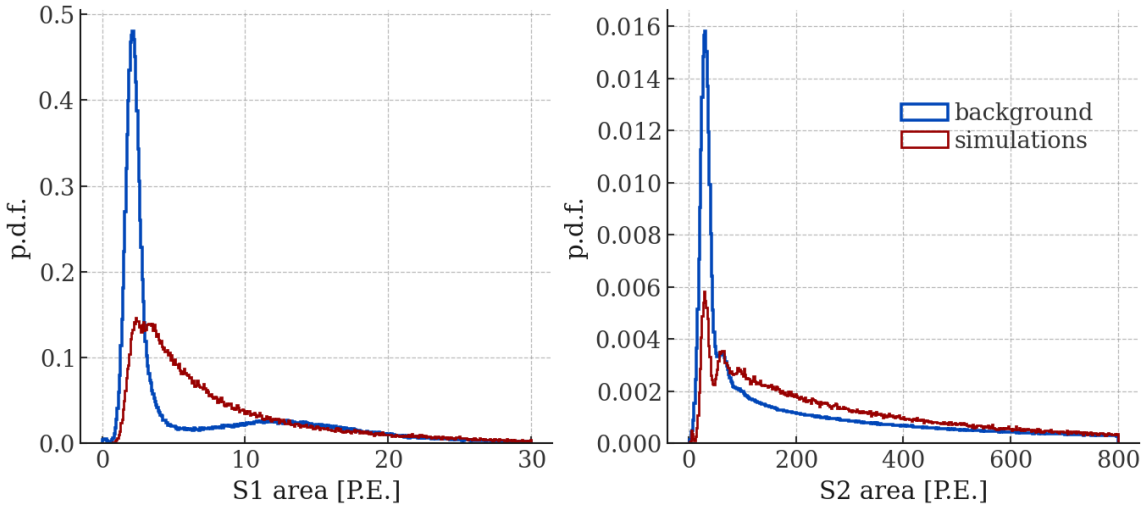

Figure 4.8: Probability density distributions for the S1 and S2 areas of a randomly sampled background data and a simulation using Bollig 2016 model with a 27 $M_{\odot}$ progenitor.

In conclusion, `fuse` offers an advanced tool for generating realistic simulations. These advanced simulations are crucial for predicting detector responses and refining our understanding of the underlying physics driving these interactions. Realistic simulations can help define data selection to discriminate against the background and study our sensitivities for detecting and identifying a potential supernova signal within the real data stream.

For detailed supernova studies within the XENONnT TPC, a custom software package called `snax` was developed. This suite integrates supernova modeling, interactions with xenon atoms, and advanced simulations to study the supernova neutrino interaction in XENONnT TPC.

4.3 SNAX Software Package

The SuperNova Analysis in XENONnT (`snax`) software package was designed to facilitate the investigation of supernova models and their interactions with xenon nuclei. It is available as an open–source software on GitHub [151], making it accessible for collaborative research. By bridging the supernova models from `snewpy` to the `fuse` simulation kit, `snax` brings together all aspects of supernova analysis in the XENONnT detector. The primary functions of `snax` include:

- **Retrieving Supernova Models**: `snax` provides an interfacte to fetch supernova models from `snewpy` and constructs a higher level `snax model` for analysis.

- **Computing Interactions**: It computes the expected interactions between the supernova neutrino flux and the target atoms by solving differential equations discussed in Section 4.2.1 and the cross–sections based on the target atom.

- **Generating Simulation Instructions**: `snax` generates simulation instruction files for the `fuse` toolkit, and provides tools to easily track and manipulate simulations.

The textttsnax was developed to assist with analysis in XENONnT, but its modular design allows for extended use. The current version implements the xenon and argon isotopes as different targets, which can be selected to study variations in experimental setups. Additionally, it supports studying different target masses and source distances, and provides an easy way to examine various form factors. Different supernova models can also be easily integrated.

4.3.1 SNAX Model and Target

Using a model from `snewpy`, a higher-level object called the snax model was built with `snax`. Once a `snax` model object is created, the fluxes in the two-dimensional parameter space (simulated time and neutrino energy) was computed for each flavor using eq. (4.3). Furthermore, the parameters of the supernova models in `snewpy` can be easily investigated within the `snax` model.

An example showing the flux density as a function of the neutrino energy and time generated by `snax` is presented in Figure 4.9 left and right panels, respectively. For each processed model, `snax` assigns a unique deterministic hash code and caches the computed fluxes, allowing easy access and retrieval at a later time.

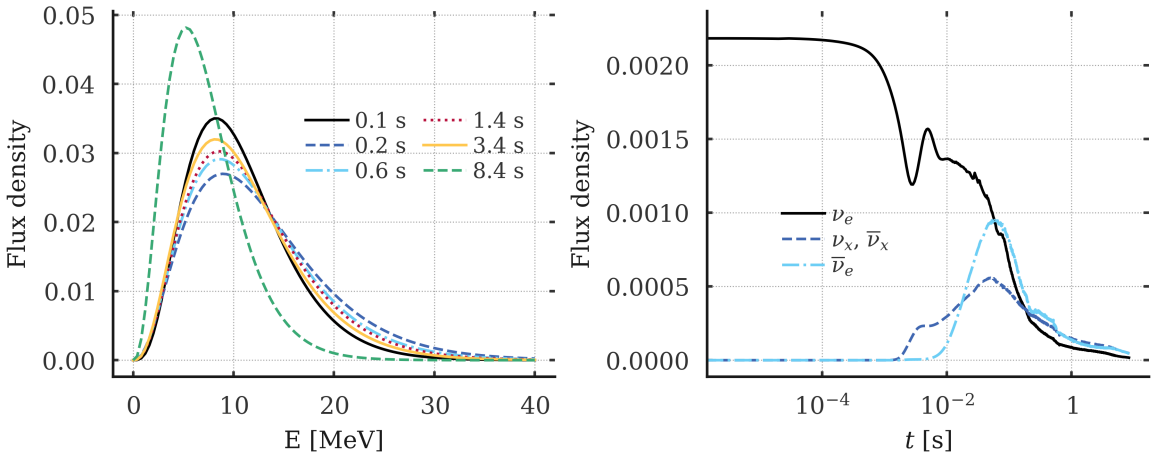

Figure 4.9: Comparison of neutrino energy spectra and time distribution of flux density for the Bollig 2016 model with a 27 $M_\odot$ progenitor. Energy spectra of electron neutrinos (*left*) at various time intervals. The plot displays the normalized flux density as a function of neutrino energy, E, for different times, ranging from 0.1 seconds to the final time step of the simulation. Each line style represents a different time, with the color gradient indicating the progression of time from early to late stages. Time distribution of the flux density for different neutrino flavors (*right*) at a fixed neutrino energy of 5 MeV. The plot shows the normalized flux density as a function of time on a logarithmic scale. Each line style and color represents a different neutrino flavor, highlighting the temporal evolution of the flux density for each flavor.

To compute the interactions of these supernova neutrinos with matter, `snax` introduces a target object. The Xe and Ar atoms are predefined along with their isotopes in `snax`.

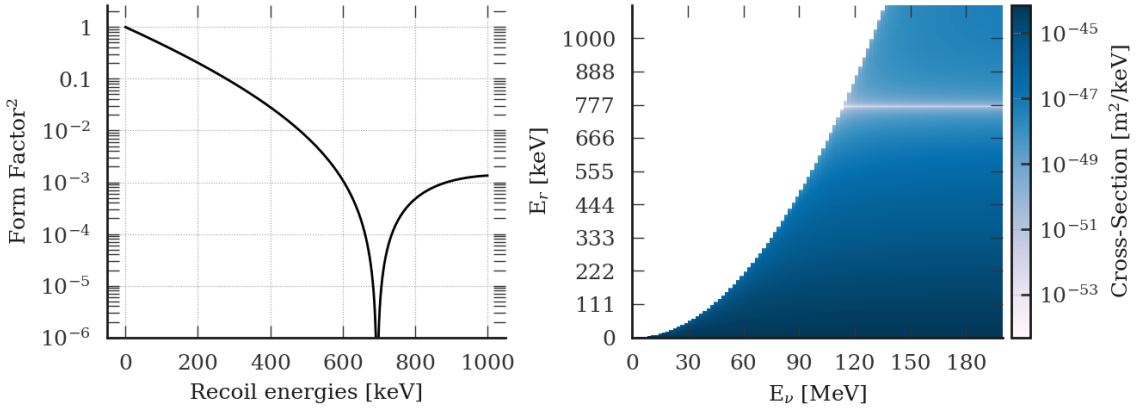

Figure 4.10: The squared Helm form factor for the most abundant argon isotope and the corresponding cross–section probabilities for different neutrino and recoil energies.

However, users can provide any other target to compute interaction cross–sections. The target isotope determines values for the parameters defined in the form factor, such as A and Z, as well as cross–section probabilities. The properties of xenon isotopes are summarized in Table 2.1.

Given the choice of the target isotope, the Helm form factor is first calculated in `snax` and subsequently used to compute the cross–sections as discussed in Section 2.1.4. An example of the Helm form factor along with the cross–section calculated for the most abundant xenon isotope (^{131}Xe) was shown in Figure 2.3. Similarly, Figure 2.4 presents the individual form factors for the isotopes listed in Table 2.1 as well as three argon isotopes. As discussed in Section 2.1.4, the specific target isotope with which a neutrino interacts significantly affects the cross–section and, consequently, the recoil energy.

For future neutrino-nucleon interaction studies in DM experiments, additional noble elements may be incorporated. Argon, for instance, is already available as an option in `snax`. The squared form factors for various argon isotopes were displayed in the bottom panel of Figure 2.4, and cross–section probabilities as a function of neutrino and recoil energies are presented in Figure 4.10. Additionally, users can select a specific isotope to isolate interactions with that isotope, enabling detailed analyses of isotope-specific effects. By default, `snax` samples targets based on the natural abundance of each isotope.

4.3.2 SNAX Interactions and Simulation Instructions

To create and investigate the interactions between a selected neutrino flux and a target, `snax` implements an *Interaction* object where a model flux and a specific target are coupled. This coupling is characterized by the specifed distance and target mass. For a given `snax` model and a specified target, `snax` can compute the differential interaction rates at each time and for each recoil energy and flavor, considering the distance to the supernova and the total target mass of the experiment, following the equations defined in Section 4.2.1. This allows tracking of the interactions on all interaction sites, including the time of interaction,

neutrino energy and flavor, the interacting isotope, and the resulting recoil energy. Having the interactions as a standalone object enables the study of different scenarios such as for different distances, different target masses, or different fiducialization volumes.

The distance scaling follows an inverse square relation to determine the expected number of interactions. For every ton of xenon, there are approximately 4.6×10^{27} xenon atoms. These parameters can be adjusted within the `snax` to explore various scenarios. This flexibility allows for the consideration of future next-generation DM observatories with larger volumes, such as DARWIN or XLZD [73, 74].

The expected number of interactions with each neutrino flavor, assuming a supernova at 10 kpc, is computed for all models considered, using `snax` and shown in Figure 4.11 for XENONnT–size experiment of 5.9 tons. The total expected interactions are also shown assuming perfect detection efficiency.

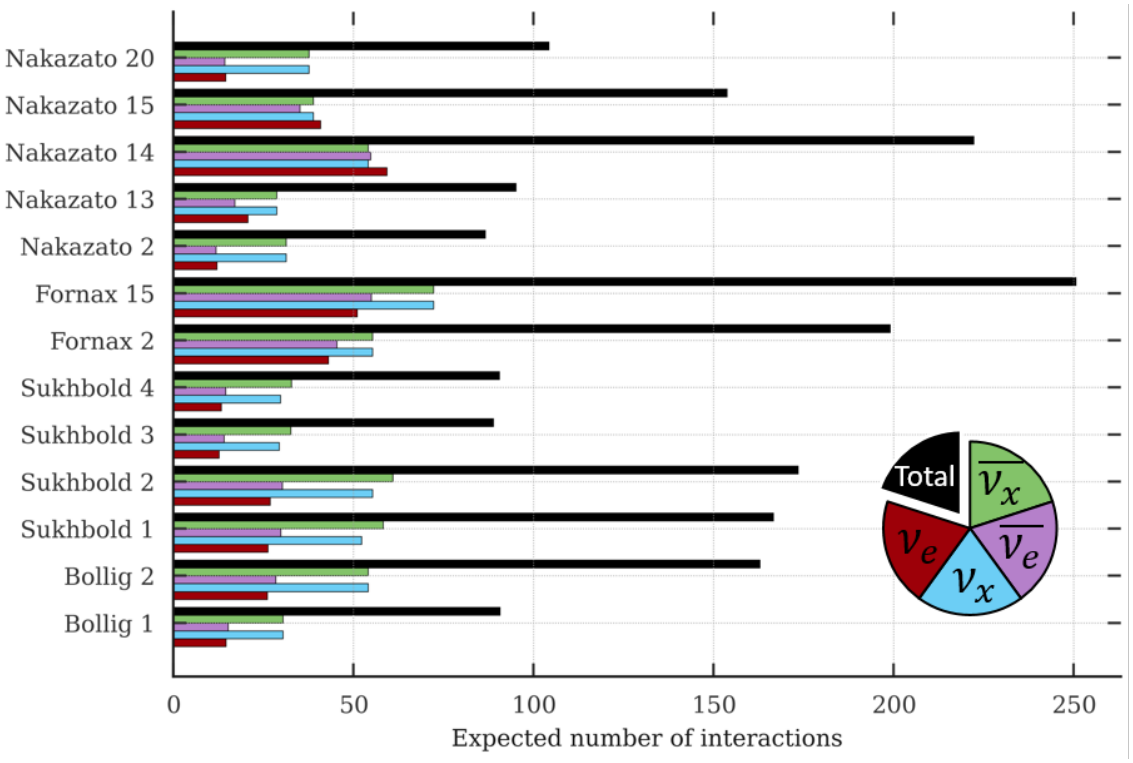

Figure 4.11: The total expected number of interactions for a 5.9 tons detector for all selected models, assuming the progenitor is located at 10 kpc. The expectations are shown for individual neutrino flavors as well as the total expected interactions. For simplicity, the `snax` index of each model is displayed on the y-axis next to the model name. These indicies can be referenced to obtain exact progenitor parameters in Table 4.1.

It is important to note that, while some models predict fewer total interactions, the energies and timing of these interactions are crucial. For example, models predicting a failed supernova, such as Nakazato models with `snax` id3 and id4, may forecast fewer interactions overall, but the time span for these neutrino interactions is relatively short. This results in higher interaction rates over a short period of time, which can make detection easier. The expected number of interactions as a function of the target mass and the distance to the supernova is shown in Figure 4.12 for the Bollig 2016 model with a 27 $M_\odot$ progenitor. The plot

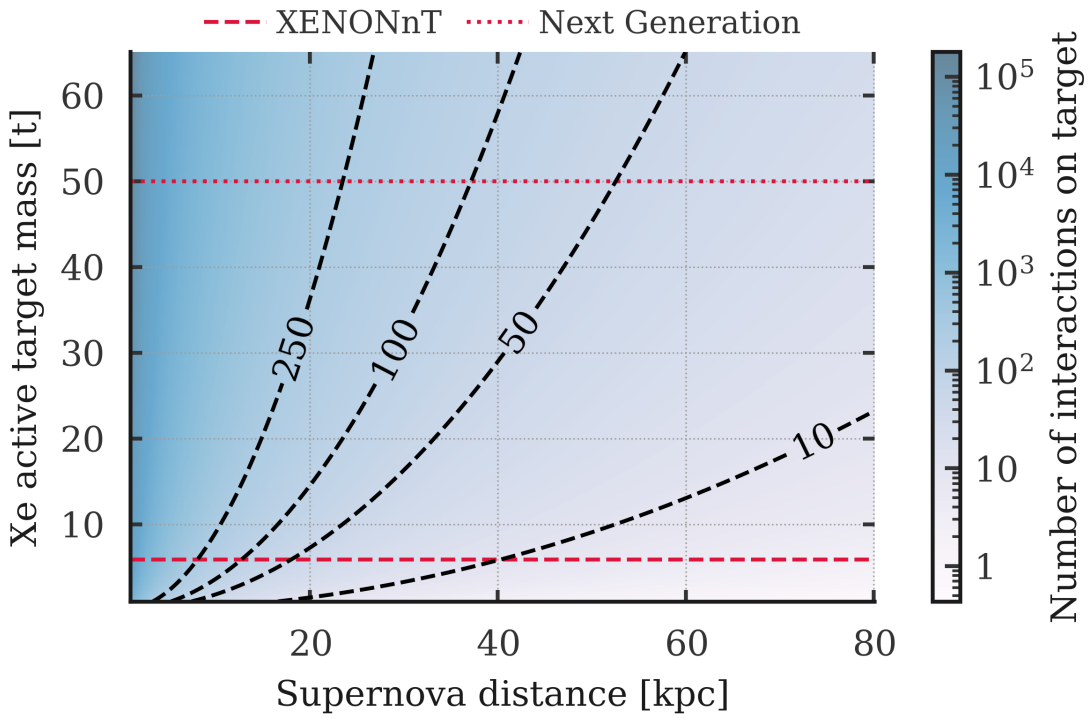

Figure 4.12: Expected interactions for the Bollig 2016 model with a 27 $M_\odot$ progenitor. The counts are scaled for distance according to the inverse square law for neutrino flux and by the target mass based on the number of xenon atoms. The masses for the current XENONnT experiment and a potential 50-ton next-generation DM experiment are indicated with red dashed and dotted lines.

includes the volume of the XENONnT experiment and a potential 50-ton volume for future next-generation DM experiments.

The interactions are computed in several steps. First, the total number of expected interactions at a given distance is calculated by integrating the individual neutrino flavor fluxes over all times and energies, considering only the active target mass of 5.9 tons of xenon.

Next, since the flux and mean energy of each neutrino flavor vary over time (as shown in the right panel of Figure 4.9), the interactions are calculated separately for each neutrino flavor. The interaction time is first sampled from the total flux distribution of a given flavor, which was integrated over all energies. Then, for each interaction, the neutrino energy is sampled based on the corresponding energy distribution for that flavor at the given time (see the left panel of Figure 4.9).

If no specific isotope is requested, an isotope is randomly sampled for each interaction based on its natural abundance. By retaining individual neutrino flavor information throughout the process, this method provides a more accurate recoil energy spectrum while preserving key distinctions in flux and energy distributions that could otherwise be lost in a flavor-blind CEνNS interaction. By performing this stepwise sampling and retaining the information for individual neutrino flavors, the expected recoil energy spectrum is ensured to be accurate, even though the final CEνNS interaction with xenon isotopes is flavor-blind.

The events from a supernova flux are expected to have uniform spatial distributions within the TPC, and `snax` provides functionality for uniform sampling within the XENONnT geometry.

Moreover, its flexibility allows for different sampling methods for position distribution, timing, and electric fields, which then determine the charge and light yields of NEST. The sampled time and recoil energy distributions can be used to create simulation instructions at the microphysics level, as described in Section 4.2.3. For convenience, `snax` also provides a method on the interaction object to generate simulation instructions and immediately call `fuse` to create simulations in a single line of code.

Furthermore, it is often desired to simulate more than one realization of the model at a given distance. Thus, `snax` offers a method to simulate multiple realizations of the same model by spacing the times between each supernova burst event by a specified interval to avoid overlap. This approach allows generating a single simulation file, which is computationally favorable, containing multiple supernova realizations to enhance statistical accuracy.

Typically, for each specific supernova model, 200 realizations were simulated within the XENONnT detector volume. In each simulation, the number of interactions was sampled based on the model's interaction rates, and the interactions were uniformly distributed in the detector volume. After running the simulation chain, the average number of generated S2 signals is shown in Figure 4.13, along with the expected number of interactions assuming 0 keV (perfect detection efficiency) and 0.7 keV (conservative cutoff step function).

As expected, the underlying energy distribution of the interactions plays a crucial role. For example, while the Nakazato model with a 50 $M_\odot$ progenitor is expected to have fewer interactions than the Bollig and both Sukhbold models with 27 $M_\odot$ progenitors with perfect detection efficiency, the observed counts with the XENONnT experiment are actually higher for the Nakazato model. This discrepancy highlights the impact of the detection efficiency and energy thresholds on the observed interaction rates.

The `snax` package also offers a function for the salting process detailed in Section 5.1.2. For this, the user specifies the simulated data together with the background data and chooses how many realizations to sample. Additionally, the user specifies a minimum separation time between different injections to avoid overlapping supernova signals. The salting function then identifies and samples the specified number of realizations from the hundreds produced for each model. For instance, the user can request to inject 50 supernova bursts into a 5-hour-long background run with a minimum separation of 2 minutes between each injection. The tool randomly selects 50 realizations from the simulated data file, with sampling allowing one to select the same realization multiple times. Then it identifies the appropriate time ranges in the background and selects injection points by sampling the requested number of times. To avoid complications from signals at the edges of a run, the tool does not place a supernova signal in the first or last minute of the background data. If the requested sampling is not feasible, such as attempting to inject 10 supernova realizations into a 1-hour run with 10-minute separations, the tool samples as many points as possible with the specified time separation. In this example, it would sample 5 time points, thus injecting only 5 realizations instead of the requested 10. Each sampled realization is then referenced to the corresponding timestamps, and the data iss sorted by time again. This ensures that each realization is salted on top of the background data.

The tools presented in this chapter, including `snewpy`, `fuse`, and `snax`, have been used to generate simulations at various distances with different supernova models. These simulations

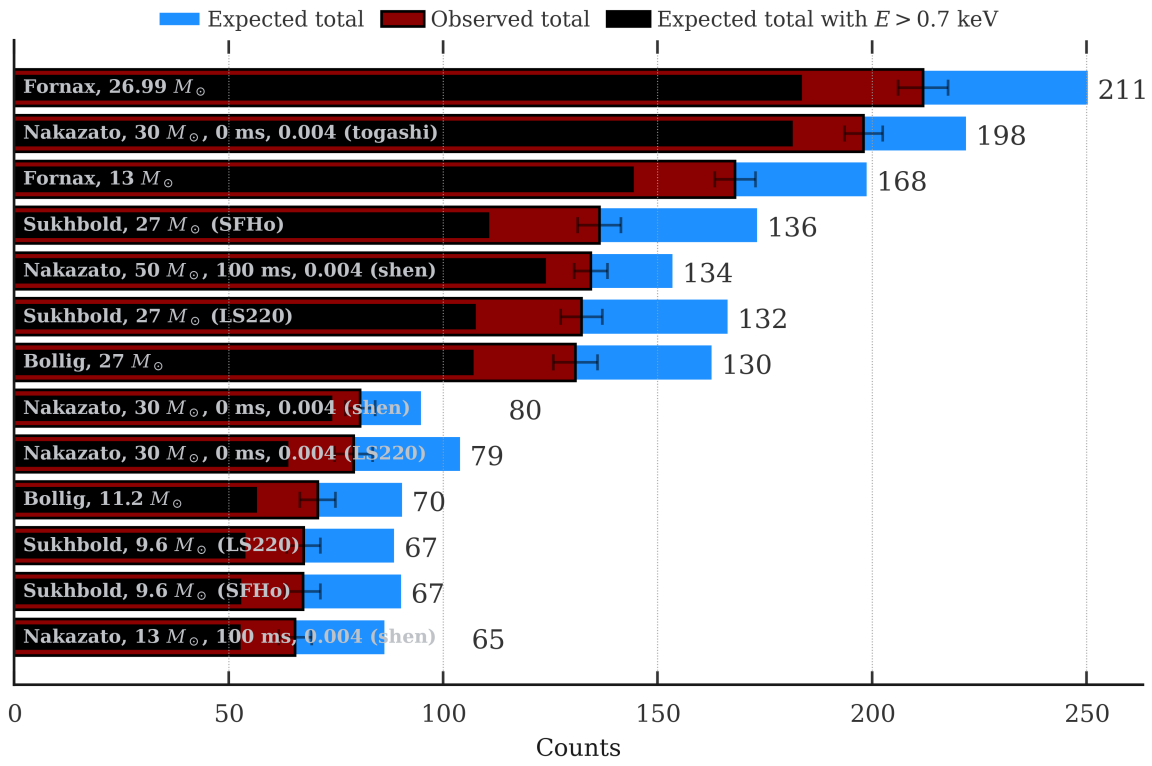

Figure 4.13: The total expected interactions are shown for both a 0 keV (perfect detection efficiency) and a more conservative 0.7 keV energy threshold, as referenced in Figure 4.4. The figure also displays the average number of detected S2 signals from the simulations representing each observed interaction, with the 1σ Gaussian standard deviation of the detected counts in 200 repeated simulations as error bars. As expected, due to the finite detection efficiency, not all interactions lead to a detectable signal. The observed counts from the simulations fall between the predicted counts for a perfectly efficient detector and those for a detector with an energy threshold as a step function at 0.7 keV. The average observed counts are also presented next to each bar.

provide the foundation for the analysis in subsequent chapters. In the next chapter, the focus shifts to examining the simulated signals and comparing them with the real detector background data. Subsequent sections explore the methods used to identify these signals within the background and evaluate the significance of the detections. Ultimately, an analysis framework for supernova detection and the sensitivity of the XENONnT experiment to various supernova models is established.

5

Probing the Cosmos: Exploring Supernova Detection with the XENONnT Experiment

"Without data, you're just another person with an opinion."
–W. Edwards Deming

Contents

In the previous chapter, we examined supernova neutrino emission models and simulated their interactions in the XENONnT detector, establishing an expected signal framework. In this chapter, we focus on developing analysis strategies to effectively identify supernova neutrino interactions within the XENONnT data, covering techniques for signal selection, background suppression, and prompt supernova identification. We also discuss methods to assess the significance of the observed signals. In Section 5.1 the background model and the salting method are introduced, followed by the analysis of the *salted data* in high–level signals. In Section 5.2 signal identification methods are discussed and finally in Section 5.3 expected signal significances are presented.

This analysis primarily focuses on promptly detecting supernovae rather than precisely identifying each interaction. However, selecting interactions and reconstructing the supernova light curve are crucial and will be addressed later, requiring careful investigation of each signal.

5.1 Analyzing TPC Data

In this section, we discuss methods for analyzing interactions within the XENONnT detector, building on the supernova signal simulations introduced in Chapter 4. To effectively capture supernova signals within the regular XENONnT data stream, we begin by investigating the background model in Section 5.1.1. Simulated supernova signals are then added to the sampled real background data following the procedure in Section 5.1.2 to allow studying the characteristic signal signatures under realistic conditions.

This section presents the analysis of the high–level data structures of XENONnT: the peak (Section 5.1.3) and the event–level data (Section 5.1.4). Peak–level data includes detailed signal features and are used for data selection using both traditional methods and machine learning approaches. This dual–level analysis approach combines the benefits of each data type, offering a comprehensive view of supernova signals.

5.1.1 Background Model

The sensitivity and detection capabilities of the XENONnT experiment for low–energy nuclear recoil interactions were discussed in Chapter 2. XENONnT, being an ultrapure and highly sensitive detector designed to search for low–energy recoils from WIMPs, is also well suited for detecting CEνNS interactions from supernovae, as these events are expected to result in nuclear recoils within the same energy range.

The background in XENONnT within the same supernova energy region is dominated by single–electron signals arising from delayed electron emissions [107, 106]. These delayed electrons are caused by photoionization of impurities caused by S1 and S2 light, producing signals with S2 areas similar to those expected from supernova interactions. Understanding

these background components is essential for mitigating their effects on the supernova detection.

As mentioned in Section 2.2.1, increased photoionization activity in SR1 caused changes in signal rates. To illustrate the changes in background S1 and S2 signal rates, two hours of data were randomly selected from both SR0 and SR1 runs, and the rates are shown in Figure 5.1. Note that in this figure, no data selection was applied; therefore, the raw signal counts were high. For the purpose of supernova neutrino search, this chapter primarily addresses the most recent detector conditions, i.e. SR1. The analysis was optimized to reflect these conditions while also monitoring the behavior under SR0 conditions.

Similarly, the evolution of how the S1 and S2 signal rate would look like during a supernova burst is shown in Figure 5.2, where the rates are calculated using a rolling window of 1-second and smoothed for illustration. The figure shows the average rate of the model with a 1σ deviation around the mean using 100 simulations. The analysis aims to develop signal selection criteria that optimize the signal-to-background ratio.

5.1.2 The Salting Method

The evolution of the signal rate is the primary indicator of a supernova burst, as the background is expected to be stable within a selected range in the S1 and S2 signal region, while in a supernova burst the signal rate is expected to exhibit a sharp increase (e.g. see Figure 5.2). The optimal way to investigate rates over time is by examining the rate evolution in a rolling time window.[i] To better understand the impact of the size of the rolling window on the examination of the rate evolution, simulated data was used and the rates in different rolling windows were calculated using the width 0.5, 1 and 3 seconds. This is shown in Figure 5.3 where the rate evolution of the signal is shown for three different sizes, and the upper panel displays the individual time stamps as simple lines. The rates were computed by opening a backward window at the time of a given peak, counting the number of peaks within the window, and then moving to the next peak. Consequently, the rates start at zero and build up to a maximum at the peak of the supernova burst. The size of the chosen rolling window is crucial; a window longer than the duration of the burst would smooth out the features and obscure the data, whereas a window too short would not represent the shape of the actual supernova flux.

The background rate is expected to be stable throughout the collection of data within the selected region. Therefore, the choice of rolling window size should not affect background rates. For comparison, the rate evolution within different rolling window sizes for a minute–long background data from SR1 is shown in Figure 5.4.

To analyze supernova signals under realistic conditions, simulated supernova signals are superimposed onto existing background data using the `snax` tool in a process called *salting*. This approach, detailed in Chapter 4, allows for a data-driven study of signal relationships

[i] Binning the data is inefficient in such studies, as the bin size and placement can significantly alter the results, potentially causing the signal to be missed. This was also discussed in [147].

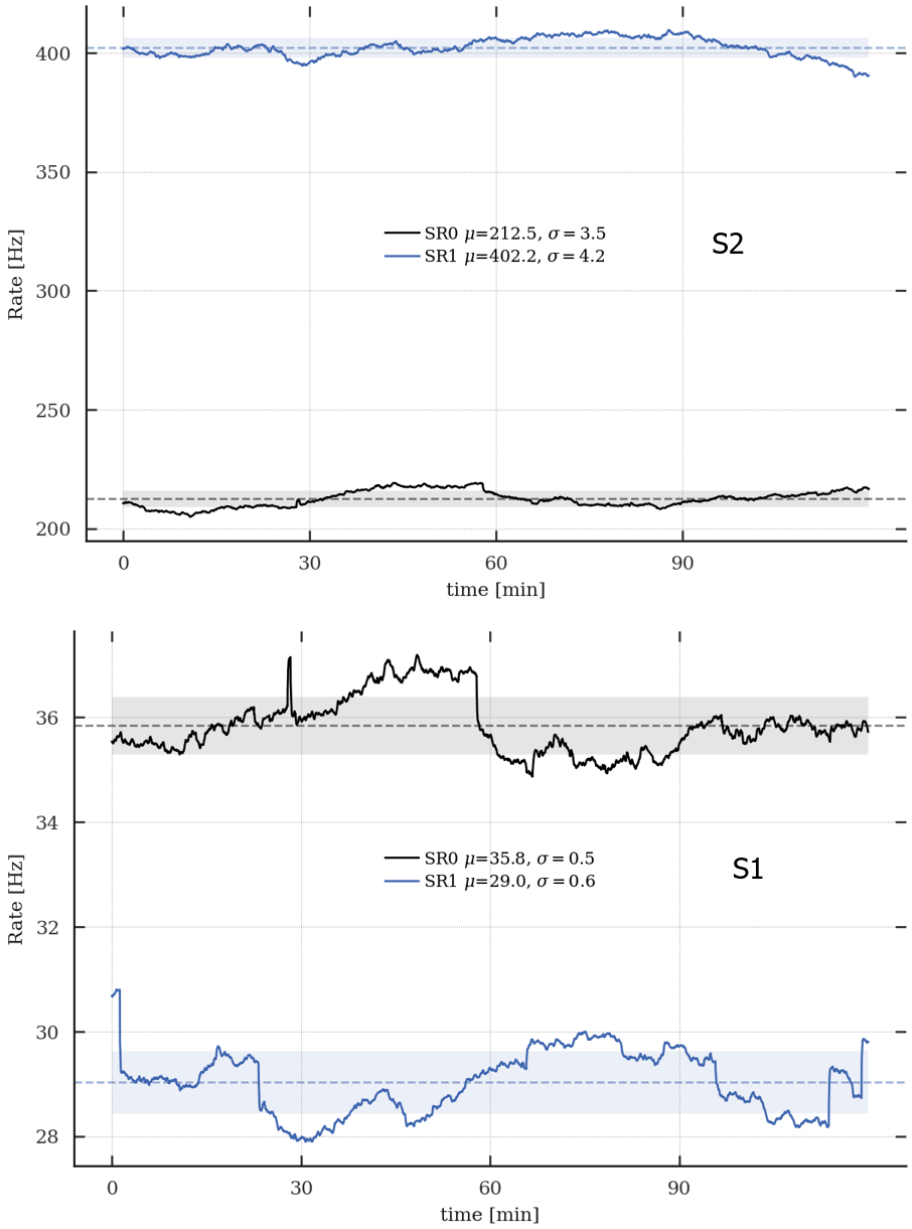

Figure 5.1: S1 and S2 signal rate evolutions over randomly sampled 2 hours long real background data in SR0 and SR1. The S2 signal rates are shown at the top figure while the S1 signal rates are shown at the bottom figure. The rates were computed in 10 second rolling windows and smoothed with a 30–minute averaging for better illustration. Also shown are the average signal rates and 1σ deviation around the mean assuming a Gaussian distribution. No data selection was applied.

and the detection sensitivity of XENONnT. The salted data, incorporating simulated S1 and S2 signals at the peak level, provide a realistic framework for supernova burst analysis.

Figure 5.5 illustrates the salting process, where the black data points represent the background around a given mean rate, and the red points indicate the supernova signals injected into the dataset. The figure shows how the superposed data would look in the data stream. In the salted data, each signal is marked to indicate which ones were part of the real background and which were originating from the injected supernova simulations.

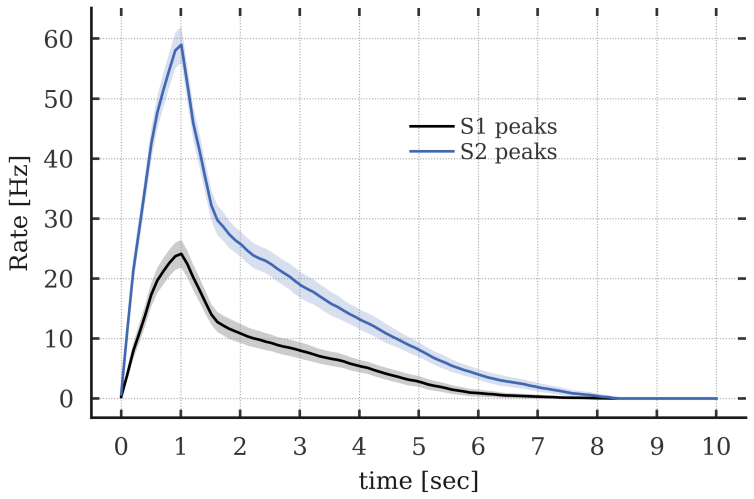

Figure 5.2: S1 and S2 rate evolution in the Bollig 2016 model with a 27 $M_{\odot}$ progenitor. The bands show 1σ deviation around the mean, extracted from the slread of 100 simulated supernova bursts.

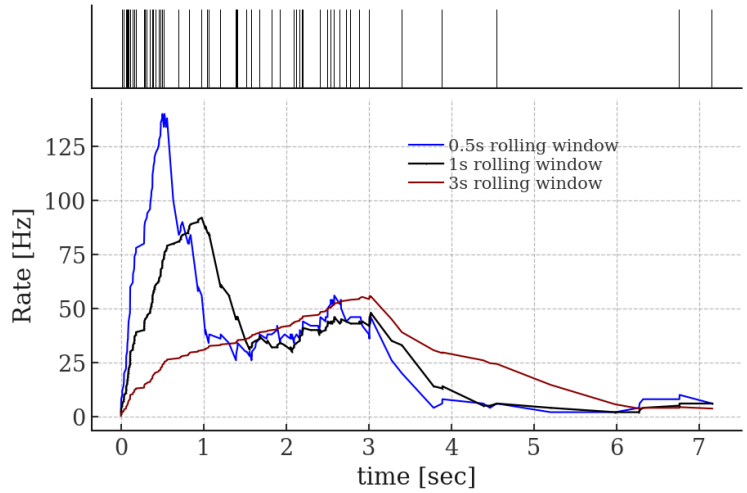

Figure 5.3: Peak rate evolution of a simulated signal. The upper panel shows the exact registered time of each peak, and the lower panel shows the peak rate evolution using different window sizes. For comparison, the rates are shown in 0.5, 1, and 3 second rolling windows. Although the total number of peaks remain the same, the shape of the signal changes based on the selected window size. For small window sizes the signal shape cannot be accurately constructed and for larger window sizes the features are washed out.

Following the salting procedure, 26 background runs that span in total about 90 hours were sampled from the SR1 data and salted with a total of 2194 individual supernova bursts across 13 different progenitors from 4 models[ii]. The process flow chart is also shown in Figure 5.6. The details of the runs and the injected supernova progenitors are listed in Table 5.1. The number of injections was automatically determined by the `snax` tool, taking into account the duration of the sampled runs and the predefined minimum interval between each supernova injection. Notice that each model was used to salt 2 randomly sampled background runs. The initial set of salted data was prepared with a minimum separation of 3 minutes between

[ii]To investigate behavior under SR0 conditions, five runs with run id 025423, 025424, 025425, 025426, and 025427 were also sampled. These runs span approximately 4 hours, from 2021-07-13 16:58 to 2021-07-13 21:04, and were salted using the Bollig 2016 benchmark model.

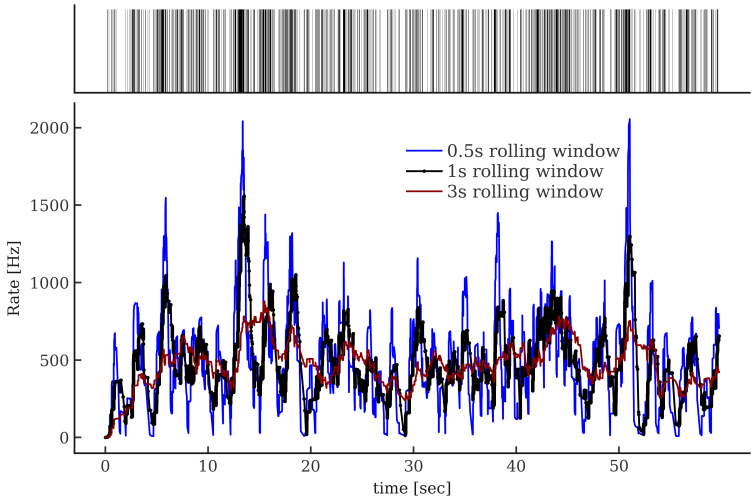

Figure 5.4: The peak rate evolution in rolling windows of 0.5−, 1−, and 3− seconds, in SR1. Notice here no data selection is applied and the rates are shown for all peak types.

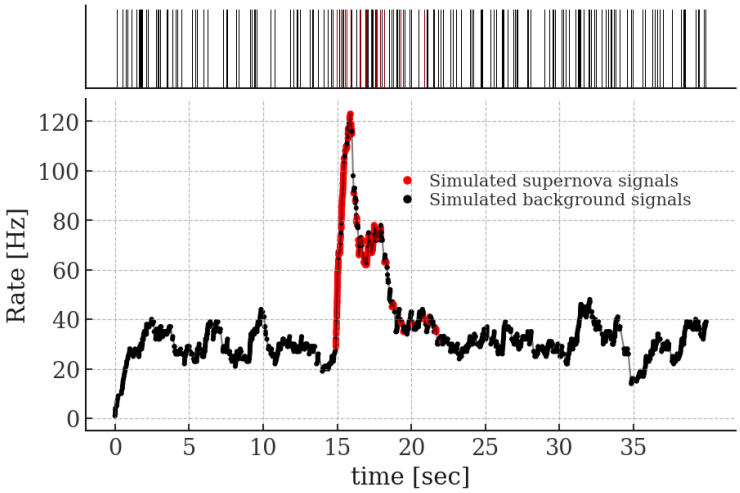

Figure 5.5: Illustration of the salting procedure. The upper panel shows the time stamps of each signal as black lines while the lower panel shows the rate evolution of these signals. For this illustration, the background (black) was also simulated by sampling times around a selected mean rate, while the supernova signal was simulated from the model and shown in red. The figure highlights the superposition of the two signals.

injections, whereas the second set was prepared with a 2–minute separation to enhance statistics.

After the salting procedure, additional features were derived from the salted data to enhance the analysis. One such feature was proximity information, which includes the number of nearby signals occurring both before and after a given signal, as well as the time intervals between them. Another enhancement involves *subtyping*, which extends beyond conventional S1 and S2 classifications by incorporating temporal and spatial relationships, along with the area of the signal. This method identifies specific interaction types, such as delayed electron signals, lone S2 signals, and photoionization signals, enabling targeted data selection. By focusing on subtypes associated with physical interactions, the analysis can more effectively isolate relevant signals while minimizing background events.

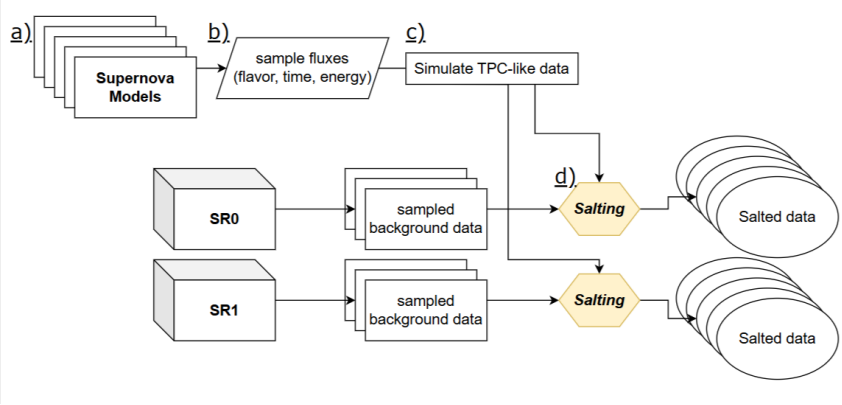

Figure 5.6: Steps to prepare salted data for analysis: *a)* Supernova models are collected from `snewpy`. *b)* Interactions are computed, and samples for multiple supernova bursts are created using `snax`. *c)* Realistic TPC simulations for each supernova realization are generated with `fuse`. *d)* Sampled background data and simulated supernova data are combined (salting) using `snax`.

Run ID	Run duration	Nr. of injections	Model name	`snax id`	Progenitor mass
053243	3.00 hr	60	Bollig 2016	1	11.2 $M_\odot$
050048	0.77 hr	23	Bollig 2016	1	11.2 $M_\odot$
052156	5.00 hr	100	Bollig 2016	2	27.0 $M_\odot$
046920	2.00 hr	60	Bollig 2016	2	27.0 $M_\odot$
053493	3.00 hr	90	Fornax 2021	2	13.0 $M_\odot$
052376	5.00 hr	100	Fornax 2021	2	13.0 $M_\odot$
053256	3.00 hr	60	Fornax 2021	15	26.99 $M_\odot$
050074	6.00 hr	180	Fornax 2021	15	26.99 $M_\odot$
053125	5.00 hr	100	Nakazato 2013	2	13.0 $M_\odot$
047661	3.00 hr	90	Nakazato 2013	2	13.0 $M_\odot$
053252	3.00 hr	60	Nakazato 2013	13	30.0 $M_\odot$
047635	3.00 hr	90	Nakazato 2013	13	30.0 $M_\odot$
053258	3.00 hr	60	Nakazato 2013	14	30.0 $M_\odot$
047344	3.00 hr	90	Nakazato 2013	14	30.0 $M_\odot$
047634	3.00 hr	60	Nakazato 2013	15	30.0 $M_\odot$
053246	3.00 hr	90	Nakazato 2013	15	30.0 $M_\odot$
052963	5.00 hr	150	Nakazato 2013	20	50.0 $M_\odot$
047574	3.00 hr	60	Nakazato 2013	20	50.0 $M_\odot$
047636	3.00 hr	90	Nakazato 2013	20	50.0 $M_\odot$
052921	5.00 hr	100	Sukhbold 2015	3	9.6 $M_\odot$
048470	2.73 hr	81	Sukhbold 2015	3	9.6 $M_\odot$
051818	5.00 hr	100	Sukhbold 2015	2	27.0 $M_\odot$
047630	3.00 hr	90	Sukhbold 2015	2	27.0 $M_\odot$
049327	6.00 hr	120	Sukhbold 2015	1	27.0 $M_\odot$
046765	2.00 hr	60	Sukhbold 2015	1	27.0 $M_\odot$
047572	3.00 hr	60	Sukhbold 2015	4	9.6 $M_\odot$
046934	2.00 hr	60	Sukhbold 2015	4	9.6 $M_\odot$

Table 5.1: Details of the salted SR1 runs, including the run ID, run duration, number of injected supernova bursts, supernova model, Snax index, and progenitor mass. For the complete parameters of the progenitor, refer to Table 4.1.

5.1.3 Analysis of the Peak–Level Signals

Physical interactions in the TPC typically generate S1 and S2 signals, if either of them is missing, a higher level *event* data can not be built. Furthermore, the construction of event–level data requires constituent peaks to have a minimum 2–fold PMT coincidence for S1 signals, and a minimum 100 PE area for the S2 signal ($A(S2) > 100$ PE). Therefore, not all interactions generate an event. However, low–energy interactions will often still generate a peak–level signal with only S2 conuterpart.

The primary distinction between the CEνNS signal from a supernova and the background is that a supernova represents a burst event, resulting in a considerable number of interactions in a short time and a specific region of interest defined by the main signal properties. In contrast, background signals in the same region are mostly due to delayed electrons, which can often be identified.

To mitigate these background signals and highlight the embedded supernova signal, salted datasets outlined in Section 5.1.2 were used. It is important to note that the data filtering procedures described in the following sections were not optimized to select supernova–like signals with the highest purity but aim to build a generic framework capable of identifying any supernova burst while yielding the lowest possible false supernova alarm rates. This false alarm rate (FAR) is influenced by both the signal selection and the burst identification (also referred to as *trigger*) mechanism discussed in the next section.

The most commonly used data cleaning techniques based on peaks are detailed in Section 5.1.3.1. Although these methods provide valuable information, advanced machine learning techniques, explained in Section 5.1.3.2, were ultimately used to categorize each signal into supernova origin or background noise. The trigger algorithm then processes the signals that were classified as supernova origins to detect supernova signatures. The signal selection threshold and trigger parameters were globally tuned using all salted data, as elaborated in Section 5.2.

5.1.3.1 Standard Methods for Data Filtering

In general, to distinguish signals from supernova interactions amid background events, various features of the data were analyzed at different processing levels, each offering unique insights. Event–level data, constructed using both S1 and S2 signals, provides information to identify physical interactions. However, for low–energy recoils, such as those expected from low–mass WIMP interactions or CEνNS from astrophysical sources, event–level reconstruction is often challenging. Traditionally, searches at these energy scales rely on peak–level data, such as S2-only signal analysis [46, 106]. However, incorporating S1 signals alongside S2 can offer an additional layer of information, particularly for detecting sharp rate increases indicative of supernova bursts. Using both types of signals, detection capabilities can be significantly improved. Consequently, as a foundational step, the properties of all signal types were

examined to identify features that effectively discriminate between supernova signals and background noise, establishing a robust framework for data filtering.

Analyzing the characteristics of the signal in the salted data indicates that background noise can be significantly reduced by selecting signals based on signal area and the relation to surrounding peaks. The signal filtering procedure should not be optimized to fit a particular model, as the main objective is the rapid identification of any supernova model, which may exhibit diverse characteristics. It is important to recognize that the salted data assume a supernova signal at 10 kpc; however, more distant supernovae will produce fewer interactions, leading to potentially sparser detections over time. Over–tuning parameters, such as the time to the next peak, could be counterproductive under these conditions. To avoid this, a minimalistic approach to signal selection is adopted.

The distinct populations identified through subtyping are shown in Figure 5.7, illustrating both background and simulated supernova signal distributions for SR0 (upper left) and SR1 (lower left) prior to any data selection. Initially, only signals with specific subtypes were considered for analysis, focusing on those most relevant to expected neutrino interactions. The subtypes included in the analysis are the following:

- *DE*: Delayed electron signals.
- *lonePS2*: Lone primary S2 signals.
- *PS2*: Primary S2 signals.
- *S1*: Primary S1 signals.

Other subtypes were excluded because of their irrelevance to neutrino interactions, rarity, or association with detector effects. A summary of excluded subtypes and their reasons for exclusion is presented in Table 5.2.

Table 5.2: Excluded signal subtypes and reasons for exclusion.

Subtype	**Reason for Exclusion**
isoDE, *isoS1*, *isoPS2*	Isolated signals unlikely to result from neutrino interactions.
loneS1, *sloneS1*, *junk*	Misclassified or background signals unrelated to physical interactions.
S1PH, *S2PH*, *slS1PH*, *fakeS2PH*	Signals related to photoionization, likely caused by detector impurities.
fakeS2olS2, *S1olS2*, *S2olS2*	Rare subtypes involving secondary signals near larger S2 signals, infrequently associated with neutrino interactions.

This approach ensures that the analysis focuses on the most relevant signal subtypes while systematically excluding those that are either irrelevant or unlikely to correspond to genuine neutrino interactions. Notice that low–energy nuclear recoil interactions primarily generate single- and double–electron signals, but this region also overlaps with delayed electron signals. Consequently, a selection criterion was employed to reduce background events by eliminating very-low area signals, despite the fact that this results in a considerable loss

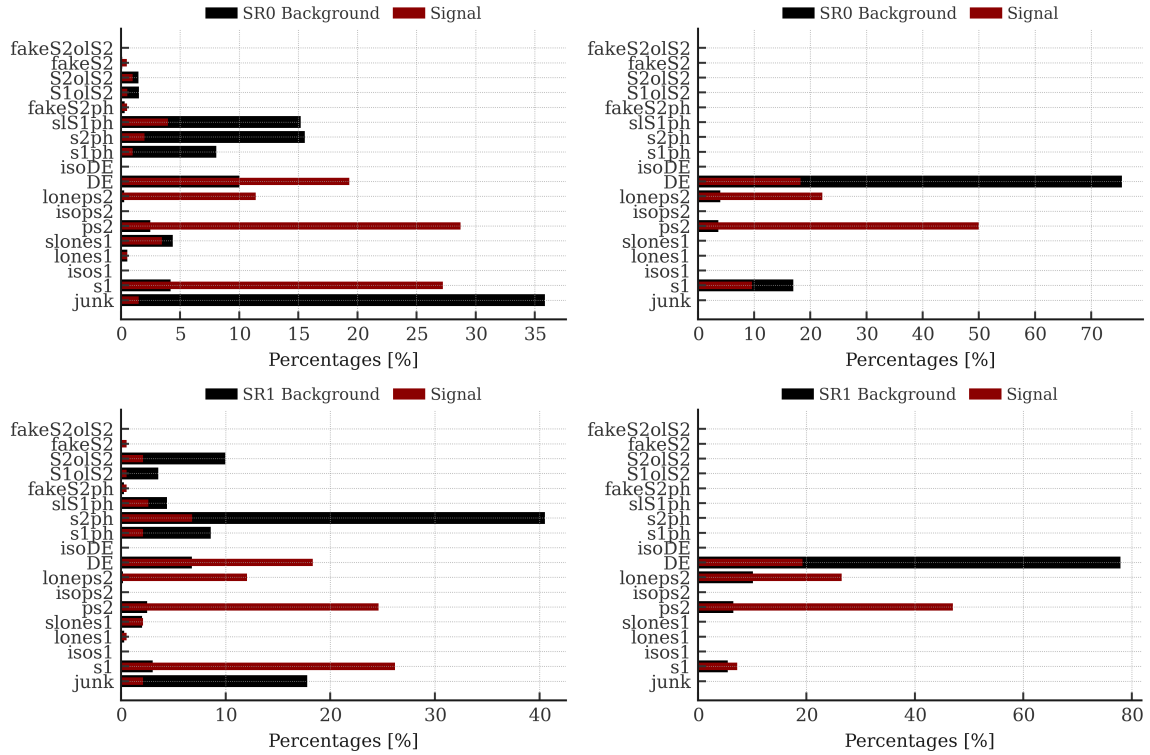

Figure 5.7: Population distribution of different subtypes before and after applying the traditional cuts. The upper panels show the different populations in SR0: before cuts (left) and after cuts (right). The lower panels show the distributions in SR1: before cuts (left) and after cuts (right). For each figure, respective data salted with the same benchmark Bollig model is shown.

of signal. Nevertheless, neutrino interactions will continue to produce some higher–energy signals that remain useful for identification purposes.

Subsequent signal selections were based on various parameters, including signal area, width[iii], time to the preceding signal (*t_to_prev_peak*), number of competing signal counts (*n_competing*), and x–y positions (through fidicuialization). Figure 5.8 illustrates the distributions of the number of competing counts and the time to the preceding signal. These distributions are shown for both background signals (black) and simulated supernova signals (red) in the salted SR1 dataset. The figure highlights the distinct characteristics of signals with background and supernova origins, providing a basis for data–selection strategies using these parameters.

Similarly, Figure 5.9 presents two–dimensional distributions of the signal area versus the signal width in the salted SR1 data set. The upper panels depict the distributions of background signals, whereas the lower panels show those of simulated supernova signals. These figures highlight the distinct signal distributions, showing how area and width parameters can differentiate background from supernova signals.

Data selections were designed to maximize the signal-to-background ratio within the chosen parameter space while avoiding over-optimization specific to the benchmark model. By maintaining generality, the criteria are kept sufficiently broad to ensure sensitivity to other

[iii]The width in the context of XENONnT analysis is defined as the range where the are of the signal reaches 50% (*range_50p_area*) or 90% (*range_90p_area*) of its total area.

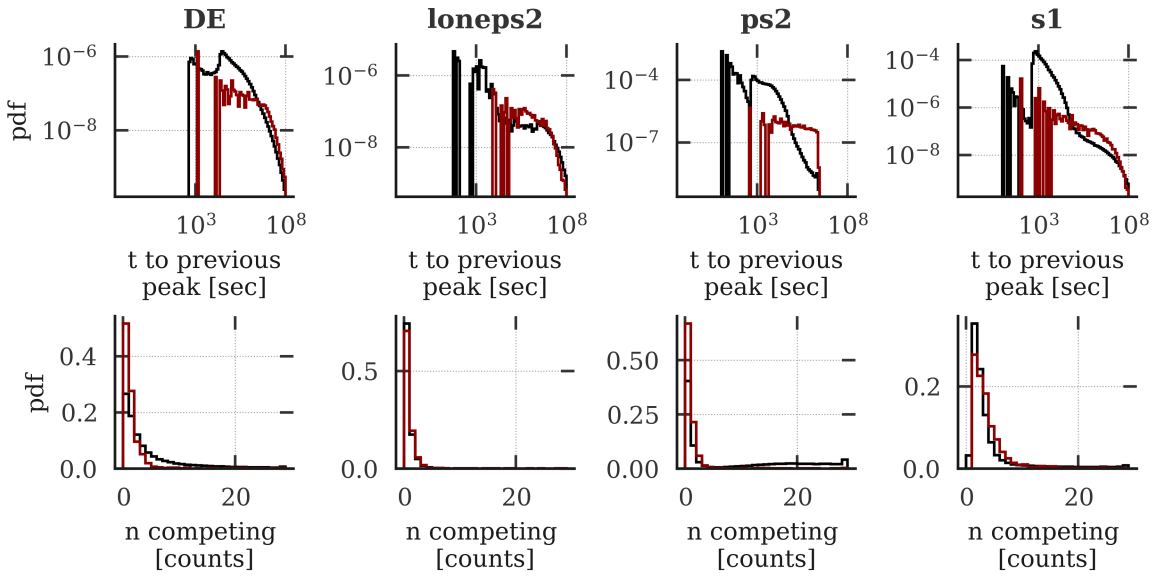

Figure 5.8: Distributions of the number of competing signals and the time to the preceding signal for selected subtypes, prior to applying any data selection. The black histograms represent background signals, and red histograms represent simulated supernova signals.

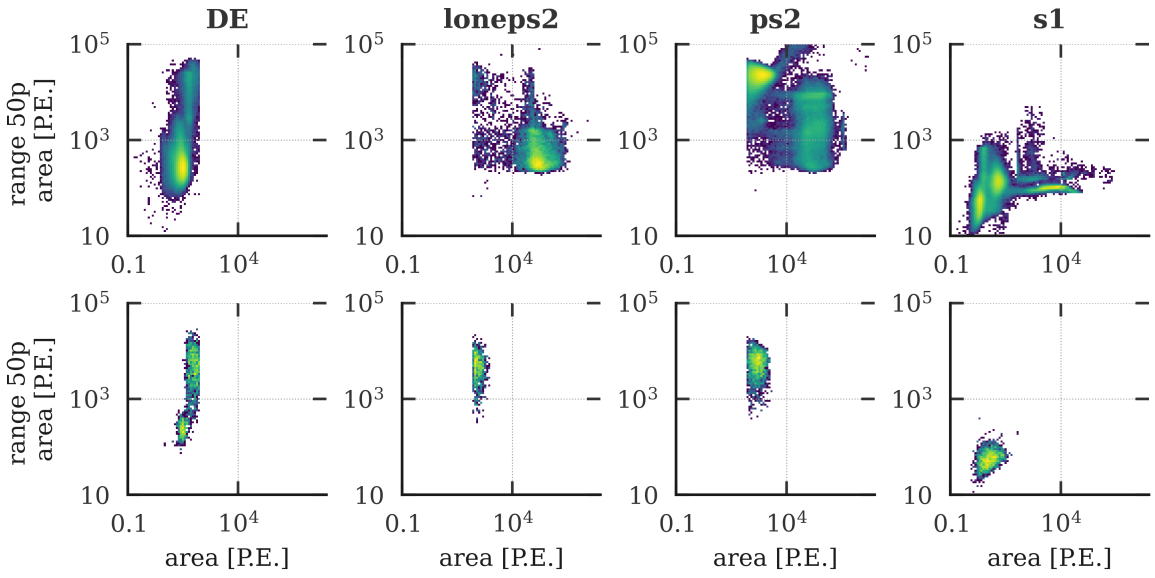

Figure 5.9: Two-dimensional distributions of signal area versus width for selected subtypes. The top panels show the distribution of background signals, while the bottom panels show the distribution of simulated supernova signals.

supernova models. Separate tuning was performed for SR0 and SR1 to account for their distinct detector conditions. The final selection criteria are summarized in Table 5.3, while Table 5.4 presents survival probabilities, indicating the percentage of signals retained for each selection step.

Before applying the signal selections, the raw signal rates for both the SR0 and SR1 data, salted with the benchmark model, are shown in the upper panels of Figure 5.10 with blue lines. After selecting specific subtypes and applying the signal selection, the final rates are shown in the lower panel of each figure with red lines. Here, the rates were computed within

Table 5.3: Final signal selections for SR0 and SR1. The described parameters are; TPC radius, the width of the signal, number of competing signals, time to previous signal, and the signal area.

Subtype	Parameter	SR0 Selection	SR1 Selection
All	Fiducial radius	≤ 58 cm	≤ 58 cm
s1	range_50p_area	≤ 85 ns	-
s1	n_competing	-	≤ 10
pS2	t_to_prev_peak	$\geq 7 \times 10^4$ ns	$\geq 7 \times 10^4$ ns
DE	area	≥ 50 PE	≥ 40 PE
DE	n_competing	≤ 4	≤ 4
DE	range_50p_area	-	≥ 800 ns

Table 5.4: Final signal acceptances for SR0 and SR1. Individual signal acceptances are calculated using $N-1$ method.

Subtype	Parameter	Signal (SR0)	BG (SR0)	Signal (SR1)	BG (SR1)
All	FV radius	71%	6%	72%	81%
s1	range_50p_area	99%	29%	-	-
s1	n_competing	-	-	99%	15%
pS2	t_to_prev_peak	97%	56%	97%	75%
DE	area	89%	3%	99%	71%
DE	n_competing	98%	54%	98%	81%
DE	range_50p_area	-	-	96%	29%
Total	-	48%	0.05%	47%	0.15%

1–second rolling windows.

Notably, the raw signal rates are quite high, primarily due to background events near the walls of the TPC, which are heavily reduced with the fiducial volume selection. After applying signal selection, the average rates in both salted datasets reduce to 8.2 Hz and 11.1 Hz for SR0 and SR1, respectively. The average rates, along with the Gaussian standard deviation $1\,\sigma$, are also shown in Figure 5.10. It is important to note that while the average rates are provided, the rate increases at the injection points are quite clear and usually exceed 40 Hz, hinting at a high detection significance for supernova signals.

Traditional data selection based on signal characteristics offers an intuitive way to distinguish the supernova signal from the background within the region of interest (ROI). This approach typically involves one–dimensional filtering, where a value is selected based on whether it falls above or below a specific threshold. Employing subtyping and applying different selections for various subtypes makes the selection process multidimensional. By tailoring selections to each subtype, it is possible to differentiate between supernova signal and background data more accurately. Conditional selection, where a selection on one feature is applied only if another feature meets certain criteria, further enhances this discrimination. For example, a cut in peak area might be applied only if the peak width falls within a specific range, leveraging the relationships between different features, but this can potentially introduce correlations in the data that should be carefully studied.

Improved signal and background discrimination can be achieved by employing machine learning methods such as decision trees, which will be explored in the next section. Decision

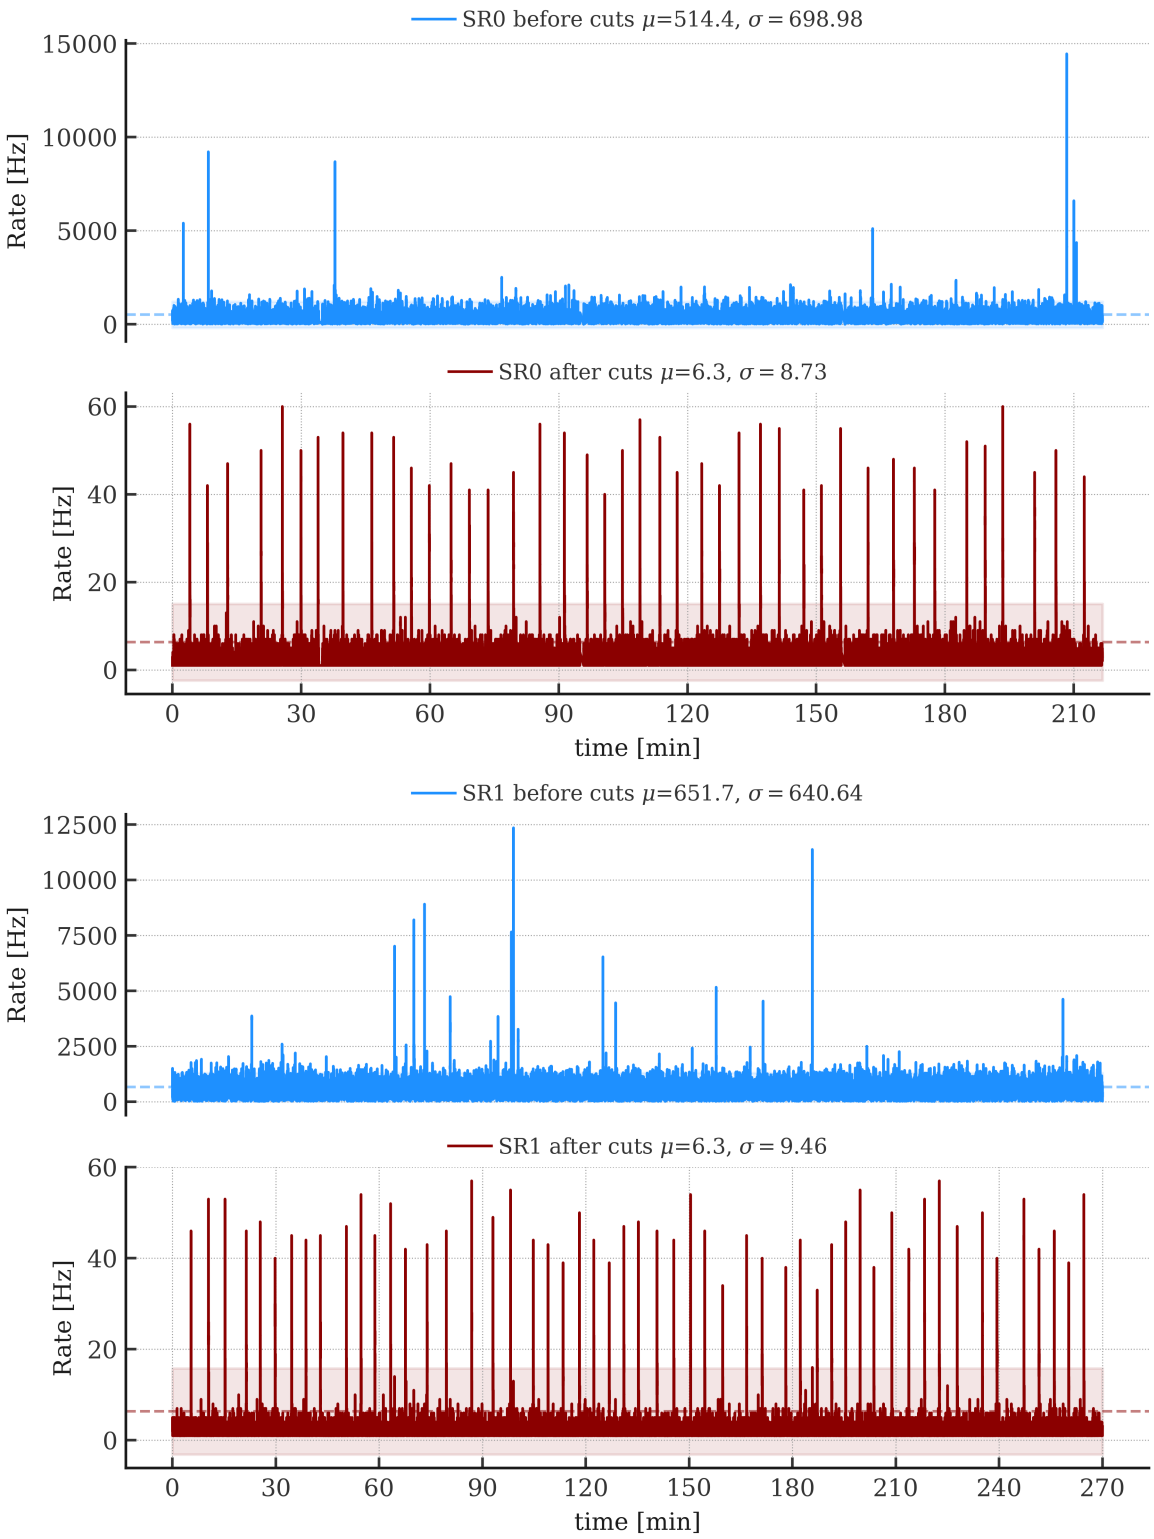

Figure 5.10: The total peak rates in SR0 (*upper*) and SR1 (*lower*) before and after applying the cuts based on traditional analysis. Notice in both cases the rates are heavily reduced and background is significantly mitigated. In the filtered data, each vertical rate increase corresponds to an injected supernova signal which was not visible before applying the data selection.

trees handle complex interactions between multiple features simultaneously, offering greater accuracy in identifying supernova signals. They build on the concept of conditional selections

by creating a hierarchical structure, where each decision node applies a selection based on one feature, and subsequent nodes refine the selection based on additional features. This approach constructs a model with high discrimination power that efficiently separates the signal from the background noise.

5.1.3.2 Data Selection with Machine Learning Techniques

Looking for efficient ways to leverage signal parameters, we note that while some features have known correlations, such as signal area and extracted electrons, others, such as delayed electron correlations with large S2 signals, are not as straightforward. Traditional methods attempt to decode these relationships, but machine learning models, particularly decision trees, offer deeper insights by classifying data from simulations and identifying complex patterns that traditional approaches might miss.

A decision tree splits data into smaller subsets (or leaves), based on parameter thresholds, continuing this process until predefined conditions, such as a minimum number of samples per leaf or a desired accuracy, are met. For this analysis, a binary classification model based on decision trees, XGBoost[iv] (eXtreme Gradient Boosting), was employed. Trained on the salted SR0 and SR1 data described in Section 5.1.2, XGBoost constructs an ensemble of decision trees sequentially, where each tree aims to correct the errors of the previous one. This iterative process minimizes the loss function and improves the overall predictive accuracy.

The salted data were organized in a tabular and *sequential* format, where each row represents a peak and includes various features such as area, width, and proximity to adjacent peaks. Temporally relevant information, such as the number of nearby signals, the temporal distances to the previous and the next signal, and the rates, was incorporated in each row to allow individual signal characterization and to facilitate *random sampling* for analysis. To avoid potential bias, certain details, such as the exact time of interaction, interaction ID, or labels added during the salting procedure, indicating the origin of the signal, were intentionally excluded. Although subtype information also contributes to classification, it was derived from other parameters that the XGB model, in principle, can independently infer. Moreover, subtype information is not part of the primary data processing pipeline in the XENONnT experiment and must be computed separately, a process that is computationally intensive. For this study, both the SR0 and SR1 datasets were analyzed with and without subtype information. It is worth noting that subtypes can still be calculated post-detection to improve the selection purity of the identified signals.

The most significant features, determined by their contribution to the total discrimination power of the model, are shown in Figure 5.11 for SR0 (top) and SR1 (bottom). In this figure, it can be seen that proximity information and subtypes (when applicable) are some of the most important features.

[iv] In this study, XGBoost and XGB are used interchangeably.

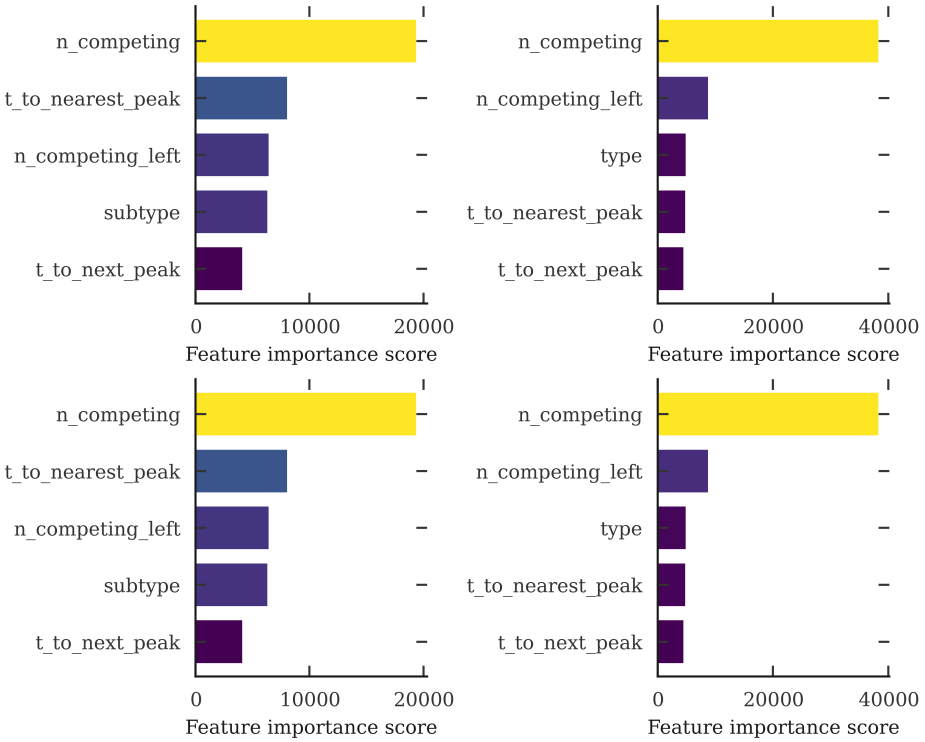

Figure 5.11: The feature importance scores from the XGB models trained on different datasets. The upper panel shows the salted SR0 data with (without) subtype information on the left (right), while the bottom panel shows the salted SR1 data with (without) subtype information on the left (right). Features like the number of competing signals (*n_competing*), time to the nearest peaks (*t_to_nearest_peak*), and signal types (*type/subtype*) are often among the most important features. The feature importances are computed based on their impact on accurately classifying a signal to its true origin. It can be seen that the proximity information is always important, and similarly, the subtype information is always in the top five most important features, when available.

Decision trees were used to classify the peaks into two labels: supernova or background signals. For SR1 data without subtype information, the XGB model was configured to perform binary classification. The model was designed to capture complex interactions within the data by setting the learning rate to 0.1 and limiting the tree depth to 12. To reduce the risk of overfitting, only 80% of the samples and features were used in each iteration of the boosting process. Additionally, regularization techniques were employed to constrain the model's complexity. A subset of the background data was randomly sampled to ensure a balanced dataset between background and supernova signals. The model reached its optimal performance after 71 iterations. The performance of the model, assessed using a logarithmic loss function, is illustrated in Figure 5.12.

The model assigns an XGB prediction score to each peak, ranging from 0 to 1, indicating the likelihood that the peak originates from a supernova (1) or a background (0). The distribution of the predicted labels is shown in Figure 5.13. This figure shows the labels in the salted SR1 data along with the distributions of the true underlying classes. Note that the y–axis is on a logarithmic scale. A selection was defined based on these predicted labels. A dashed vertical line at 0.9 is shown to indicate a conservative selection in which the peaks with an XGB prediction of 0.9 or greater were kept. Allowing for lower confidence predictions increases coverage and the true positive rate in the final data; however, this

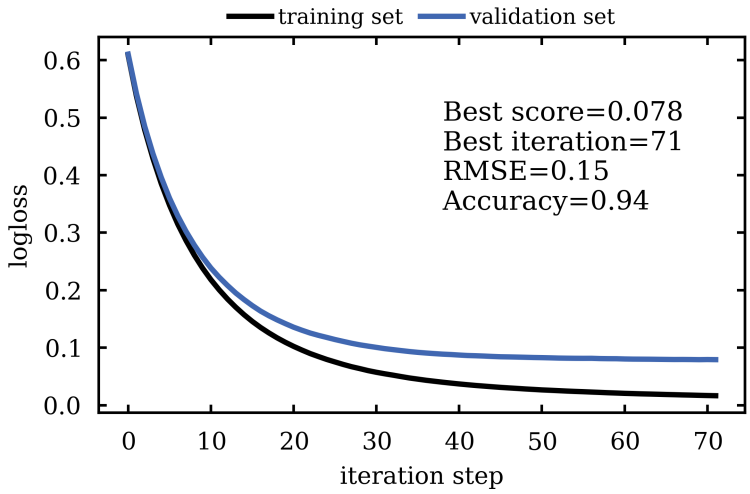

Figure 5.12: The evolution of the logarithmic loss function that XGB model is minimizing. The curves are shown for both the training set and the validation set. The model trains on the 70% of the total data containing a balanced number of each class, at each iteration another 15% of the data is used to validate and penalize the model's performance. The classification parameters are taken as final when the performance in the validation set stops improving.

also increases the false positive rate as more background events are identified as signals. In contrast, strict selection limits background leakage and improves the purity of the selected data, but also misses some potential signals, leading to an increase in false negatives. The true positive and false positive rates as a function of the XGB prediction cut are also shown in Figure 5.13. The chosen value of the XGB prediction cut is discussed in the following sections together with the trigger parameters.

As shown in Figure 5.11, proximity–related features of a peak play a key role in distinguishing supernova signals from background noise. This observation is consistent with the expected characteristics of a supernova signal, which manifests as a burst of neutrino interactions occurring within a short time frame. Consequently, the time intervals between registered signals are significantly shorter compared to those of regular background events.

The signal and background acceptances for each model are summarized in Table 5.5. For comparison, the signal acceptances shown in Table 5.4 for the traditional selection, are also listed. In both cases, the XGB model performs better in identifying the supernova signal, but also allows for a greater percentage of the background. Although the difference may seem minor, it can still be significant given the high signal rates.

Table 5.5: Final signal acceptances for SR0 and SR1 using XGB-based data selection and the traditional data selection.

	SR0		SR1	
model	**Signal**	**BG**	**Signal**	**BG**
XGB with subtype	82.3%	0.1%	78.8%	0.08%
XGB without subtype	80.9%	0.5%	81.7%	0.8%
Traditional Selection	48%	0.05%	47%	0.15%

The confusion matrices, which illustrate the predicted percentages for each class, are shown in Figure 5.14 for SR0 (left column) and SR1 (right column), with the results from the

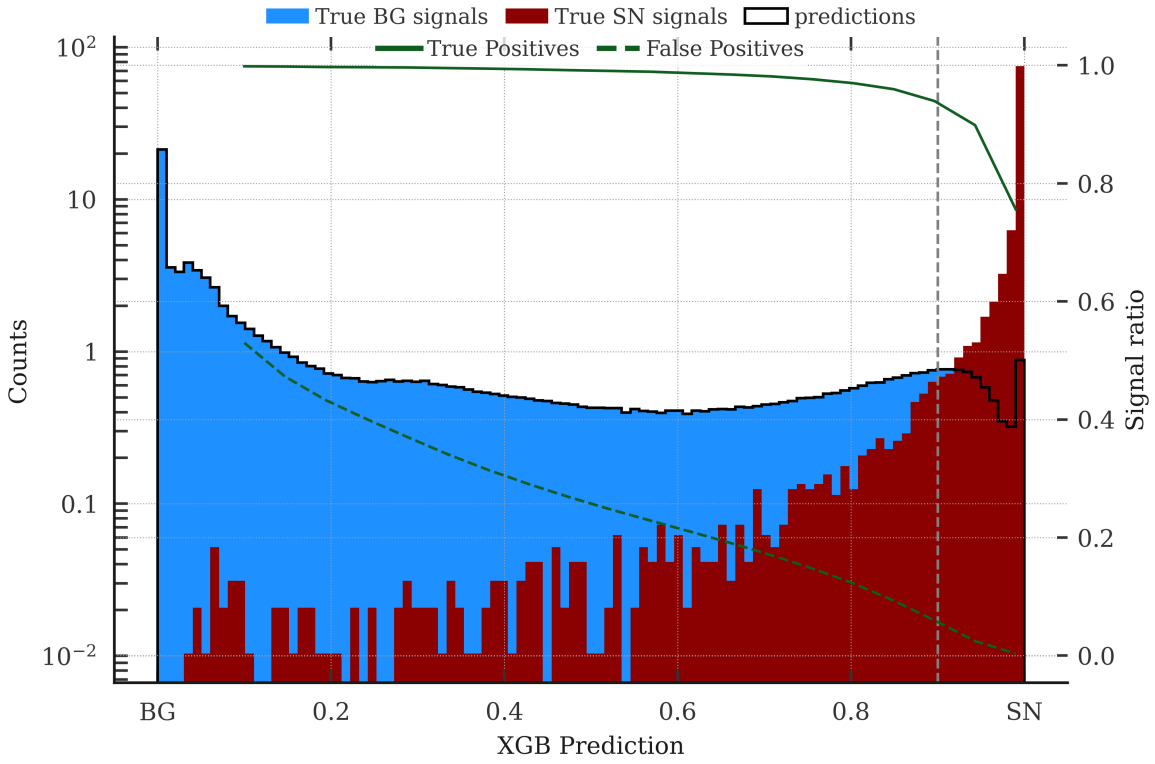

Figure 5.13: The distribution of predicted labels along with the true population. The model shown here does not use subtype information. Shown for the SR1 data salted with the benchmark Bollig model. The XGB model is trained with a subset of the same data. The true positive and false positive rates are shown for XGB predictions between 0.1-0.99.

XGB model with subtype information (upper row) and without (lower row). In all cases, more than 90% of both supernova and background signals were correctly identified. For SR1, incorporating subtype information improves background identification accuracy by approximately 5% and supernova signal identification by about 3%. Unexpectedly, for SR0, the inclusion of subtype information slightly reduces the accuracy of true label predictions. This discrepancy may be due to the comparatively smaller sample size of the salted SR0 dataset.

Comparing the performance of models is not as straightforward, and another metric, called the precision–recall curve, is used. The Precision–Recall (PR) curve is a graphical representation that is used to evaluate the performance of a binary classification model, especially when dealing with unbalanced datasets[v]. It plots precision (the ratio of true positive predictions to the sum of true positive and false positive predictions) against recall (the ratio of true positive predictions to the sum of true positive and false negative predictions) at various threshold settings. The PR curve is particularly useful for understanding how well the model performs in terms of both identifying positive cases and minimizing false positives. A high area under the curve (AUC) in a PR graph indicates both high recall and high precision, suggesting that the classifier effectively identifies positive instances while minimizing the number of incorrect positive predictions. The PR curve for the XGB prediction for salted

[v]In classification models, an unbalanced dataset occurs when one class significantly outnumbers others, potentially biasing the model to predict the over-represented class. This can result in high accuracy for the dominant class, even without meaningful parameter analysis

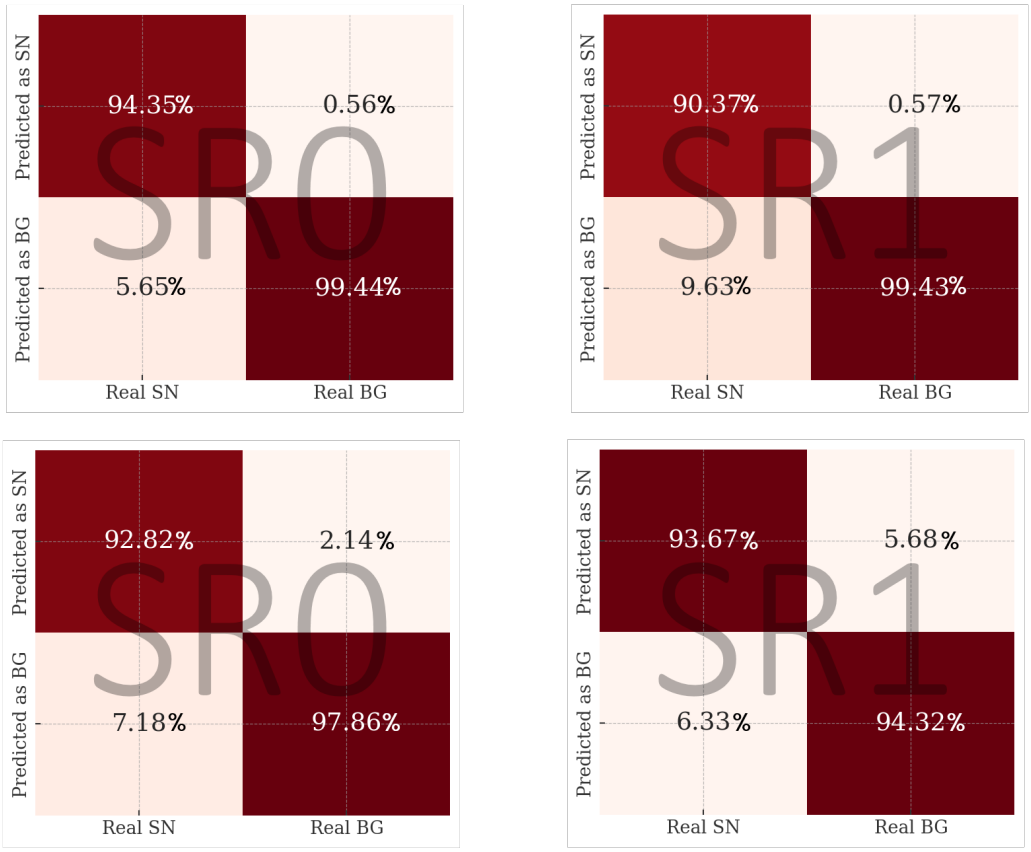

Figure 5.14: Top: Confusion matrix for SR0 (left) and SR1 (right) with subtype information. Bottom: Confusion matrix for SR0 (left) and SR1 (right) without subtype information.

SR0 data with and without subtype information is shown in Figure 5.15, again indicating superior performance when subtype information is provided.

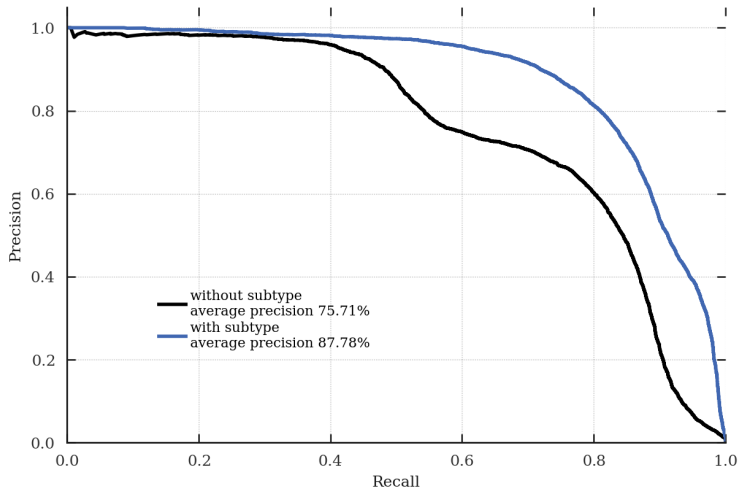

Figure 5.15: Precision–Recall (PR) curve shown for two different types of treatment; with subtypes and without subtypes. The precision is defined as the ratio of total true positives to sum of all positive predictions, and the recall is defined as the ratio of true positives to sum of all real true labels. The average precision calculated as the AUC are also shown, indicating a better classification performance when the subtype information is included.

The models are tuned to recover all artificially injected supernova signals while minimizing FAR, without trying to achieve the highest purity in the selected data. Once the signal is detected, a detailed event-by-event selection can be performed. Injections used in model tuning are sampled from the Bollig 2016 model with a 27 $M_{\odot}$ progenitor at 10 kpc. Since different models, progenitors, or distances can vary the signal shapes, the objective was to develop a generic XGB model that identifies supernova–like signals.

After selecting the signals with a conservative XGB prediction$\geq$ 0.9 threshold, the remaining signals were sorted by time to investigate the rate evolution. It should be noted that even if a high number of background peaks were labeled with a high XGB prediction, they might not be in the same time–vicinity and therefore do not mimic a supernova burst. The rates before applying the selection are shown in blue in the upper panels of Figure 5.16 for the XGB models using subtype information for the SR0 and SR1 data. The final rates after applying the XGB prediction cuts are shown in red in the same figure. In both SR0 and SR1, the injected supernova signals are clearly visible as a sharp rate increase over distinct times after applying the XGB prediction cuts. The same comparison is made for the XGB model without subtype information and is presented in Figure 5.17. It can be observed that, in both cases, the injected supernova signals remain clearly identifiable after applying the XGB selection. Furthermore, the computational cost of processing subtypes for SR1, which was not readily integrated into the existing workflow, makes its inclusion impractical for this stage of the analysis. Therefore, the XGB model without subtype information was chosen for further analysis. However, subtype information can still be computed posteriorly to enhance confidence in data selection if needed.

The rate evolution of the surviving peaks, along with their XGB prediction scores, was used to identify supernova signals in the data stream. For this purpose, the software trigger described in Section 5.2 was used. Once a potential supernova signal region is identified, the raw data within the vicinity of the trigger can be obtained and further investigated. Hence, precise peak selection and trigger settings were tuned together to efficiently detect the hidden signal within the noise.

5.1.4 Event Level Analysis

As described in Chapter 2, the higher–level processing of the data from the XENONnT experiment leverages the full potential of dual–phase TPCs. *Events* are reconstructed by pairing of S1 and S2 signals, providing a complete picture of the interaction. This allows for a discrimination between electronic recoils and nuclear recoils based on the signal ratio of S2/S1. Here, the characteristics of the supernova and background signals at the event–level are described and a basic area selection and filtering is introduced.

The area distribution of the supernova signals in the simulations compared to the background is shown in Figure 5.18. The ROI at the event level is defined in the area distribution, considering only events with an 1 PE $< A(S1) <$ 100 PE paired with an S2 peak with

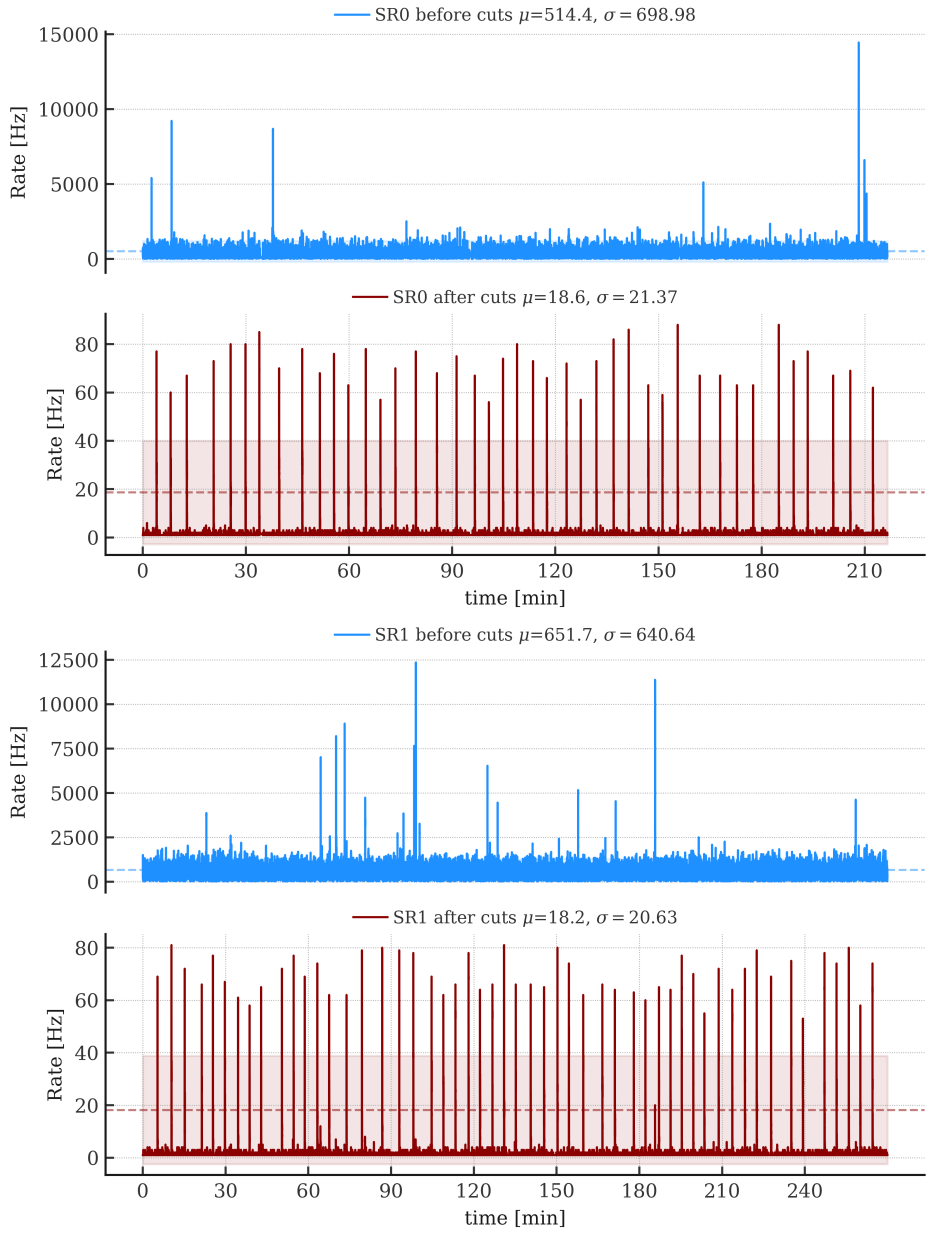

Figure 5.16: Total peak rates before and after applying the XGB model with subtype information for SR0 (top) and SR1 (bottom). The blue lines represent the peak rates before applying the XGB selection, while the red lines show the rates after applying the XGB prediction cut at 0.9. The horizontal dashed lines and the shaded areas correspond to the mean, μ, and 1σ standard deviation reported in the figure. The sharp rate increases indicate the presence of injected supernova signals.

100 PE $< A(S2) <$ 3000 PE[vi]. This region overlaps with the nuclear recoil signals expected from low–energy interactions, where the expected number of background events is low.

The average number of simulated events for each model is shown in Figure 5.19 along with the average number of events after the area selection. The average number of event counts for the SR1 background is also presented for a 10–second window, using the background event data for the 26 runs listed in Table 5.1 used for salting. Before selection, the background

[vi] Note that, prior to any selection, events are reconstructed only if their S2 signal area exceeds the lower threshold of 100 PE.

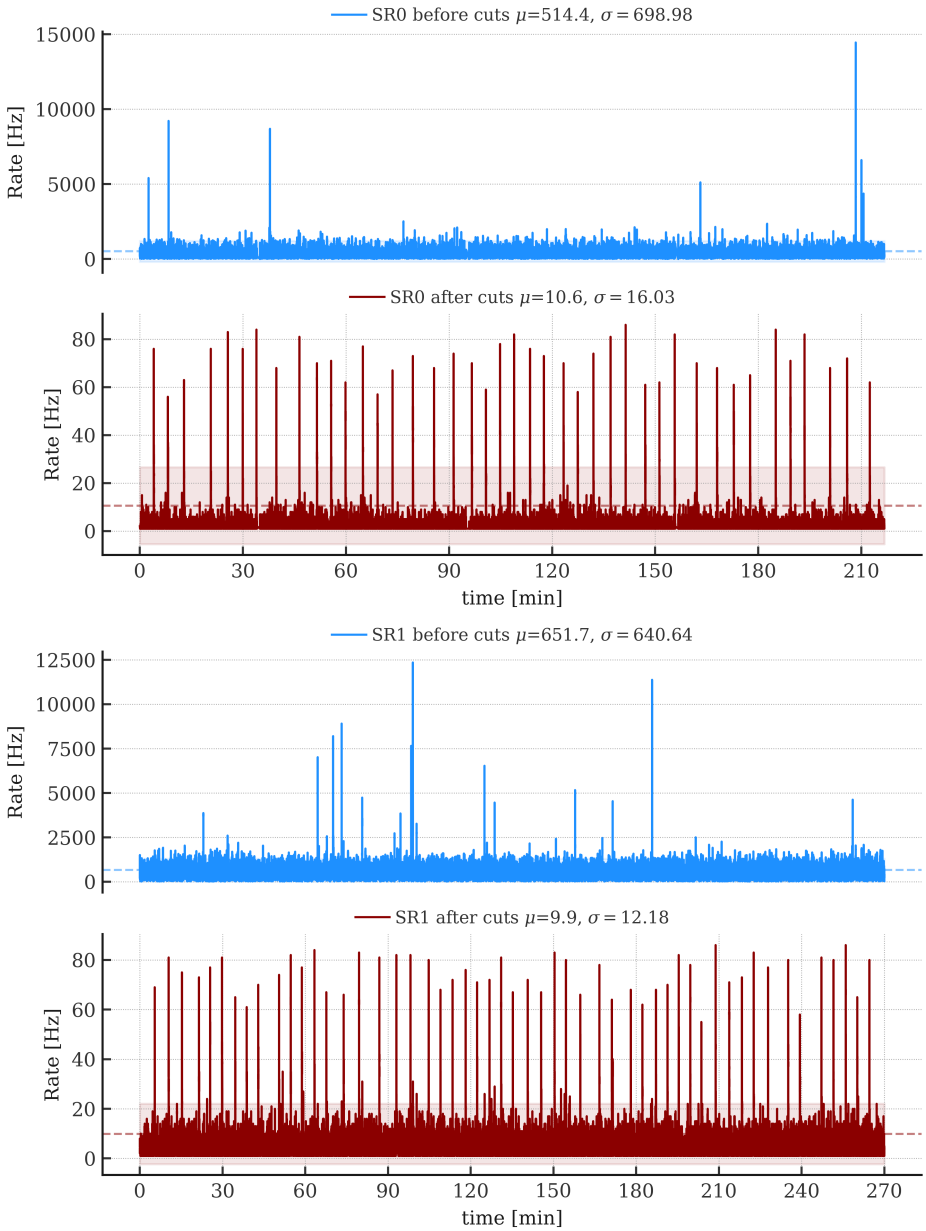

Figure 5.17: Total peak rates before and after applying the XGB model without subtype information for SR0 (top) and SR1 (bottom). The blue lines represent the peak rates before applying the XGB selection, while the red lines show the rates after applying the XGB prediction cut at 0.9. The horizontal dashed lines and the shaded areas correspond to the mean, μ, and 1σ standard deviation reported in the figure. The sharp rate increases indicate the presence of injected supernova signals.

rates were relatively high, but selecting the appropriate region of interest significantly reduces the event rates.

The average event count distribution within a 10–second window in all these runs is shown in Figure 5.20 both before and after the area selection. This figure shows that in the region of interest, the counts of background events were on average $\mu = 1.98 \pm 1.33$ in a 10–second window for the 26 runs considered. This is significantly lower than the expected counts from the models at 10 kpc. Note that the supernova models usually have shorter times, thus the expected counts register in shorter times, which means that the compared background

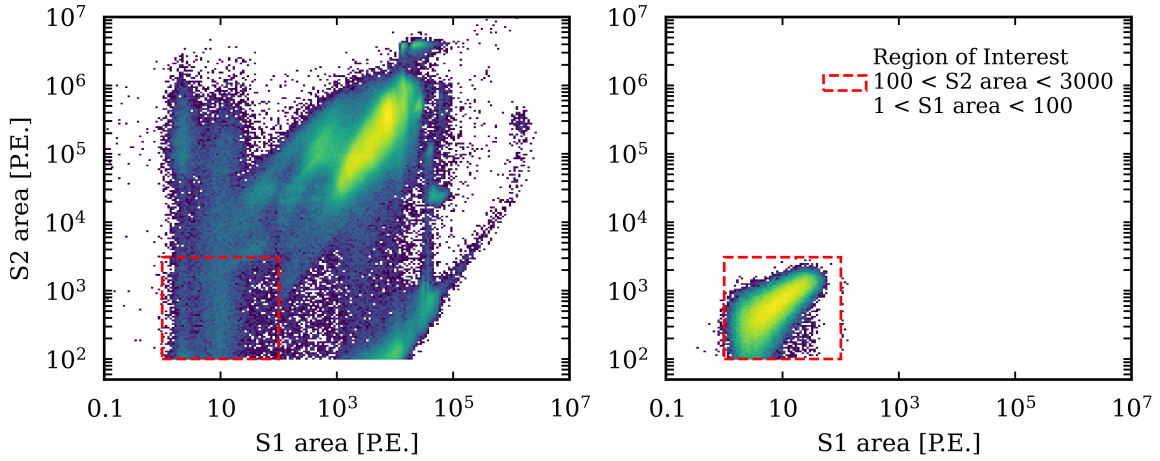

Figure 5.18: Comparison of the area distributions of S1 and S2 peaks that compose an event for background (*left*) and signal (*right*) events. The signal events combine simulated signals from all models.

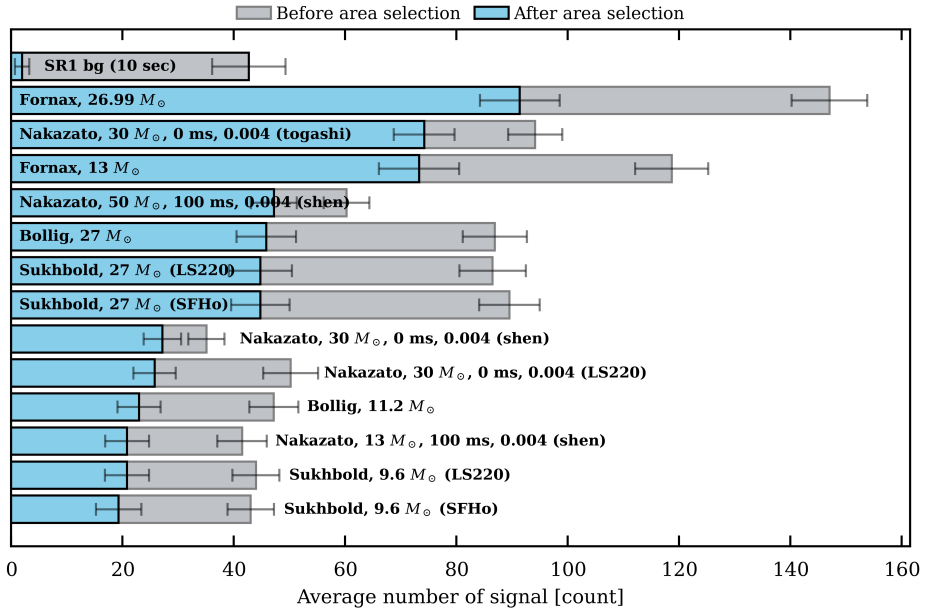

Figure 5.19: Average number of simulated events for each model and the average number of events after area selection, compared with SR1 background counts for a 10–second window. The error bars indicate the 1σ standard deviation from the mean computed over the different simulations of the same model. For the background the error bars represent the standard deviation over all the rolling windows in the considered data.

counts within 10–seconds are conservative.

As discussed in the previous sections, the detected signal will be a superposition of the supernova signal and background. Therefore, the salted datasets described in Section 5.1.2 were used to build events with the XENONnT standard processing framework. This event–building process pairs the S1 and S2 signals based on their timing proximity, which means that processed events may include signals of different origins (i.e. pairing a simulated supernova peak with a real background peak). However, it is easy to identify the time regions where supernova injections occur by examining the event–rate evolution over time.

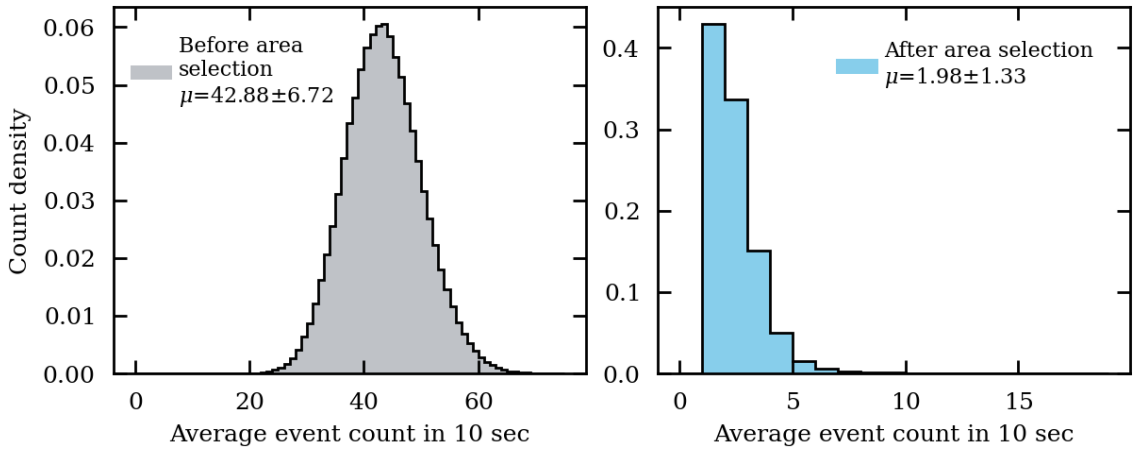

Figure 5.20: Average background event count distribution within a 10–second window for SR1, shown both before (*left*) and after (*right*) area selection.

Figure 5.21 shows this evolution of the event rate within a 19–minute interval, highlighting seven injected supernova bursts. The figure compares the event rates before any selection (*top*) with rates after applying a basic area selection (*bottom*), where only events with 1 PE $< A(S1) <$ 100 PE and 100 PE $< A(S2) <$ 3000 PE were considered.

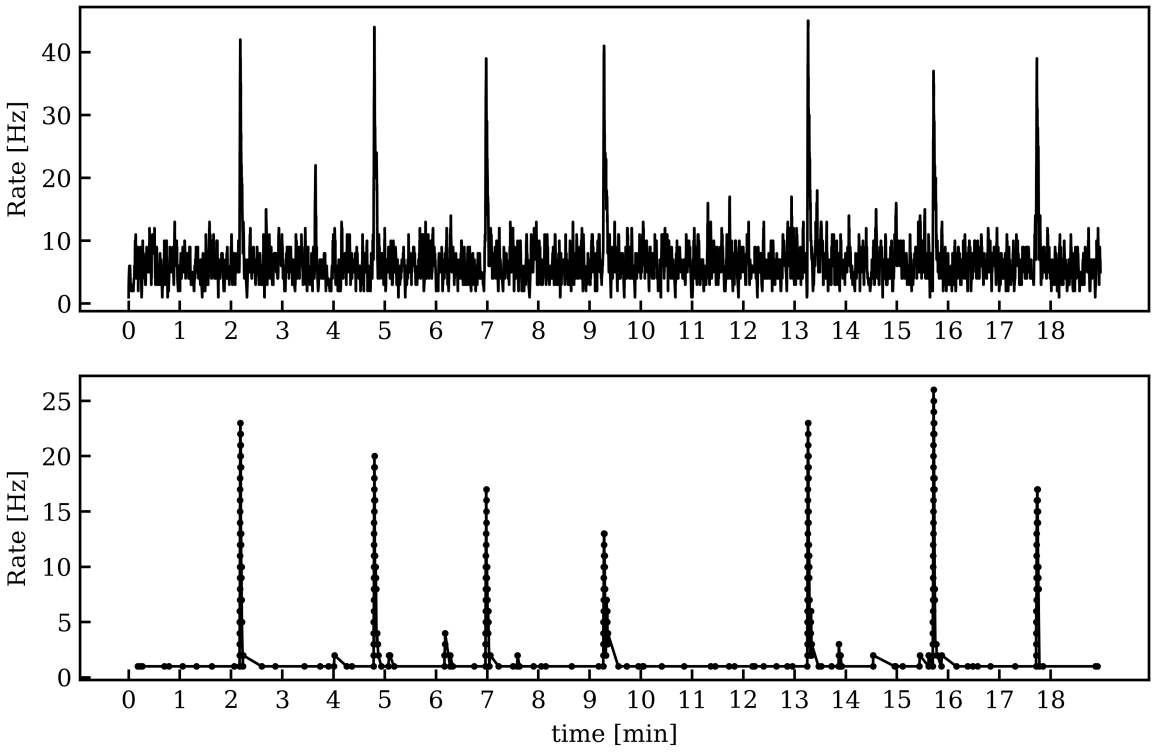

Figure 5.21: Event rate evolution within 19–minute interval showing seven injected supernova bursts. The top panel shows the event rates before any event selection, while the bottom panel displays rates after applying a basic event selection of 1 PE $< A(S1) <$ 100 PE and 100 PE $< A(S2) <$ 3000 PE.

The results demonstrate that the supernova signal stands out against the background, even with the basic area selection criterion. However, to reduce the false positive supernova

identification, an additional data selection is applied to eliminate events with excessive noise preceding large S2 signals. This selection focuses on rejecting events where significant noise, R, occurs before the main S2 signal, which can distort the true S2 area and misrepresent the event. The selection was based on the ratio of noise before the S2 defined as the total area before the main S2 signal, $A(\mathrm{bm}S2)$, subtracted by the S1 area, to the total area, adjusted by the time difference between the S2 signal, $t(S2)$, and the event's start, t'. This is expressed as

$$\mathrm{R} = \frac{(A(\mathrm{bm}S2) - A(S1))}{(A(S1) + A(S2)) \cdot (t(S2) - t') \cdot 10^6}. \tag{5.1}$$

The selection threshold of 0.04 was determined using radon calibration data. Events with a relative noise value divided by the time range less than or equal to 0.04 are kept for further analysis.

5.2 Supernova Hunt: Prompt Identification Techniques

Prompt signal identification within the data is essential to recognize possible supernova occurrences. Signal processing methods are used to locate **signal–peaks**, i.e. supernova bursts, in one-dimensional data arrays, providing a reliable and adaptable approach to identify potential supernova signals. This prompt identification process, often referred to as the *trigger*, initiates further analysis by flagging periods of interest as *potential supernova signals* for closer investigation.

The peak finding functions identify signal–peaks[vii] by analyzing local maxima that satisfy specific conditions. An illustration of the trigger mechanism applied to a supernova signal superimposed on background data is shown in Figure 5.22. These functions allow for the customization of various parameters to tailor the peak finding process to the specific characteristics of the dataset. The key parameters for peak detection include:

- **Height**: Specifies the required height of the signal. Peaks that do not reach the specified height are ignored.
- **Threshold**: Sets the required vertical distance to neighboring points to qualify as a signal–peak. This filters out smaller identified peaks that are not significantly different from their surroundings.
- **Distance**: Defines the minimum horizontal distance (in terms of sample points) between neighboring signals. This helps avoid detecting multiple closely spaced signal–peaks as separate events.
- **Prominence**: Measures how much an identified signal–peak stands out due to its intrinsic height and relative position to other peaks. This helps in identifying peaks that are prominent relative to their surroundings.

[vii] To avoid confusion between the *peak–level* signal of the experiment such as S1 and S2 type signals, the overall shape of the burst is referred to as *signal–peak*.

- **Width**: Specifies the width of peaks at a specified height relative to their prominence. This is useful for identifying broader peaks.

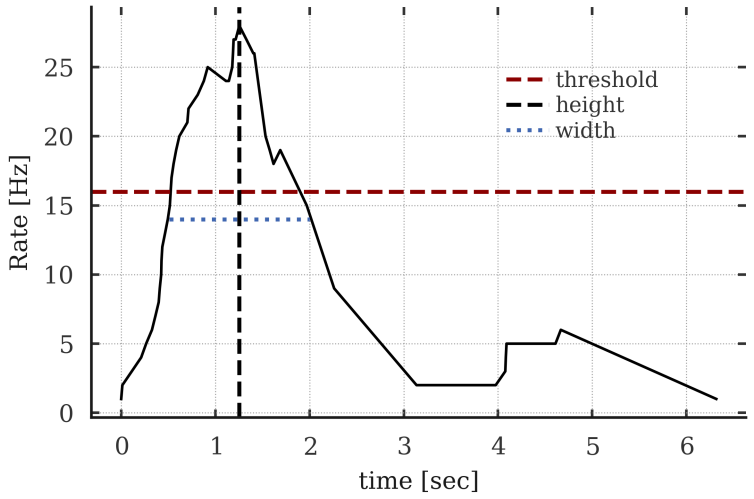

Figure 5.22: Illustration of the trigger mechanism based on the evolution of the signal rate over time. A signal–peak must exceed a specified threshold to pass the trigger algorithm. Additional parameters, such as the height of the highest peak and the width of the signal, are also evaluated to determine whether the signal–peak qualifies as a potential supernova candidate.

The trigger algorithm aims to capture potential supernova signatures in both peak–level and event–level data with slightly different treatments discussed in Sections 5.2.1 and 5.2.2 for peak–level data and in Section 5.2.3 for event–level data. The rate evolutions that contain the time and rate information at each detected data point were used for the trigger. The trigger algorithm identifies the highest rate signal that satisfies the determined set of conditions within a specified range. Following the identification of a trigger time, all data around this time are *clustered* using a window that extends 5 seconds backward and 7 seconds forward, selecting all data within a 12-second interval around the trigger time.

An example of the clusters generated by the trigger algorithm in the benchmark–salted dataset using peak–level data is shown in Figure 5.23. Here, all triggered times are aligned so that their highest points coincide. Similarly, the stacked rate evolution plots for all other models are shown in Figure 5.24.

Large S2 signals can saturate the data buffer, causing the DAQ to release this overflow and register artifact signals that mimic single–electron signals, potentially resulting in false positives. These artifacts are identified by the high–energy vetoes of the detector and labeled as DAQ busy. Upon inspection of the false positive triggers within the background data, a correlation between DAQ busy times and false positive identifications was observed. Therefore, to further refine the detection process, DAQ busy veto cleaning is applied. The DAQ busy-veto-tagged peaks were clustered in 2 second intervals, and the data in this interval were rejected, effectively removing false triggers caused by this effect. An illustration of this cleaning process is shown in Figure 5.25.

The analysis framework was optimized to accurately identify signal–peaks corresponding to potential supernova bursts while minimizing false positives. Although the initial parameters are tuned using the salted data for supernova bursts simulated at 10 kpc, a loose selection

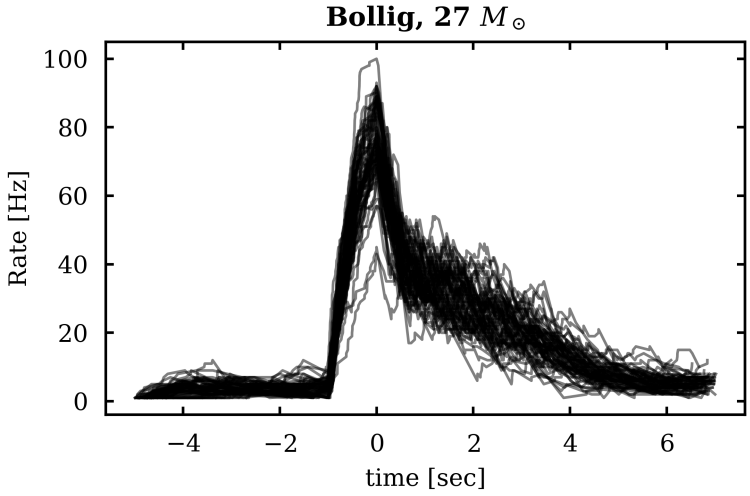

Figure 5.23: Rate evolution curves for the peaks from the benchmark model. The signal–peaks are triggered using the algorithm defined in text, then the rate curves are stacked on their highest peak rates.

was made to account for fainter signals. This ensures that all potential supernova signals were identified even if it produces a higher false–positive rate. Further mitigation of these false positive signals is discussed in Section 5.2.2.

5.2.1 Triggering with the Peak–Level Data

This section describes how the supernova signal–peak was identified and false–positive triggers were mitigated using the peak–level data. At the peak–level all salted data containing different models were used to tune the XGB prediction and prominence parameters to ensure that all injected supernova bursts were captured with the highest accuracy. The prominence parameter of the peak-finding algorithm was chosen as the main parameter to tune due to its definition, which includes the comparison of the relative height of adjacent data points, an advantageous characteristic for handling variable backgrounds.

The cumulative duration of all the *salted background data* is about 90 hours, and a total of 2194 supernova realizations were injected from the 13 models. A range of distinct values for the XGB prediction and the prominence parameter of the trigger was systematically examined using a sample of 328 hours of *raw background data* and all salted background data, and the resulting true positive and false positive counts are subsequently compared in Figure 5.26.

Finally, a selection was made to keep the entire signal with an XGB prediction equal to or greater than 0.95, and to trigger the activation, the prominence value was set to 16, with a width parameter corresponding to 100 ms. This configuration is intentionally loose to ensure high coverage, allowing for 121.6 false positives per week in this dataset, which will later be reduced.

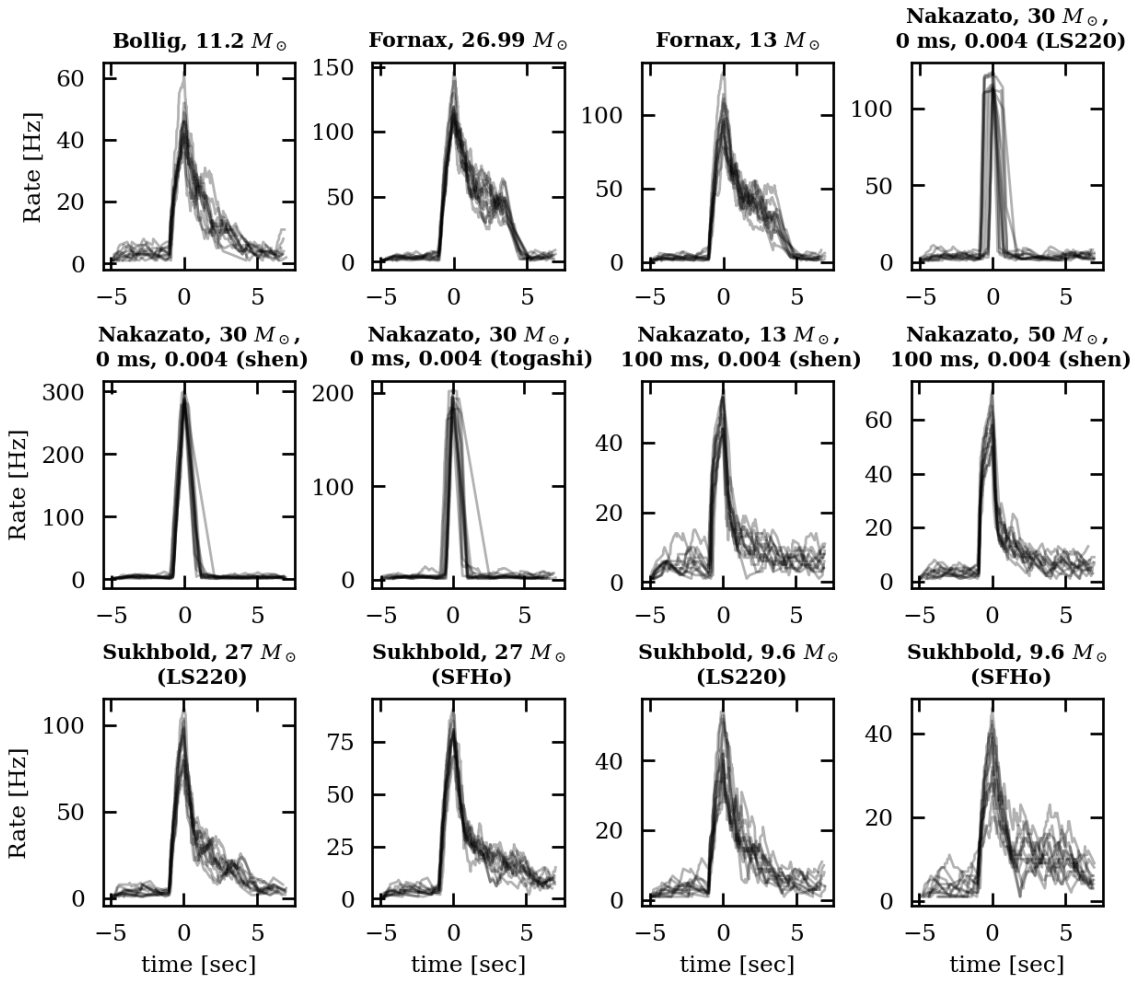

Figure 5.24: Triggered clusters across all salted datasets using the models considered in this study, excluding the benchmark model, which is presented separately. Notably, the failed Nakazato supernova models, characterized by their zero revival times, exhibit sharper and more distinct shapes. This behavior arises from the rapid increase and abrupt cutoff in neutrino fluxes for these models.

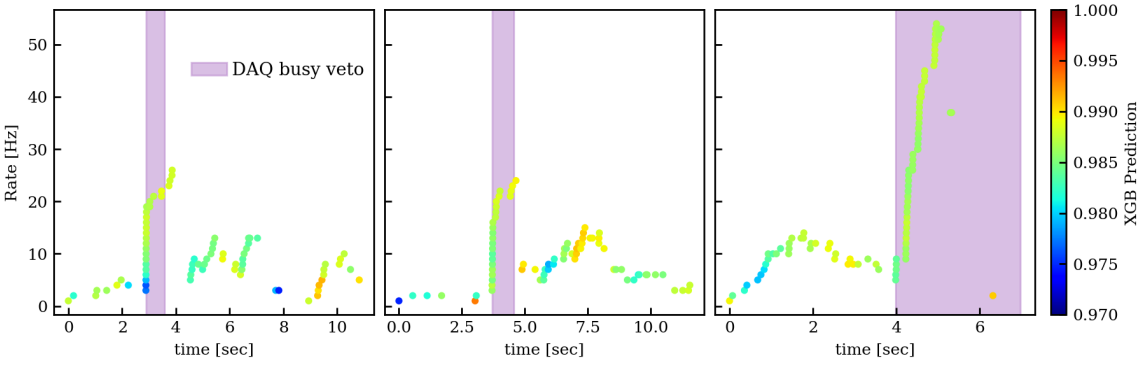

Figure 5.25: The peak rate evolution of three false positive peak clusters are shown along with the XGB prediction of the label of each peak. The purple shaded region corresponds to the clustered DAQ busy veto time intervals. If these signals are clustered within 2 seconds, they are treated as one busy region, and the data falling in this region are removed.

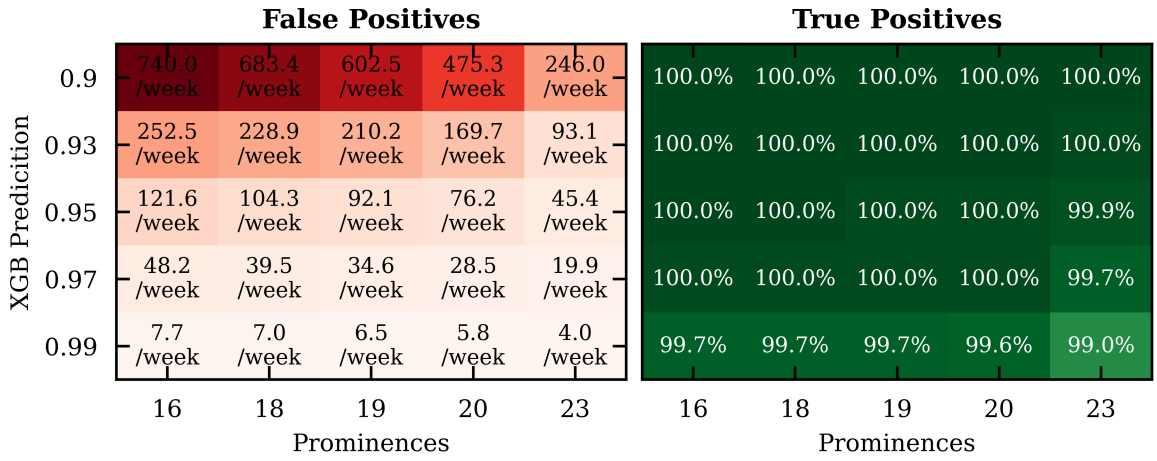

Figure 5.26: Grid scan analysis illustrating the performance of prominence and XGB prediction. The left panel quantifies the weekly false positive counts for each configuration assesed by the sampled ~ 2 weeks of background data, while the right panel depicts the true positive rate as a percentage of successfully triggered supernova out of the total 2194 injections.

Examples of correctly identified supernova signals are shown in Figure 5.27, while three false positive samples in the salted dataset are shown in Figure 5.28. By examining the characteristics of false and true positives, a secondary filtering can be performed on the clustered data by the trigger algorithm. The identified characteristics and the secondary false positive mitigation layer are discussed next.

5.2.2 Post-Trigger Refinement Using Convolutional Neural Networks

Despite passing the XGB selection criteria and satisfying the peak finding algorithm's conditions, false positives often exhibit distinct characteristics compared to true positives, making deep–learning based classification a valuable tool for enhancing trigger reliability. To improve the reliability of supernova signal detection and reduce false positive rates, a Convolutional Neural Network (CNN) architecture was trained using 91 SR1 background datasets combined with 26 salted datasets. These datasets, which amounts to more than 15 days of livetime, was used to study the characteristics of false positive triggers.

The examples of true and false positive identifications shown in Figures 5.27 and 5.28 highlight key characteristics that aid in discrimination. The upper panels show the rate evolution for all S1 and S2 signals that meet the XGB selection criteria, while the lower panels show the spatial distribution of S2 signals. In particular, false positives exhibit a higher ratio of S1 signals within the clustered data. Furthermore, examination of the third example in the false positive selection shown in Figure 5.28 reveals that it likely results from muons crossing the TPC and mimicking a supernova–like signal. Unlike true supernova signals, which are uniformly distributed, this false positive is spatially concentrated in a single location. Finally, the average XGB predictions for the signals in false–positive clusters are significantly lower, as indicated by the color bars. Although tightening the XGB prediction threshold could reduce the false–positive rate, this approach risks missing true signals, particularly from

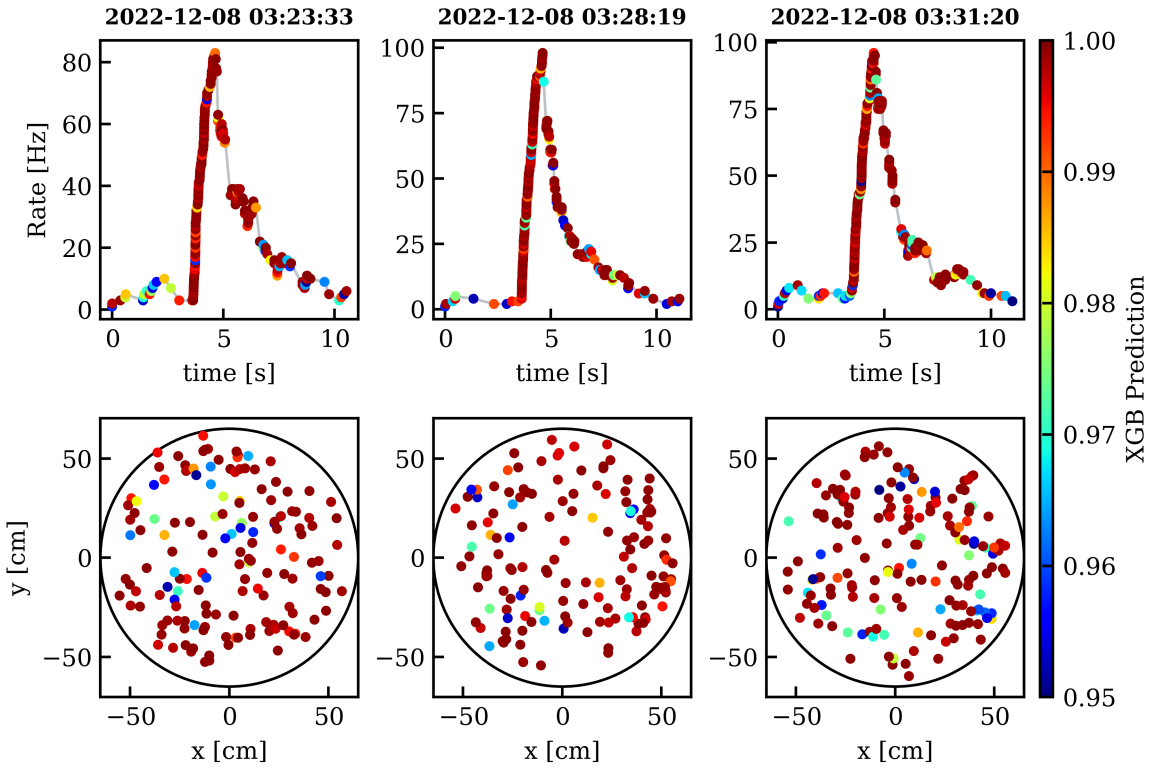

Figure 5.27: Examples of correctly identified supernova signals in the salted datasets using the benchmark model. The top panel shows the rate evolution for three injected signals, while the bottom panel displays the x–y distribution of the S2–type peaks. The XGB prediction for each peak's label is color–coded in both figures. The x–y distribution of the supernova signal is uniformly distributed, as expected.

more distant supernovae. Instead, the lower–XGB prediction signals can be kept, but the information can be used in further classification. These observations suggest that further discrimination within the triggered signals is possible. To train the CNN model, the following steps were taken:

1. **Data Preparation:** The rate evolution data were interpolated to maintain uniform data shapes across samples and included a standardized XGB prediction label.

2. **Position Distribution:** The spatial distribution was quantified by counting peaks in 16 angular slices for each sample.

3. **Peak Types:** The S1/S2 signal number ratio was included as an additional feature, providing information on signal characteristics.

In the initial configuration, the XGB selection threshold was intentionally set at a loose value of 0.95, with the peak finding algorithm configured with a prominence value of 16 and a width of 100 ms. This configuration successfully identified all 2194 injected supernova signals at 10 kpc with 100% accuracy, but also resulted in 520 false positives in these 91 background runs. Then these true and false positives were used to train the CNN model.

The training process involved splitting the data into a ratio of 80/20, with 80% used for training and 20% reserved for testing. During each training epoch, 80% of the training data was used to learn, while the remaining 20% validated the model. The training continued

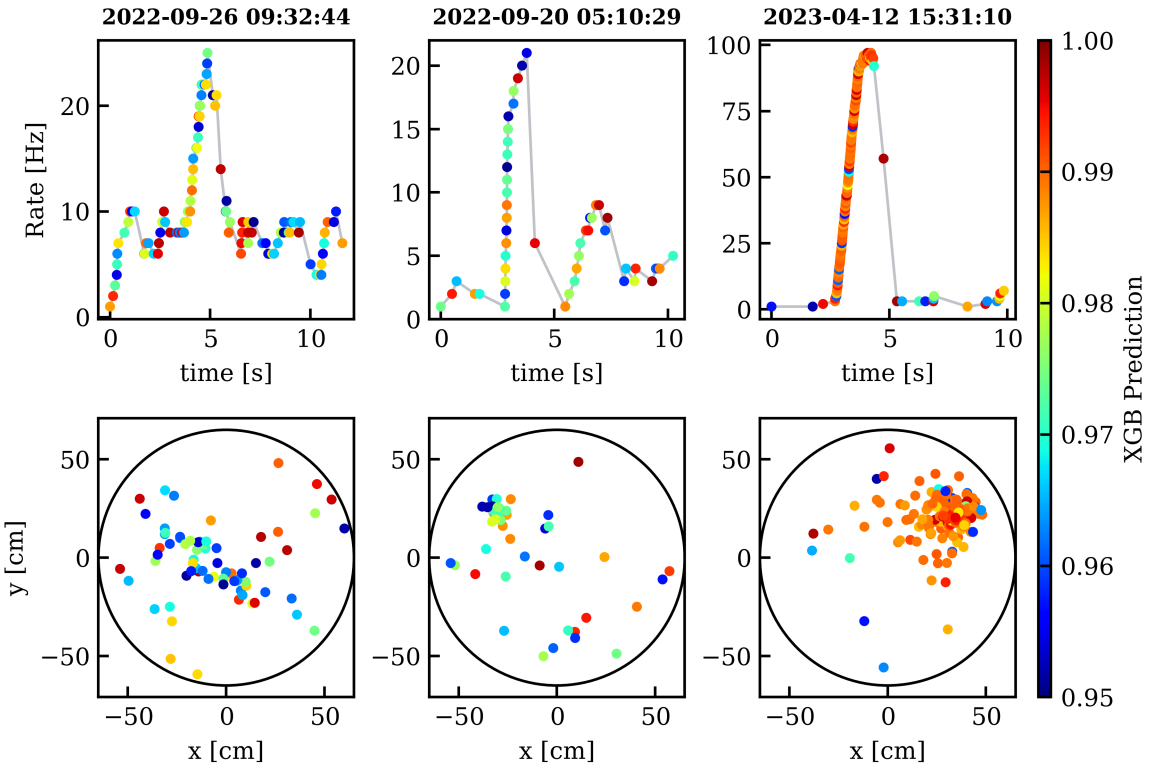

Figure 5.28: Examples of false positive identifications in the combined salted datasets. The average XGB predictions are lower than those in Figure 5.27. Additionally, the number of peaks shown in the bottom panel is fewer for the first and second examples, indicating that most peaks in the cluster are S1 type since the positions are constructed for S2 signals. The third selected cluster at time 15:31:10 is most likely a muon, with the x–y distribution clustered at the position of interaction.

until either a predefined number of epochs was reached or the validation performance plateaued. To prioritize the accurate detection of supernova signals, a custom loss function was defined, assigning a weight of five times to true positives compared to false positives. This approach penalized the misclassification of true supernova signals more heavily than the misclassification of background signals.

The custom loss function, $\mathcal{L}_{\text{weighted BCE}}$, was defined as:

$$\mathcal{L}_{\text{weighted BCE}}(y_{\text{true}}, y_{\text{pred}}) = \frac{1}{N}\sum_{i=1}^{N} \text{weight}(y_{\text{true},i}) \cdot \text{BCE}(y_{\text{true},i}, y_{\text{pred},i}), \tag{5.2}$$

where the weight function was defined as:

$$\text{weight}(y_{\text{true}}) = \begin{cases} 1.0 & \text{if } y_{\text{true}} = 0 \text{ (background)} \\ 5.0 & \text{if } y_{\text{true}} = 1 \text{ (supernova)} \end{cases} \tag{5.3}$$

and the binary cross–entropy (BCE) loss is given by:

$$\text{BCE}(y_{\text{true}}, y_{\text{pred}}) = -\left(y_{\text{true}} \cdot \log(y_{\text{pred}}) + (1 - y_{\text{true}}) \cdot \log(1 - y_{\text{pred}})\right). \tag{5.4}$$

In this formulation, y_{true} represents the true class (0 for background, 1 for supernova) and y_{pred} represents the predicted probability that the sample is a supernova. The loss is computed for each sample in the batch, weighted by the class–specific penalty, and then averaged over the batch to obtain the final loss value.

The model consists of the following key components:

- **Input Layers:** The network has three input layers:
  1. **Main Input Layer**: Processes the rate evolution data, consisting of three features; time, rate, and XGB prediction labels, each sampled at 200 intervals within the triggered cluster.
  2. **Additional Input Layer**: Captures the spatial distribution of the signals, represented by a 16–dimensional vector containing the x-–y position information of the signals within the cluster.
  3. **Secondary Input Layer**: Provides a single feature representing the S1/S2 signal ratio for the cluster.
- **Convolutional Layers:** The main input data was processed through two 1D convolutional layers, each followed by a pooling layer:
  1. **Conv1D Layer 1:** 32 filters, kernel size of 3, Rectified Linear Unit (ReLU) activation, followed by max pooling layer.
  2. **Conv1D Layer 2:** 32 filters, kernel size of 3, ReLU activation, followed by batch normalization and max pooling layers.
- **Dense Layers:** The output of the convolutional layers is flattened and combined with the additional input features through fully connected (dense) layers:
  1. The flattened output from the Conv1D layers is passed through a dense layer with 32 units and ReLU activation, followed by a dropout layer.
  2. The additional x–y input features were processed through two dense layers with 32 and 16 units, respectively.
  3. The S1/S2 ratio was processed through a dense layer with 32 units.
- **Concatenation and output:** The outputs from the dense layers are concatenated and passed through a final dense layer with 1 unit, using a sigmoid activation function to produce the binary classification output (supernova or background signal).

For the final test 20% of the total data is used. The classification of the CNN model is shown in Figure 5.29, after applying a selection at a conservative 0.5 the model successfully classifies all supernova signals in the correct category and in the final dataset correctly identifies 515 of the 520 false triggers as background. The remaining 5 false positives are shown in Figure 5.30. Following this, the 592 background runs for SR1 were evaluated, covering over 83 days of livetime from 2022-07-29 to 2023-08-08. Initially, there were 1395 false positive signals; however, after applying CNN–based selection, 1221 of these false alarms were filtered out, resulting in 174 remaining triggers.

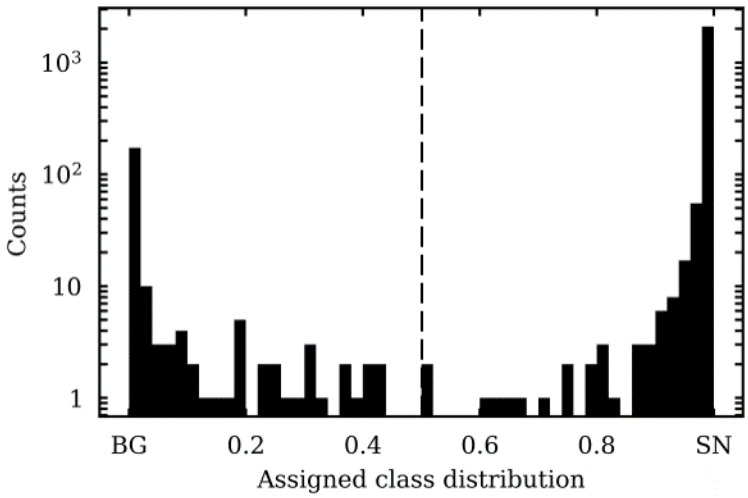

Figure 5.29: Predicted classes for initial triggers. The logarithmic y-axis shows a clear distribution, with class 1 representing supernova signals and class 0 representing background signals, accurately categorized by the CNN model.

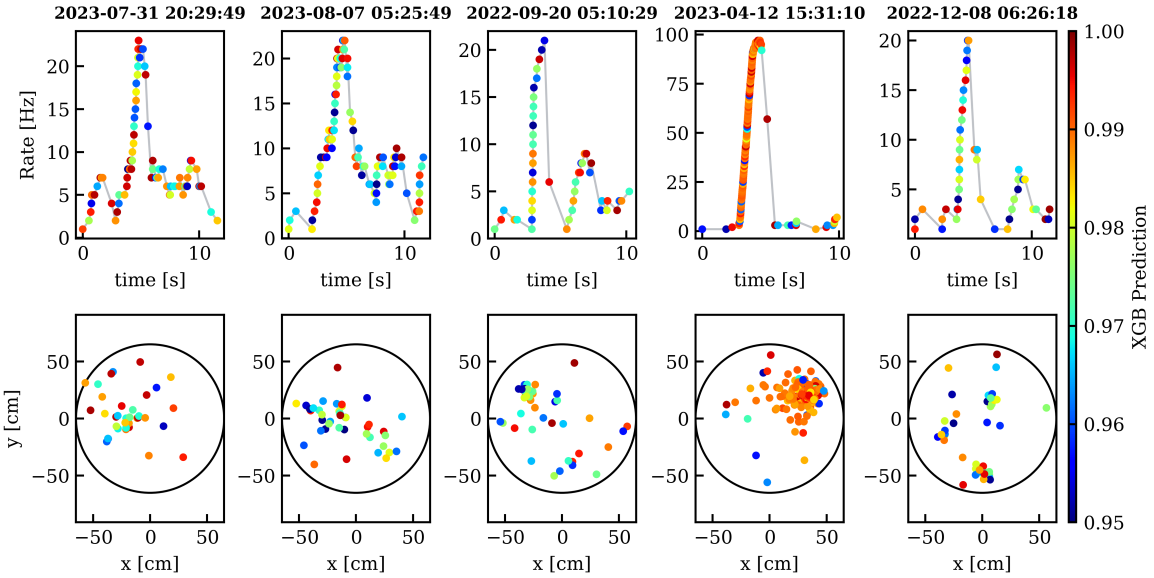

Figure 5.30: The remaining clustered data after peak–level triggering and CNN–based selection. These background signals persist due to the peak shapes and features resembling those of supernova signals.

The filtering results, true positives and false positives as a function of the XGB prediction threshold are shown in Figure 5.31, both before and after applying the CNN filter. The left y-axis represents the true positive rate as a percentage, while the right y-axis indicates the total number of false positives.

It is important to note that the CNN model was trained on data with an XGB prediction threshold of 0.95 or higher, making it less effective when applied to data selected with lower thresholds. The figure demonstrates that stricter peak selection criteria (higher XGB prediction thresholds) reduce the number of false positives but at the cost of lower true positive rates. For example, up to a threshold of 0.95, the true positive rate remains at 100%, while false positive rates gradually decrease. However, even at a stringent threshold

of 0.99, some false positives persist. Additionally, there was one missed supernova signal at a threshold of 0.94 after CNN classification, probably due to additional allowed signals influencing the rate evolution curves, causing the signal–shape to resemble the background, thus being rejected by the CNN model.

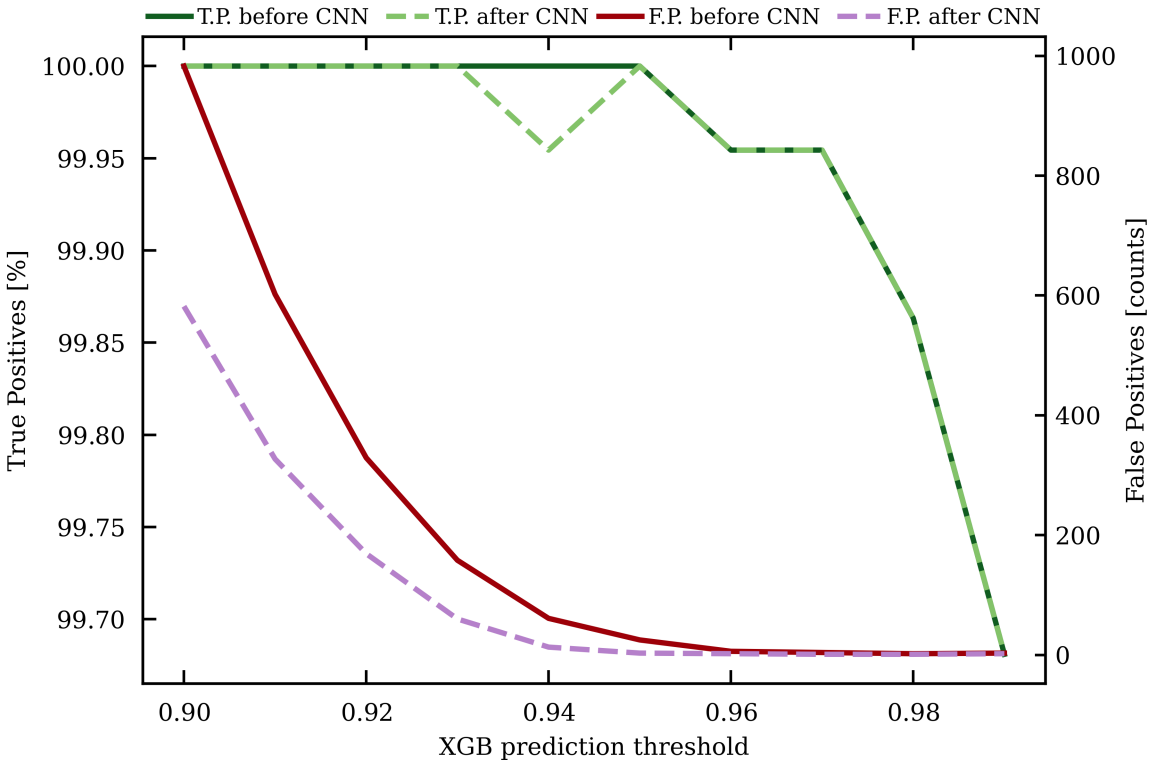

Figure 5.31: True positive (TP) rates and false positive (FP) counts as a function of the XGB prediction threshold, shown before and after the application of CNN classification. The left y-axis represents the TP rate as a percentage, while the right y-axis shows the total FP count. The CNN model was trained using data with an XGB threshold of 0.95 or higher.

It is worth mentioning that although significant improvements were made to the initial triggers, incorporating more data in training could further fine–tune the model and improve its discrimination ability. The role of the event–level trigger in further minimizing the false positive rate will also be addressed in Section 5.2.3. In Section 5.3 we also discuss how the combined analysis at different processing levels and the data from the outer detectors can increase the significance of the detection.

5.2.3 Triggering with the Events Level and Time Matching

A similar triggering mechanism can be applied to event–level data. As discussed in Section 5.1.4, while signal rates are lower compared to peak–level data, the use of event–level information provides significantly better background suppression. The events were constructed from salted peaks to explore the characteristics of supernova signal at the event–level. Since event–level data are inherently cleaner, no strict tuning is required. A basic selection based on the area of signals that make up the event, along with a selection to remove noise as

described in Section 5.1.4 is performed. Finally, before searching for the supernova signature in the event rate, the DAQ busy regions were filtered out.

To detect signals despite the reduced rates, the prominence parameter of the trigger was decreased. The SR1 background runs were analyzed at the event–level using a trigger width of 0.5 seconds and a prominence value of 2. The trigger was applied to the salted events, and the detected clusters were aligned at their highest points and presented in Figure 5.32 for the benchmark model. Aligned signal rates for both peak–level and event–level data are displayed for comparison.

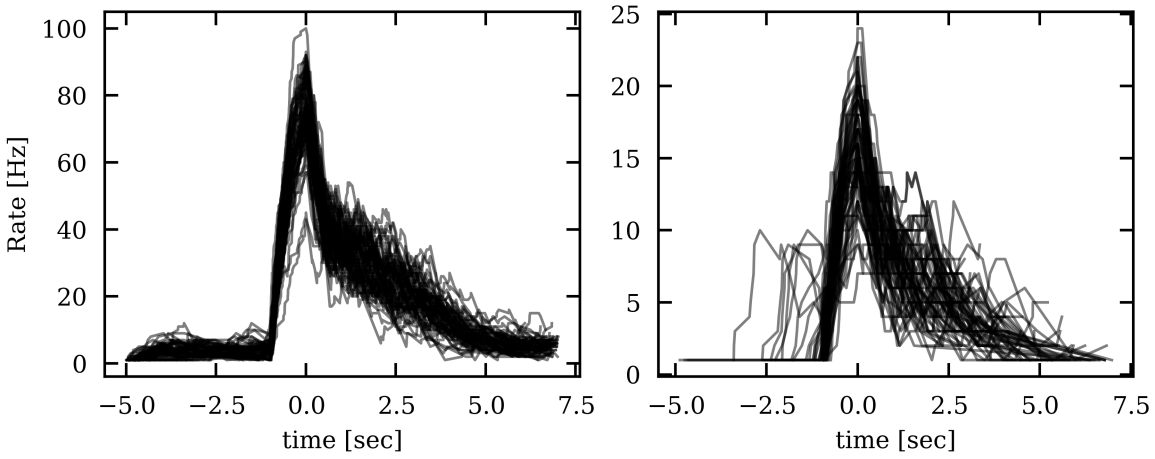

Figure 5.32: Stacked rate evolution of the benchmark model (Bollig 2016) showing peak–level signal (*left*) and event–level signal (*right*) after applying the data selection. Notice that the rates are much lower at higher level data, as only combined signals are allowed to build events, and further selection is applied aftwerward.

Initially, 4 false positive triggers are identified, along with 2193 true positive detections, with one true injection missing. The false positive triggers are shown in Figure 5.33.

Subsequently, the entire set of triggers at both the peak and event levels is temporally aligned within a 5–second margin of time to analyze trigger times and identify matching triggers, that is, to find common triggers. The initial peak–level trigger dataset comprises 2368 triggers, of which 2194 are true positives and 174 are false positives. The initial event–level trigger dataset includes 2197 triggers, including 2193 true positives and 4 false positives. When aligning these datasets within a 5-second tolerance, 2194 matching triggers are identified. The temporal discrepancies between these matching triggers are illustrated in Figure 5.34 for all data, and also only for the benchmark Bollig model with 27 $M_\odot$ progenitor.

Although the triggering times may align, the construction of the events from the S1 and S2 signals can cause the assigned time of the event, and consequently the trigger time, to differ by up to a few seconds. In addition, various data selections yield different final signal shapes in which the position of the highest point varies.

In the final matching triggers, 3 out of the 4 false event level triggers are rejected. The surviving trigger is matched with a false positive trigger in the peak level data, as shown in Figure 5.35, which most likely corresponds to the muon crossing discussed earlier.

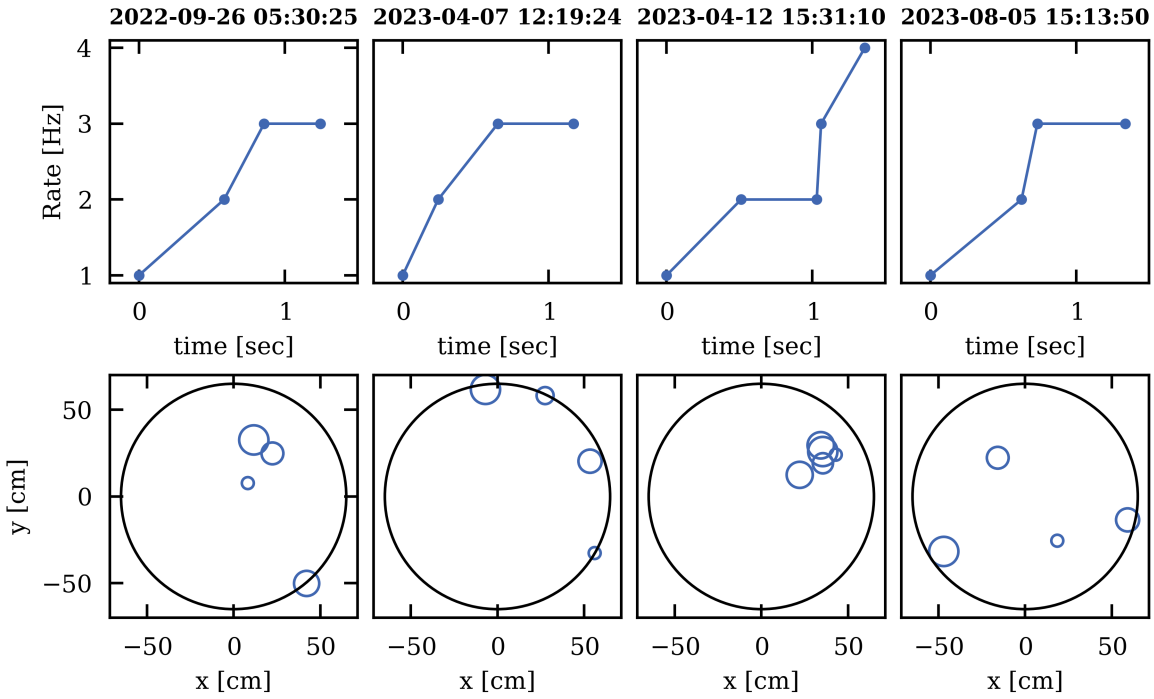

Figure 5.33: False positive signals in events. The area of the circles indicating the positions of the events in all panels are scaled by the interaction times, where the smaller circles indicate earlier interactions. It can also be seen that the majority of the surviving events are at the edges of the detector indicating possible emissions from the walls of the detector.

Additionally, we examined the injected supernova burst recognized at the peak level but absent at the event level. The triggering cluster at the peak level and the corresponding region at the event level are presented in Figure 5.36. The events that accumulate around this region largely fail to meet the event selection criteria and are significantly diminished. The remaining events do not form a pattern that activates the signal detection algorithm.

In conclusion, supernova models located at a distance of 10 kpc from Earth are typically detectable within the XENONnT TPC data. This study demonstrates that there are two primary approaches to investigating the signatures of neutrino interactions, each offering distinct advantages. Event–level analysis provides a cleaner dataset with minimal false positives. However, it is important to note that events are only constructed from peaks that meet specific criteria, such as a minimum S2 area of 100 PE, and a minimum three–fold detection of the S1 signal. As a result, when the interaction rates are lower, such as those that are expected from distant supernovae, or the S1 and S2 signals do not meet these conditions, the events may be significantly reduced, potentially leading to challenges in detection. In such cases, peaks that do not form events can still be utilized to identify the supernova burst. To ensure robust detection, a smart strategy involves monitoring both peak- and event–level data and comparing them.

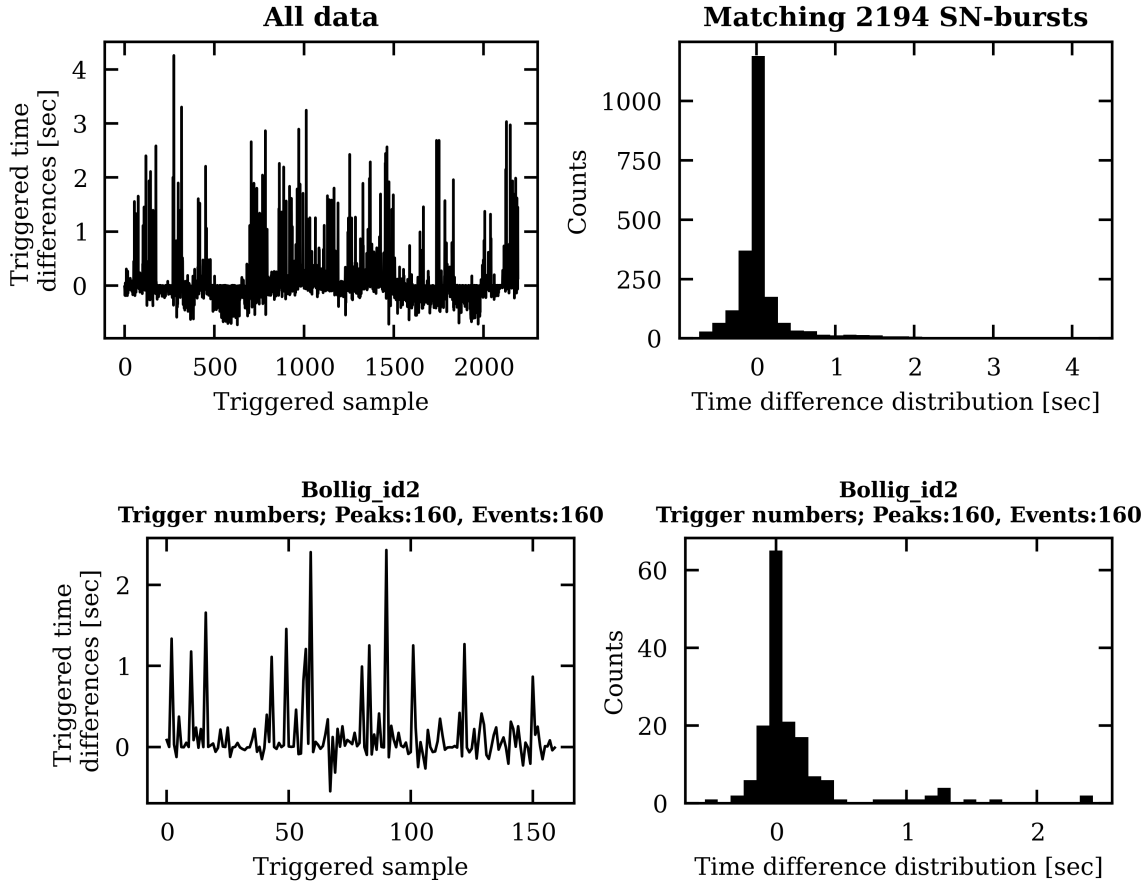

Figure 5.34: Time differences between matched peak and event level triggers. *Top* panel shows the matched data for all identified signals, *bottom*, shows only the identified signals from the benchmark Bollig 2016 model with a 27 $M_{\odot}$ progenitor. Typically, times align within half a second, but discrepancies can extend to a few seconds. For the benchmark model discrepancies between the matching signals are less.

5.3 Quantifying the Significance of Observations

This section discusses the significance of each observation around the identified supernova burst. Local detector signal activity was estimated using background signal rates just before the identified supernova region. The clustered region around the trigger time was compared with this average background rate. Local significances, which highlight the likelihood of random background fluctuations in the signal region, were derived using the likelihood ratio test (LRT).

LRT compares the likelihood of two hypotheses: the null hypothesis (H_0), where only background is present, and the alternative hypothesis (H_1), where both signal and background are present.

The general expression for significance Z within the LRT framework is given by

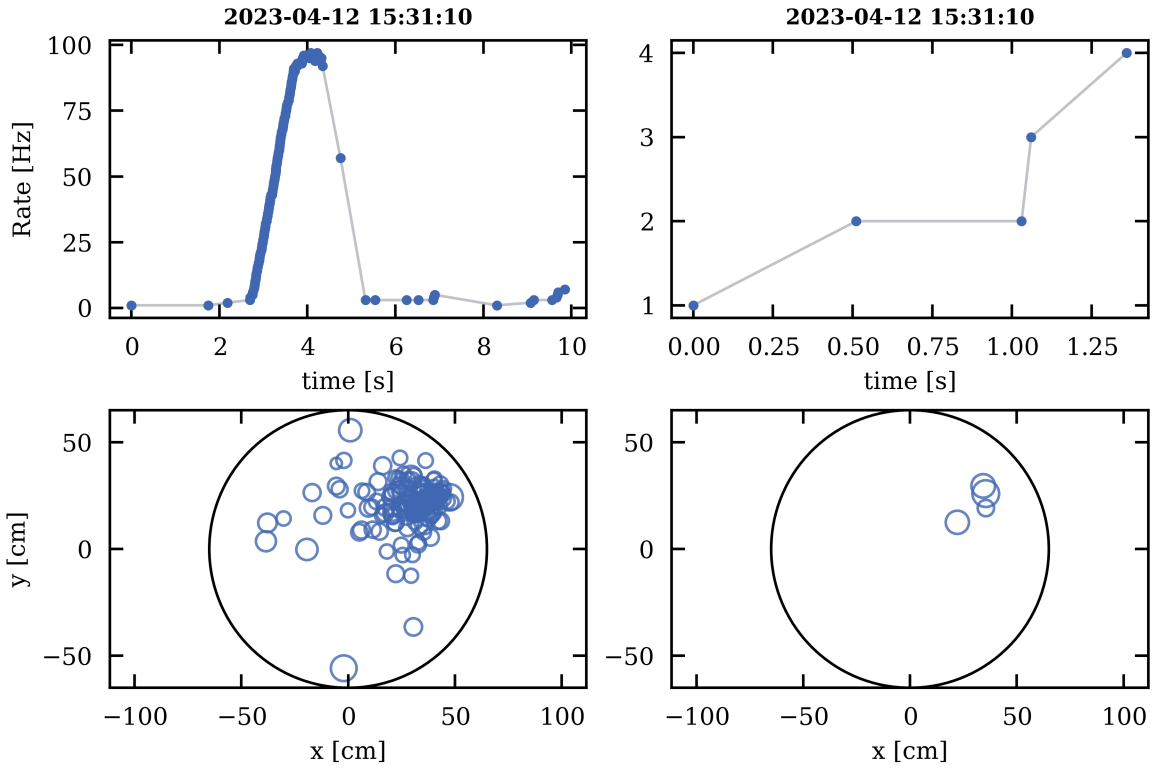

Figure 5.35: The peak (*left*) and event (*right*) rate evolution and the corresponding x–y distribution of the only surviving false positive event. It can be seen that the clustered data is a muon crossing. The area of the circles indicating the positions of the S2–peaks in both panels are scaled by the interaction times, where the smaller circles indicate earlier interactions.

$$Z = \sqrt{2 \ln \left(\frac{\mathcal{L}(\text{data} \mid H_1)}{\mathcal{L}(\text{data} \mid H_0)} \right)}, \tag{5.5}$$

which quantifies the extent to which the observed signal diverges from the background in terms of Gaussian standard deviations σ. Typically, the value of the specified test statistic is conveyed through Gaussian deviations employing the normal distribution. A deviation exceeding $3\,\sigma$, which encompasses 99.87% of the signal, would signify the presence of an *evidence* of non–standard signal within the data. Consequently, an observation exceeding $5\,\sigma$ (corresponding to 99.99% of the whole signal) from the mean indicates a *discovery*. However, in this context, the LRT serves as an intermediate test statistic that is subsequently converted to Gaussian significances for each observed data point, thereby allowing for an interpretation of the magnitude of deviation from the background, and should not be treated as evidence or discovery in the general sense.

In cases where both signal and background are present, the Poisson likelihood for the data under the alternative hypothesis is:

$$\mathcal{L}_{\text{alt}} = \frac{(s+b)^n e^{-(s+b)}}{n!} \tag{5.6}$$

where s is the expected signal, b is the expected background and n is the number of observed

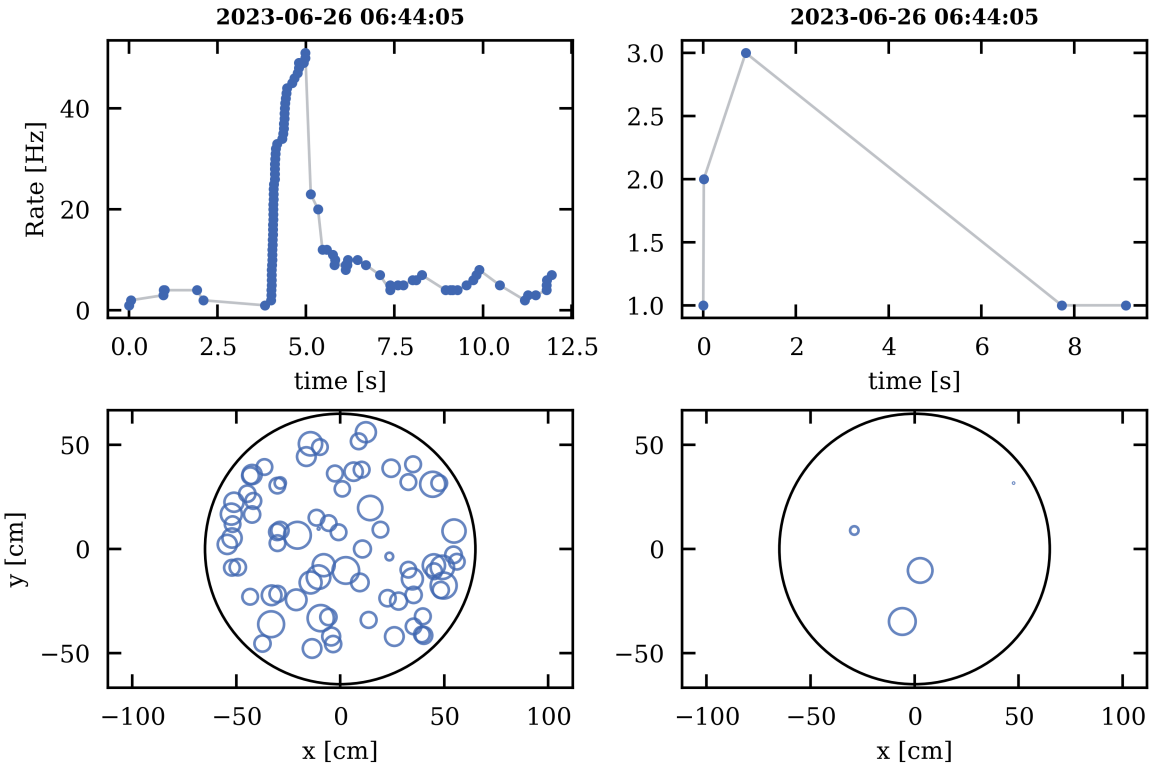

Figure 5.36: The true signal injection using the Nakazato 2013 model with the 50 $M_\odot$ progenitor. The left panel shows the peak rate that triggers at the peak level, while the right panel displays the resulting events in the same region, where many events do not meet the trigger conditions at the event level and thus are not displayed here.

events. For the null hypothesis, the Poisson likelihood is:

$$\mathcal{L}_{\text{null}} = \frac{b^n e^{-b}}{n!}. \tag{5.7}$$

Calculating the likelihood of each data point ($n = 1$) and defining the statistic of the likelihood ratio test statistic λ as the logarithmic likelihood ratio of these two hypotheses, the relation in eq. (5.9) is obtained.

$$\lambda = -2 \ln \left(\frac{\mathcal{L}_{\text{null}}}{\mathcal{L}_{\text{alt}}} \right) \tag{5.8}$$

$$= 2 \left[(s+b) \ln \left(\frac{s+b}{b} \right) - s \right] \tag{5.9}$$

This test statistic λ quantifies the likelihood of the observed data under the alternative hypothesis relative to the null hypothesis. In practical applications, background uncertainties (σ_b) frequently arise. To address this issue, Glen Cowan's formula [152] expands the standard LRT by integrating the σ_b,

$$\lambda = 2 \left[(s+b) \ln \left(\frac{(s+b)(b+\sigma_b^2)}{b^2 + (s+b)\sigma_b^2} \right) - \frac{b^2}{\sigma_b^2} \ln \left(1 + \frac{\sigma_b^2 s}{b(b+\sigma_b^2)} \right) \right]. \tag{5.10}$$

This modification adjusts the likelihood ratio to account for the uncertainty in the background. The significance Z is then defined as the square root of the test statistic. Substituting the expression for λ into the equation for Z, we obtain the following formula for significance:

$$Z = \sqrt{2\left((s+b)\ln\left(\frac{(s+b)(b+\sigma_b^2)}{b^2+(s+b)\sigma_b^2}\right) - \frac{b^2}{\sigma_b^2}\ln\left(1+\frac{\sigma_b^2 s}{b(b+\sigma_b^2)}\right)\right)} \tag{5.11}$$

This expression provides a detailed explanation to which extent the observed signal deviates from the null hypothesis, characterized by the anticipated background signal and its associated uncertainty, thus offering a thorough method for rejecting the null hypothesis in the case of significant deviations.

The process is illustrated in Figure 5.37 where the 12 seconds around the identified supernova burst peak was selected. The mean local background signal rate and a signal region containing 68% of the total signal (corresponding to the Gaussian $1\,\sigma$ interval) were computed within the last ~ 1.5 minutes before the identified supernova region, shown with the dashed line and the shaded area. For visualization, only the last 20 seconds are shown. The significance values shown on the right y–axis were calculated using eq. (5.11) to estimate the probability that the observed signal deviates from the observations of only the background scenario. The maximum local significances of each observation at different data processing levels are discussed next.

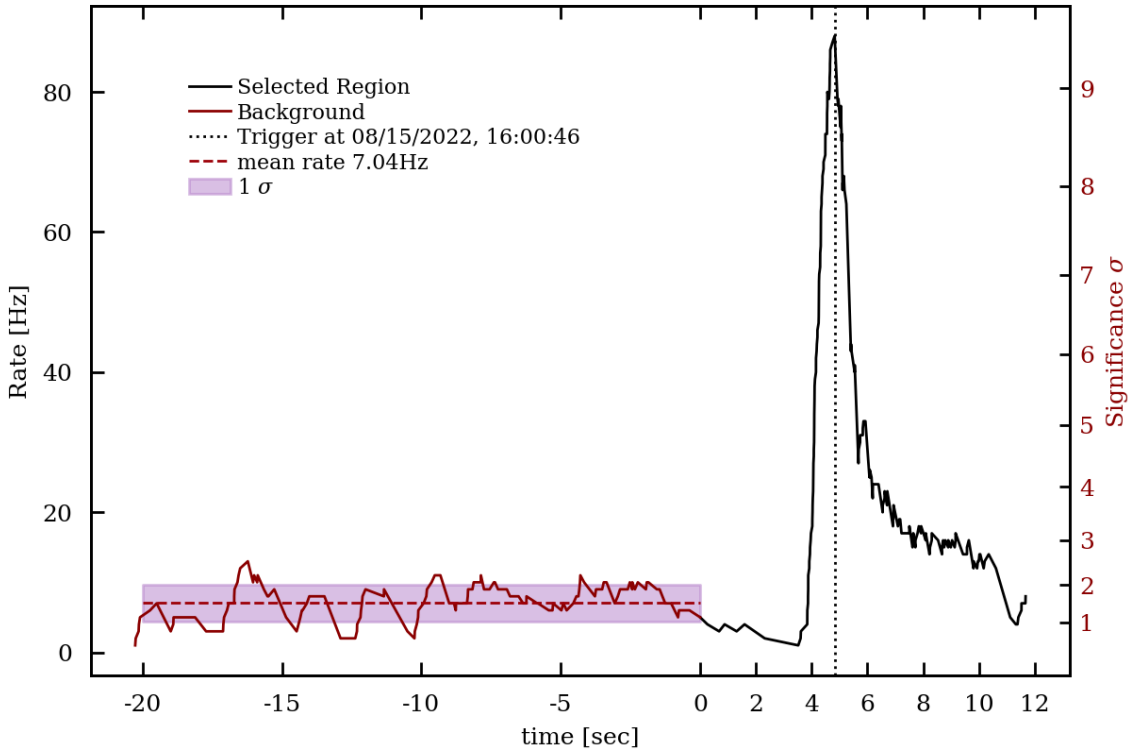

Figure 5.37: The peak rate evolution around an example trigger region is depicted. The mean and standard deviation were calculated using 2 minutes of background data leading up to the signal; however, only the last 20 seconds are illustrated. The signal region was selected as the 5 seconds before and the 7 seconds after the triggering peak. The rates and their corresponding significance, Z, are shown on the left and right y–axes, respectively.

5.3.1 Analyzing Peak and Event Level Significances

By employing signal rates obtained from both peak and event–level analyses as data inputs, we ascertain the local significances of the detected signals. A comparasion of the identified signal–peaks at both levels, is presented in Figure 5.38. This figure additionally provides the computed local significance value for each signal. The comparison shows that the signal rates for both data levels have comparable values of local significance.

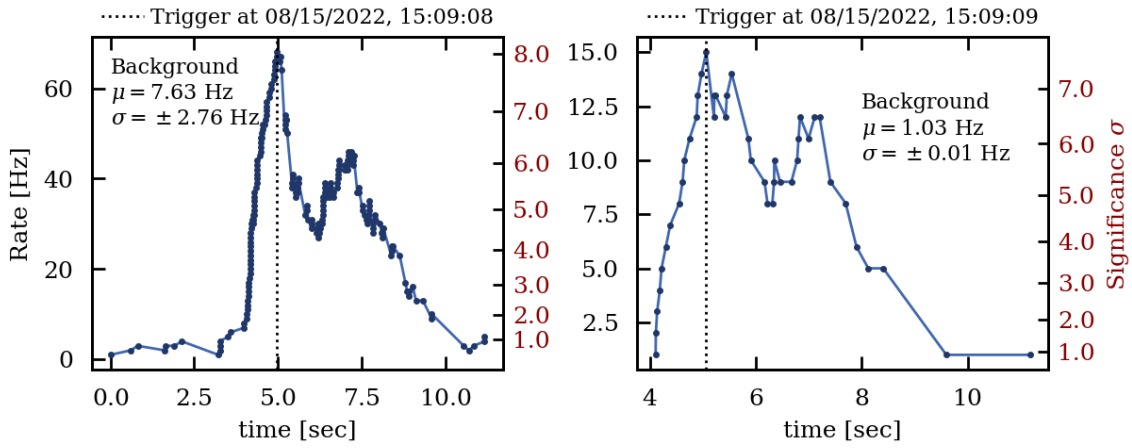

Figure 5.38: Example of the signal rate evolution and the calculated local significances at peak (*left*) and event (*right*) level of the same detected signal–peak. The local mean background signal rates and $1\,\sigma$ deviations are also shown.

At peak–level, background rates were measured locally within 1.5 minutes preceding the signal, and a Gaussian error was calculated based on the background signal in this interval. An illustration of this process is presented in Figure 5.37 for a triggered signal in salted data using the Bollig 2016 benchmark model of with a 27 $M_{\odot}$ progenitor. Local significances are calculated for all true triggers at the peak level for each model. The significance values at the highest peak rates are collected and presented in Figure 5.39.

For event–level analysis, a simpler approach was adopted using a fixed mean and standard deviation calculated from all runs and used to compute significances for all triggered clusters. Upon applying the selection, the event–level background rates become significantly lower and more stable. The background rate for each run was calculated individually, and the mean rates were compiled and presented in Figure 5.40. This figure also displays the overall mean and standard deviation across the considered background runs. It is important to note that while the runs are carried out in sequence, they are not necessarily consecutive due to interruptions and calibrations between different runs.

Analogously to Figure 5.39, the local significances derived from the event–level analysis were compiled, and the corresponding highest values were aggregated. The distribution of these significance values is illustrated in Figure 5.41, providing a visualization of the highest local significances observed in all events analyzed.

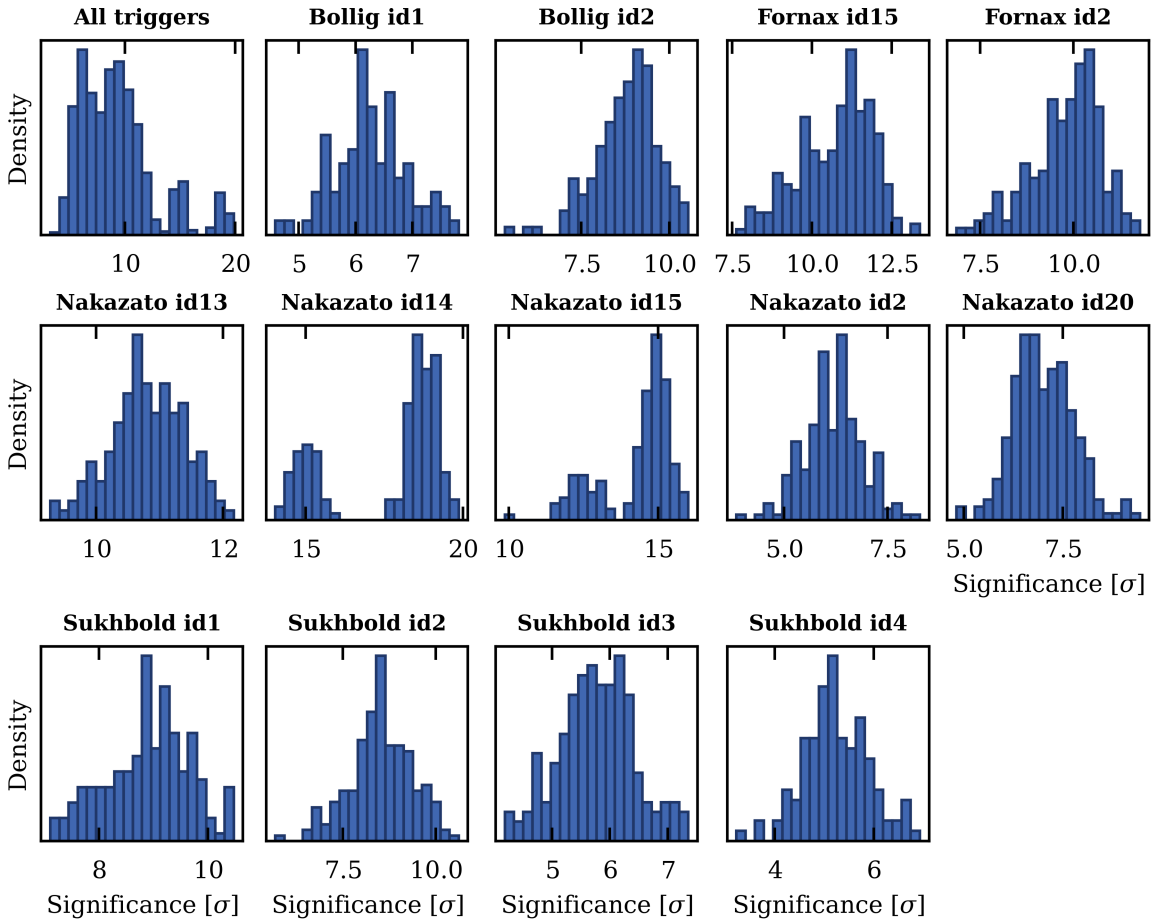

Figure 5.39: Distribution of maximum local significances within clustered data around peak–level triggers for various models. Black hole models show the highest significances due to their sharp rate rise within less than a second, but high significances are observed across all models.

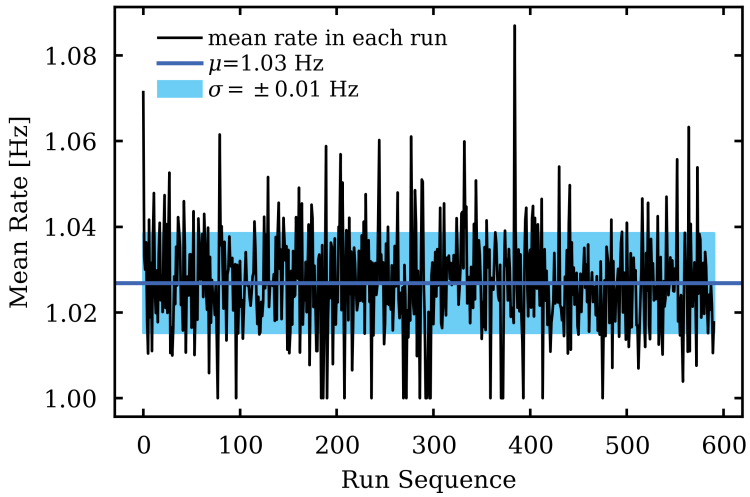

Figure 5.40: The mean rates in events for each run after applying the basic event selection. The average rates per background run are displayed individually together with the average of all mean rates (blue line) and the standard deviation througout the considered runs (blue band).

In summary, the results from both peak–level and event–level analyses demonstrate that statistically significant deviations from the background can be observed across all supernova

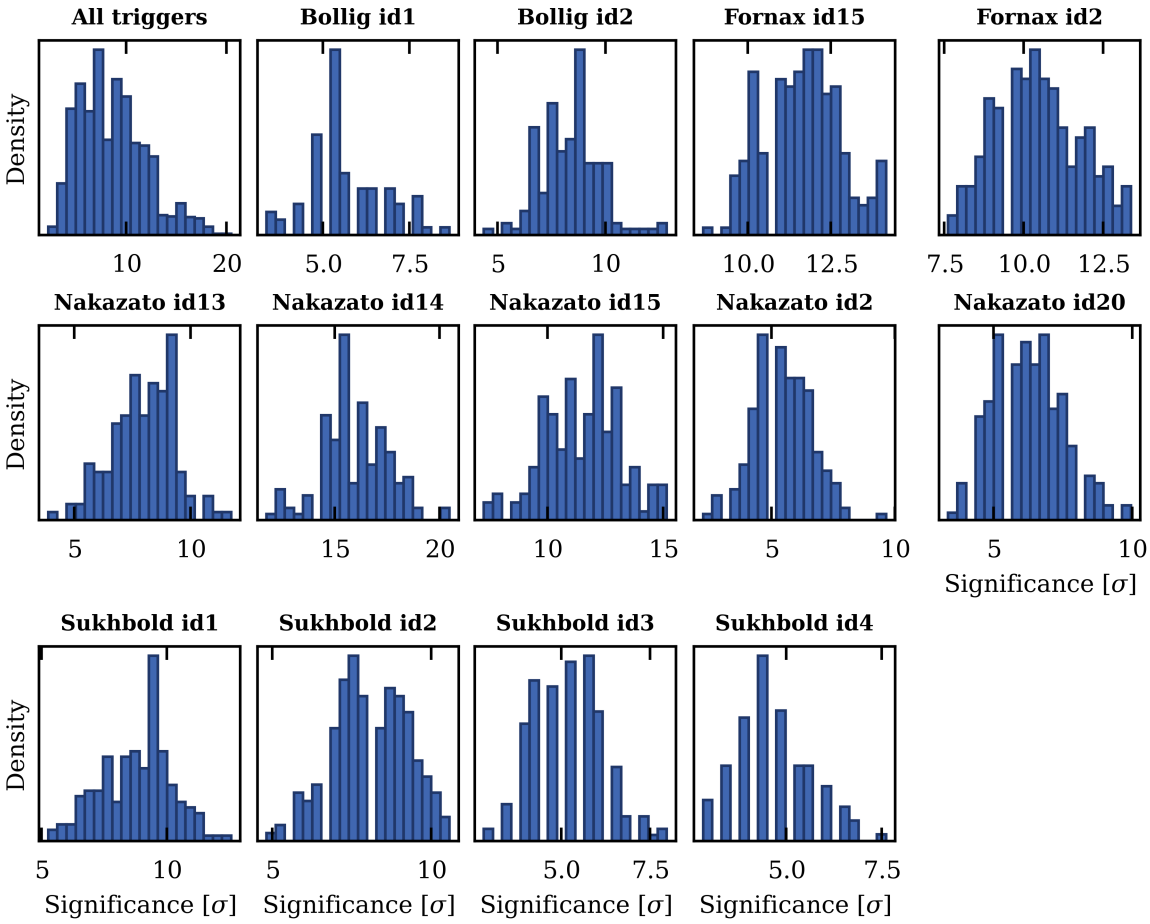

Figure 5.41: Distribution of maximum local significances within clustered data around event–level triggers for all models (*upper left*), as well as for each model considered in this work.

models. The local significance values for the various models consistently show signals that stand out well beyond the typical background fluctuations, suggesting that the XENONnT detector is capable of identifying supernova neutrinos under realistic conditions.

At peak level, models involving black hole formation, in particular the Nakazato model (*id=14*) with the 30 $M_\odot$ progenitor with Shen EOS, exhibit the highest significance due to their sharp rate increase within a short time frame, as observed in Figure 5.39. The total expected interaction rates from this model can reach up to 300 Hz (see Figure 5.24) as all expected interactions are registered in less than a second.

In a similar manner, event–level analysis, while exhibiting lower overall rates and more stable background conditions, provides an additional layer of verification for the detected signals. The compiled background rates, as shown in Figure 5.40, confirm that the event–level data is less prone to background fluctutations. This results in consistently high local significance values, as demonstrated in Figure 5.41.

Overall, the XENONnT detector demonstrates the ability to reliably distinguish neutrinos from a supernova at 10 kpc from the background signal, whether through peak or event–level analyses. Future improvements, such as optimizing the event selection criteria and extending the analysis to additional models, could further enhance the detector's capabilities and

potentially increase the significance of the observed signals.

5.3.2 False Alarm Rates and Combined Significances

It was demonstrated that the simulated supernova signals from all models at 10 kpc can be successfully identified within the background data. However, despite the application of data selection criteria, supernova burst identification methods, and the rejection of false positives using a CNN model, a residual number of false positive triggers persists. In peak–level analysis, 174 false positives remain after the final CNN-model selection. At the event level, only one of these 174 false positives remains, while three additional false positives were found exclusively in the event–level data.

Although there is potential for further refinement in background mitigation and trigger algorithm optimization, anticipated FAR, based on current background observations, can be evaluated across various distances. In addition, alternative strategies to reduce false alarms, such as simultaneous analyses of the three detectors (TPC, muon veto, and neutron veto), warrant exploration.

While the extended LRT provides a method for quantifying the significance of individual signals in the data, combining these significances across different detectors and analysis levels requires a unified statistical framework. To address this, we transition to an approach based on the FAR. This perspective not only ensures consistency in combining significances from different analysis methods, such as peak and event–level data, but also facilitates the integration of expected observations from the muon and neutron veto detectors.

The anticipated FAR for any detector can be estimated based on the observed background signal rates. In scenarios with low counting statistics, if the expected background signal rate is given by λ, the probability of observing k events can be described by the Poisson distribution. This probability, denoted as $P(\geq k, \lambda)$, represents the likelihood of observing at least k events when the expected mean number of events is λ.

In scenarios where the experiment is repeated multiple times, such as counting the signal every second and expecting a consistent background rate around the mean, a correction factor known as the *trial factor* (N) must be introduced. The trial factor accounts for the number of independent trials conducted during the observation period. The frequency with which the background process mimics the observed signal, known as the *imitation frequency* (F_{im}), (or FAR), can be calculated as the probability $P(\geq k-1, \lambda)$ multiplied by the trial factor N:

$$F_{im} = P(\geq k-1, \lambda) \cdot N. \tag{5.12}$$

The use of $k-1$ in this equation is due to the nature of the counting process, which typically begins once the first data point is observed. Therefore, the probability calculation is based on observing at least $k-1$ events.

The inverse of the imitation frequency, as a function of the total observed signal counts, is shown in Figure 5.42 for both peak–level and event–level analyses. The average background rates in each case were estimated from the entire SR1 dataset, after applying data selection and excluding regions where the CNN model rejected a trigger. Although local background rates may vary slightly from these estimates in the peak–level analysis, this assumption provides a reliable approximation of the imitation frequency.

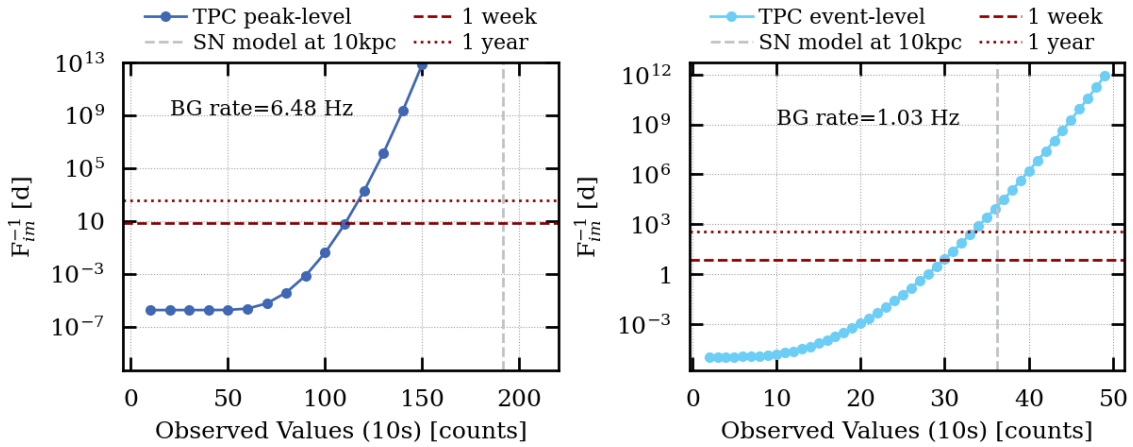

Figure 5.42: The inverse of the imitation frequency for the observed counts in both peak- and event-level analysis for a 10 seconds rolling window. The dashed and dotted red lines show the 1 week and 1 day marks. The shown data points for each panel represent the time it would take for the given detector background to imitate the observed counts. The expected total counts for the benchmark Bollig model with 27 $M_{\odot}$ progenitor is also shown with the vertical gray dashed line.

In Figure 5.42, the dotted and dashed lines represent the time intervals of one year and one week, respectively. The inverse imitation frequency at a given observed count represents how long it would take, on average, for the background to mimic a signal with that many observed counts, based solely on upward Poisson fluctuations at the specified mean rate.

For example, as shown in the right panel of Figure 5.42, observing 30 counts in a 10-second window, when the average background rate is 1.03 Hz, could occur approximately once per week without any additional signal. This result arises because of the large number of trials in a day, where the number of trial windows is given by the total seconds in a day multiplied by the mean background rate. Thus, even a seemingly high count could be due to background fluctuations alone, depending on the rate and number of trials. However, for considerably larger observed counts, the required imitation time is greater than a few decades, indicating that the observed signal almost certainly contains an additional signal source or a systematic effect.

It is important to note that this study employs a 10-second window to estimate the total expected background counts as well as the total counts within the same window size around the supernova signal. This implies that, with an optimized time window, the observed signal counts in the vicinity of the supernova region would yield a higher signal-to-background ratio and align with lower imitation frequencies. Furthermore, it should be noted that while the study presumes that the background rates adhere to a Poisson distribution around the observed mean, this may not consistently be the case due to the varying detection conditions.

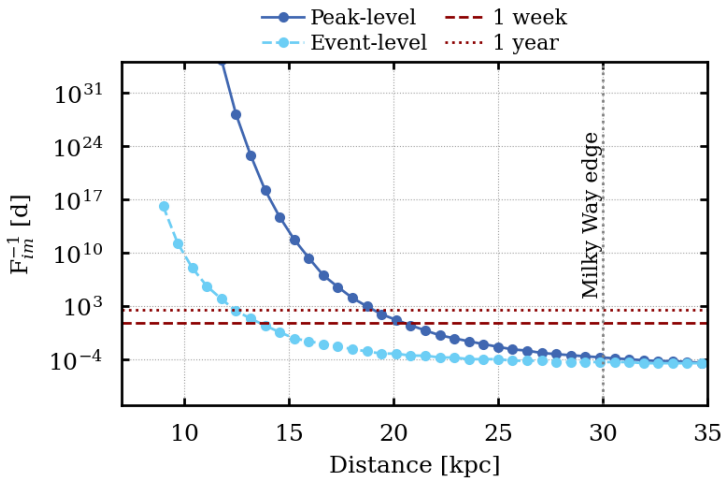

Figure 5.43: The inverse imitation frequency as a function of supernova distance, shown for both peak–level and event–level counts. The observed counts are averaged over all injected supernova signals in the background within a 10-second window, which is positioned around each trigger to maximize the counts within the window. As the distance to the supernova increases, the expected counts decrease, making it more likely for the background to mimic the observed signal at higher distances.

As the expected number of neutrino interactions decreases with increasing distance of the supernova, the imitation frequency can be recalculated for different distances. Using the Bollig reference model with a 27 $M_{\odot}$ progenitor, the observed counts are scaled accordingly for supernovae at varying distances. The inverse of the imitation frequency as a function of distance is shown in Figure 5.43.

This figure should be interpreted as follows: for a supernova located at 20 kpc, the counts it would generate correspond to an imitation frequency of approximately once per week at the nominal background rate. The closer the supernova, the higher the expected counts, leading to longer time intervals before the background could mimic such counts. This implies that if the desired FAR is not more than once per week, supernovae located beyond 20 kpc become increasingly difficult to distinguish from the background at the peak–level, as the background itself can produce similar counts more frequently than once per week. This underlines the expected fact that for distant supernovae, distinguishing true signals from background fluctuations becomes more and more challenging.

For a specified background rate, there exists a threshold for the observed count that determines the false alarm limit. This study establishes the limit to one false alarm per week following the suggestions of SNEWS. Such rates as demonstrated in Chapter 3 correspond to very low false coincidences to issue a SNEWS alert to the network. In the context of peak–level analysis, this threshold is equivalent to observing 110 counts within a 10-second interval. Exceeding this count within the specified time frame implies the potential presence of an additional signal, such as a supernova, beyond the ambient background noise.

A probability can be defined to quantify whether the observed signal is more likely from an additional source rather than Poisson fluctuations in the background. This probability increases as the observed count exceeds the set threshold. In contrast, if the total observed count N_{tot} falls below the threshold, the probability of detecting a true signal decreases.

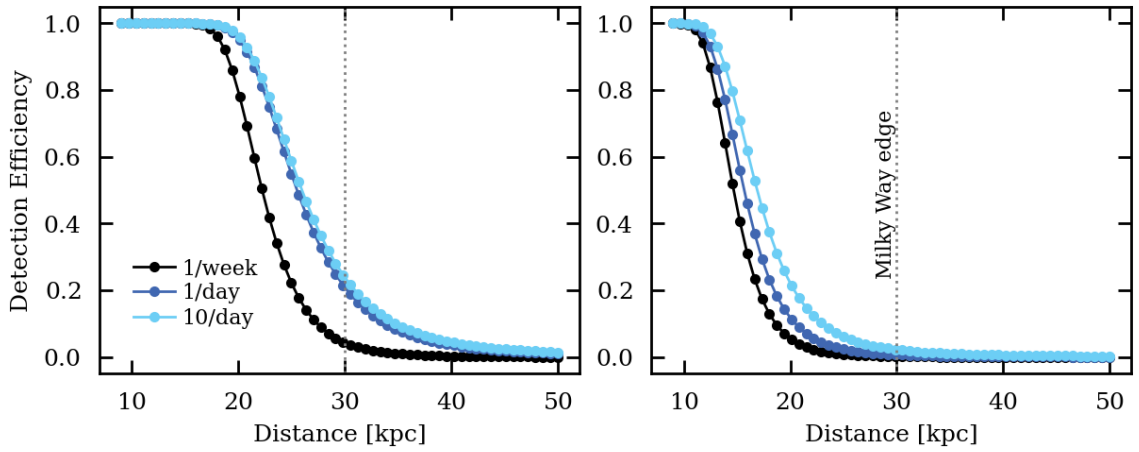

Figure 5.44: Detection efficiencies for various FAR, once per week, once per day, and ten times per day, based on observed signal counts in 10-second windows. The left panel represents the peak–level analysis, while the right panel shows the event–level analysis. Expected signal counts are derived from the Bollig model with a 27 $M_{\odot}$ progenitor and scaled for different distances. A higher FAR (e.g., ten times per day) results in a lower detection threshold, allowing supernovae at greater distances with fewer counts to trigger detection, thus increasing the probability of detection. In contrast, a stricter FAR (e.g., once per week) requires a higher minimum count, making detection more likely only for closer supernovae.

The probability of detecting a signal greater than the threshold count c (in this example, 110 counts) is given by $P(c \geq N_{\text{tot}})$. This probability provides a measure of confidence in distinguishing a supernova signal from background fluctuations.

The detection efficiencies for different FAR, for example once a week, once a day, and ten times a day, are illustrated in Figure 5.44. The analysis is presented for data from both the peak level (left panel) and the event level (right panel). The expected signal counts for each case were derived from the Bollig model with a 27 $M_{\odot}$ progenitor and scaled according to the distance within the 10-second observation windows.

A higher FAR (e.g., ten times per day) implies a lower detection threshold, meaning that even supernovae at greater distances, which generate fewer signal counts, could still trigger an alert. This results in a higher probability of detection. Conversely, a stricter FAR (e.g. once per week) requires a higher minimum number of observed counts, limiting detection to supernovae that are closer and produce stronger signals.

This suggests that the detection of supernovae at greater distances is possible by lowering the required total observed counts and permitting a higher FAR, such as ten times per day. However, setting the FAR at such a high threshold would result in an impractically frequent number of alarms, many of which may not correspond to true supernova signals. To address this, a more effective strategy is to incorporate a coincidence analysis across multiple detectors.

Since event–level data are a higher level of processing derived from peak–level signals, they are not treated as completely independent signals. However, as discussed in Chapter 2, the XENONnT detector is equipped with two outer veto detectors, muon veto and neutron veto, which can also detect neutrino interactions from supernovae via inverse beta–decay, IBD. In

this analysis, the estimated background event rates after data selection for the muon and neutron veto detectors derived in [110] are used.

By forming a coincidence between the three available detectors, it is possible to significantly reduce the FAR while still maintaining high sensitivity to distant supernova events. This approach is analogous to the r–fold coincidence technique discussed in Chapter 3, where the combined FAR F_{comb} is calculated by considering simultaneous detections across multiple detectors. The combined rate for an r–fold coincidence between n detectors can be expressed as:

$$F_{\text{comb}} = (C(n,r)+1) \times F_{d1} \times F_{d2} \times \ldots \times F_{d_r} \times \delta t^{(r-1)} \tag{5.13}$$

$$= (C(3,3)+1) \times F_{TPC} \times F_{MV} \times F_{NV} \times \delta t^2 \tag{5.14}$$

where $C(n,r)$ is the binomial coefficient representing the number of ways to choose r out of n detectors, $F_{d1}, F_{d2}, \ldots, F_{d_r}$ are the individual imitation frequencies for each detector, and δt is the coincidence time window. By combining the information from all three detectors, the strategy greatly enhances the overall detection reliability and reduces the chance of false alarms while still preserving the ability to detect supernova signals from greater distances. For example, by anticipating one alarm per week in each of these detectors and requiring that all alarms coincide within a 10-second window, the combined FAR would be reduced to once about every 351 million years. Even for an alarm rate of once a day, the coinciding false alarm in a 10-second window is expected to occur once every hundred thousand years assuming that upward fluctuations are purely of statistical origin.

For this study, the values for the expected background rates in each of the two veto detectors and the expected average supernova neutrino signal in a 10-second window were taken from [110] and are listed in Table 5.6. Consequently, the inverse of the imitation frequencies for the veto detectors was computed and shown in Figure 5.45.

	μ veto	n veto
BG rates	3 Hz	0.05 Hz
SN counts	70	8

Table 5.6: Background rates and expected average supernova counts within a 10-second window for muon (μ) and neutron (n) vetoes. The numbers are from [110].

Ultimately, the combined FARs are shown in Figure 5.46. It can be observed that the inverse of the imitation frequency decreases with increasing distance for all analysis, indicating that the time needed for the background to statistically fluctuate upward and imitate the expected signal count at the given distance becomes shorter. Until about distances of 23 kpc, combined analysis has a lower FAR, beyond which the FAR of the TPC dominates.

At large distances the expected counts from the reference supernova model become increasingly negligible, meaning that most of the observed counts are from the background. In this scenario, where the expected signal is extremely low due to the great distance, it becomes

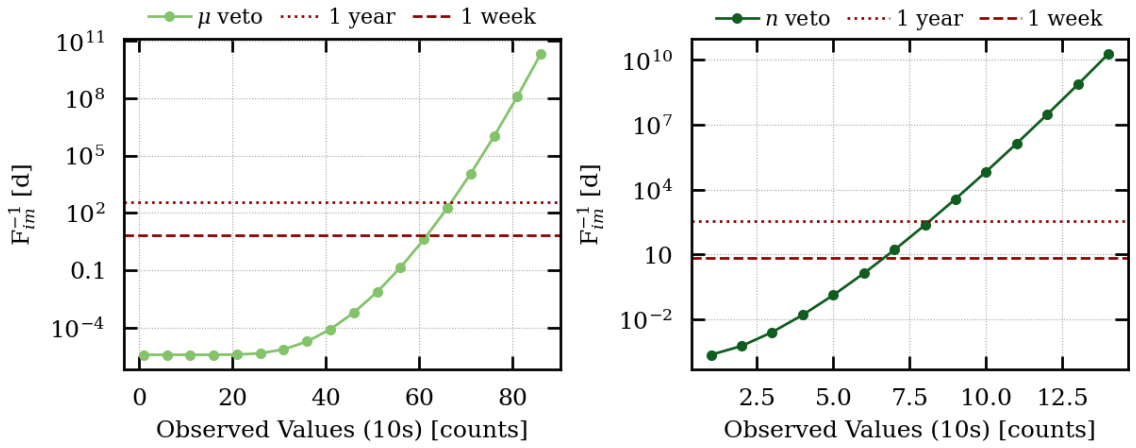

Figure 5.45: The inverse of the imitation frequency for different observed values in a 10 seconds rolling windows for both muon and neutron vetoes. The dashed and dotted red line shows the 1 week and 1 day marks. The shown data points for each detector represent the time it would take for the given detector background to have a statistical fluctuation up to the observed counts.

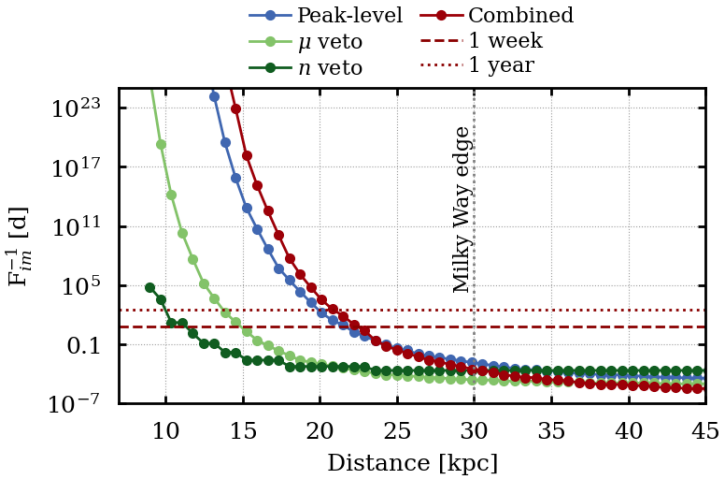

Figure 5.46: The FAR as a function of distance shown for TPC peak–level data, muon and neutron veto data, and the combined analysis. The combined analysis shows a reduction in the FAR until about 23 kpc compared to individual detectors. At larger distances the FAR of the TPC dominates, as the expected counts in the two veto detectors at these distances become lower than the mean background rates.

more probable for background fluctuations to mimic a signal. Background events, which naturally fluctuate, are more likely to form clusters with at least the average number of counts, leading to an increase in the FAR. As a result, the FAR tends to flatten at these greater distances. When the expected FAR for an individual detector falls below the duration of the coincidence window, multiple false alarms can occur within the same coincidence window. This causes the combined analysis curve to fall below the individual detector curves, as overlapping false alarms reduce the effectiveness of coincidence–based suppression.

This study demonstrates that with a requirement of one false alarm per week, all three detectors can independently identify the benchmark supernova model at distances up to 10 kpc, and only the TPC can detect it up to nearly 20 kpc with a high confidence when

the expected FAR is set to once a week. A combined analysis, which requires a coincidence between the three detectors within a 10-second window, can extend the detection range and effectively reduce the frequency of false alarms caused by background fluctuations. The distances at which each analysis method expects counts that are comparable to the counts expected from the background once a week, once a day, and ten times a day are shown in Table 5.7. The values show that when a combined analysis is performed, the supernova can be confidently identified up to 22.5 kpc when the expected FAR is set to once per week.

	TPC	μ veto	n veto	Combined
1/week	21.7 kpc	15.1 kpc	11.3 kpc	22.5 kpc
1/day	22.7 kpc	15.7 kpc	12.2 kpc	23.4 kpc
10/day	21.5 kpc	14.9 kpc	10.9 kpc	22.3 kpc

Table 5.7: The expected distances at which each detector and the combined analysis generate given FARs purely from the background. For distances closer than stated, the analysis method would identify the observed signal as a supernova. For greater distances, the observed counts may be confused with background fluctuations.

Furthermore, the signal might be missed on either of the detectors for several reasons, such as calibration or maintenance. In these cases, a coincidence between the remaining two detectors can still be formed. This is shown in Figure 5.47.

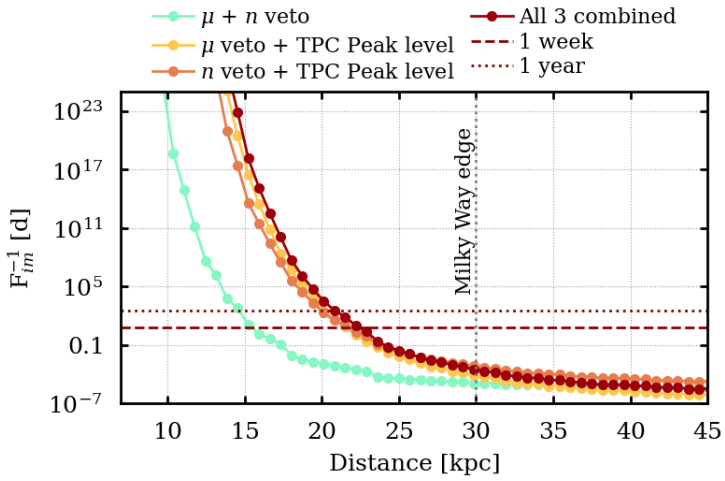

Figure 5.47: Imitation frequencies for 2- and 3–fold combination of detector systems.

It is crucial to emphasize that the TPC remains central to supernova detection due to its sensitivity to the expected burst of neutrino interactions. The analysis conducted at the peak–level archival data reveals a higher FAR than predicted by the average background rates in the TPC. However, event–level analysis demonstrates a significantly lower FAR, with only four irreducible false alarms observed over 85 days of data. This contrast underscores the inherent trade-offs between peak–level and event–level analyses: the former's intentionally relaxed trigger parameters, designed to maximize detection efficiency, and the latter's more stringent criteria for enhanced precision. Temporary detector conditions, capable of generating nonstandard signals that mimic a supernova, also contribute to this discrepancy. While there is room for refining the data selection criteria and enhancing the discrimination

capabilities of the CNN model, the overall results indicate that even when one detector exhibits a higher false positive rate, the probability of coinciding false positives across multiple detectors remains low at close distances.

The identification of increased signal rates within the TPC represents an improtant step in supernova detection. As demonstrated, TPC data alone can provide a statistically significant signal that stands out against background fluctuations, offering a robust mechanism for identifying supernova neutrinos. However, combining the observed signals from other detectors, such as the muon and neutron vetoes, or integrating results from separate analyses of TPC data products and different processing levels, can further enhance detection significance. This approach not only strengthens confidence in the detection of supernova events but also extends the reach of the experiment to observe lower-count supernova signals with comparable significances, solidifying the TPC's position as a crucial instrument in the field of multimessenger astrophysics.

6

Dark Matter Detectors and Supernova Neutrinos: XENONnT and Beyond

"The important thing is not to stop questioning. Curiosity has its own reason for existing."
–Albert Einstein

Contents

The previous chapters discussed the physics of supernovae, describing how neutrinos from these cataclysmic stellar explosions travel across the cosmos to be detected on Earth. In particular, the focus was on how direct DM detectors like XENONnT are sensitive to these neutrino signals through CEνNS. Detailed signal analysis methods within the XENONnT framework were presented, including the development of a robust software trigger designed to promptly identify potential supernova signals in the data stream and report them to SNEWS.

This chapter builds on these findings, exploring avenues for improvement, alternative scenarios, and the exciting prospects of future DM experiments. In Section 6.1, the multimessenger role of the XENONnT experiment in the SNEWS network is discussed. Possible extensions and enhancements of the current study, including advanced signal selection and modeling techniques, are discussed in Section 6.2. In Section 6.3, the potential of next-generation DM experiments to contribute to supernova neutrino detection is highlighted.

6.1 XENONnT as a Multimessenger Player

As previously discussed, the imminent Galactic CCSN will be pivotal for multimessenger astrophysics, providing signals such as neutrinos, gravitational waves, and electromagnetic radiation. Neutrinos will reach Earth first, well ahead of electromagnetic signals, allowing the formation of an early alert. In preparation of capturing this cosmic event, several neutrino observatories and DM detectors, including XENONnT, have been developing analysis frameworks.

As established in Chapters 4 and 5, XENONnT can contribute to the formation of coincidence alerts within the SNEWS network by providing observation timings for the detected CEνNS signal from a CCSN. This capability underscores the multifunctionality of the experiment. As SNEWS is evolving from SNEWS 1.0 to SNEWS 2.0, the XENONnT detector is fully integrated into the network and ready to receive these alerts. Soon, it will also be able to provide prompt automated observations of potential supernovae and take part in the global effort of the early detection of a supernova.

The SNEWS network integration involves three aspects: reacting to SNEWS alerts, sending periodic heartbeats to maintain server connection, and sending an observation message when a potential supernova signal is detected. The details of these three aspects are as follows:

- ***Reacting to alerts:*** XENON Collaboration has established a protocol for the actions required when an alert is received. As part of this thesis, a wrapper software package, `supernova_communications`, was developed to interface with `snews_pt` and `snews_cs`. This package manages credentials and establishes connections to SNEWS alert topics, enabling the monitoring and retrieval of supernova alerts. Upon receiving an alert, it generates an internal notification and sends messages to shifters on-site and relevant online channels. The shifters on site should label and promptly process the data at the time of the alert for further analysis.

- ***Sending heartbeats:*** The XENONnT experiment is the first DM detector to send frequent signals (*heartbeats*) to SNEWS 2.0. These heartbeats help evaluate system latency and identify network and Internet stability issues during development. For this, we first identify the operational mode of the experiment; modes such as calibration or maintenance indicate inactivity, while others indicate activity. Then, this information is automatically converted into a heartbeat message and sent to SNEWS every 10 minutes.

- ***Sending observation messages:*** XENONnT already uses an *online data monitor* for data quality checks. This tool is being expanded to include automated, prompt supernova identification. The new functionality will simultaneously search event- and peak-level data for supernova neutrino interactions following the framework described in Chapter 4 . If detected, an observation message will be sent to the SNEWS network using the `supernova_communication` package. Based on this study, the recommended strategy is to monitor the TPC data at both high and low levels, together with the data from the outer veto detectors. Any coincidence between peak- or event-level data and an outer veto detector signals high confidence and triggers an automatic message to SNEWS. If only one detector triggers while others remain operational, the tool calculates local significances to assess whether the signal is due to background fluctuations. A message is sent to SNEWS only if the background-only hypothesis is rejected. Every message to SNEWS also notifies shifters on site and the XENONnT multimessenger team by text, email, and Slack. False observations can be retracted from the SNEWS server.

6.2 Possible Improvements of the Study with XENONnT

The analysis of supernova neutrino interactions in XENONnT was covered in Chapter 5. Although the analysis framework presented here demonstrates efficacy in examining models in current settings, there is still potential for improvement and additional discussion. On the basis of experience gained throughout this thesis, this section presents several recommendations for potential future extensions of the study, along with some preliminary assessments for some of these extensions.

6.2.1 Determining the Onset of the Signal

Identifying the onset of the signal, t^*, is important, especially for multilateration purposes [135]. The local significances of each observed signal around the triggered time were discussed in Section 5.3, where the asymptotic formula for the likelihood-based test statistic was used for the calculation. Within [153], the authors propose alternative models that aim to improve the quantification of significance while simultaneously offering a robust method to identify the onset of the signal. The method uses hypothesis testing to identify the signal and its

onset. For hypothesis testing, a test statistic ℓ was computed on the observed data. The significance level of observing the test statistic under a specified hypothesis H is given as

$$p(\ell) = \int_{\ell}^{\inf} P(\ell'|H)\mathrm{d}\ell', \tag{6.1}$$

and the corresponding *z–score*

$$z(\ell) = \Phi^{-1}(1 - p(\ell)), \tag{6.2}$$

where Φ is the cumulative normal distribution, indicating the number of Gaussian standard deviations from the mean expected signal.

There are two proposed test statistics; simple number count in a given time bin (Counting Analysis, CA), and Shape Analysis (SA), which is model dependent. The test statistic for the CA, ℓ_{CA}, is defined as the total counts, n, as

$$\ell_{CA}(t^*, \Delta t, t_i) = n \equiv \sum_i w(t_i - t^*, \Delta t), \tag{6.3}$$

where w is a window function equal to 1 within the time window Δt and 0 outside, here i indexes each detected event in the dataset. The test statistic for the SA is defined through a log likehood ratio test between the alternative (H_{alt}) and the null hypothesis (H_0) as

$$\ell_{SA}(t^*, t_i) = \log \frac{P(t_i|H_{\mathrm{alt}})}{P(t_i|H_0)} = \sum_i \log\left(1 + \frac{S(t_i - t^*)}{B(t_i)}\right) \tag{6.4}$$

where S and B are signal and background counts at a given time t_i respectively.

To evaluate each method, a dedicated software [154] was used. Initially, an average representation of the supernova models was constructed by employing all simulated signals derived from the respective model. Since each burst simulation was drawn from the selected supernova model incorporating a Poisson fluctuation, and the detector effects were propagated in the `fuse` simulation, this representation accounts for both statistical and systematic fluctuations. Two examples of average model representations are shown in Figure 6.1 for the Bollig model with a 27 $M_\odot$ progenitor, and for comparison another from the Fornax model with a 26.99 $M_\odot$ progenitor. Note that for both models, a signal distance of 10 kpc are used.

Subsequently, the test statistics ℓ_{CA} and ℓ_{SA} are evaluated in proximity to the detected time of the signal–peak. The left panel of Figure 6.2 displays the observed rates along with the expected background-only hypothesis and the alternative hypothesis, which anticipates the average model counts in addition to the expected average background counts. The right panel of Figure 6.2 illustrates the evolution of the test statistic (eqs. (6.3) and (6.4)) and the z–score (eq. (6.2)) as a function of time.

When using SA, if the observed signal shape matches the underlying model of the SA, the significance will improve greatly compared to that of CA (lower right panel of Figure 6.2). However, this approach introduces model dependence. The scope of such a study can be broadened by developing a generic supernova model to explore the evolution of the test statistic across various models within this framework. CA, in contrast, is similar to

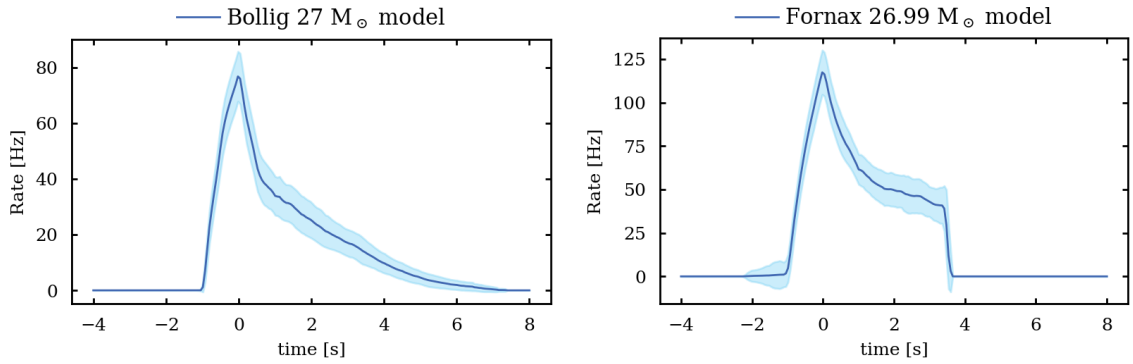

Figure 6.1: The average model inferred from the simulated data after all data selection is applied. The error bands represent the 1σ standard deviation around the average based on the all the simulations of a given model.

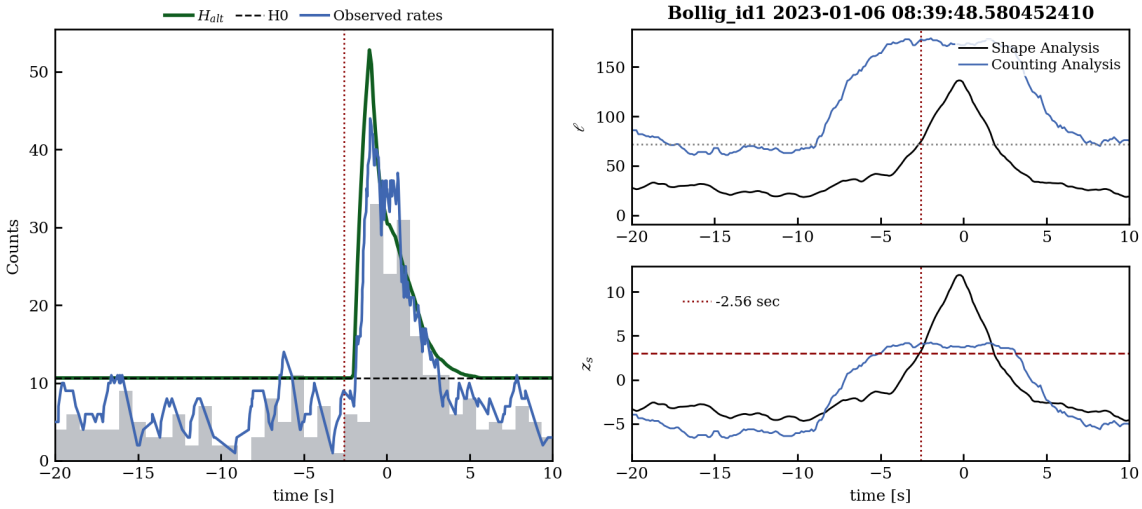

Figure 6.2: *Left*: The observed signal counts shown as grey bins and rate evolutions (solid blue line) along with the expected rate evolution from the null (dashed green line) and alternative hypothesis (solid green line) are overlaid on top. Rates are not shown on the y–axis explicitly. *Right*: The test statistic and the z-score as a function of time for both CA and SA. The Bollig model with 11 $M_{\odot}$ progenitor is used in this example. The dashed horizontal line and the intersecting vertical line represent a predetermined threshold of 3σ z-score and the associated time value at the intersection of the SA z–score.

the asymptotic likelihood ratio tests discussed in previous chapters but tends to suffer from upward fluctuations in response to increases in detector background rates. Moreover, the software presently exhibits certain limitations, particularly in the form of numerical computation errors occurring at high levels of significance. However, both methodologies provide a robust and reproducible mechanism for the quantification of signal significance. By establishing a predefined z–score threshold, a coherent understanding of the onset time of the signal can be achieved. This approach can be further refined in future studies by applying corresponding SA to improve the reliability of signal identification and onset determination.

6.2.2 Trigger Performance at Different Distances

This work focuses mainly on supernova signals at 10 kpc and builds a framework for the identification of the signal at this distance. While 10 kpc is a likely distance for the next galactic CCSN, it is important to expand the study by evaluating the detection capabilities at farther distances as well. The trigger algorithm is based on the rate of count evolution rather than on the signal amplitude alone. Although the number of interactions decreases with distance, detection can still be achieved through rate evolution curves, allowing the system to trigger supernova alerts even for more distant sources.

To test this, the observed counts in the 2194 simulated supernova bursts from all the models used in this study were scaled to mimic observations from further distances, adjusting for the expected decrease in event counts. The same trigger algorithm developed for 10 kpc was applied to assess detection performance at these larger distances. The results are presented in Figure 6.3 as a bar graph for each model and in Figure 6.4 as a percentage of the signals recovered as a function of different distances. At distances up to 20 kpc, all Nakazato black hole models (id13, id14, and id15, corresponding to the 30 $M_{\odot}$ progenitor with LS220, shen, and togashi EOS, respectively) are identified with 100% efficiency. At 20 kpc, Nakazato id13 misses one injected supernova burst signal, while id14 and id15 maintain 100% efficiency. At 25 kpc, Nakazato id15 identifies all but one signal, while id14 achieves the same level of efficiency at 30 kpc. The identification of other models gradually decreases.

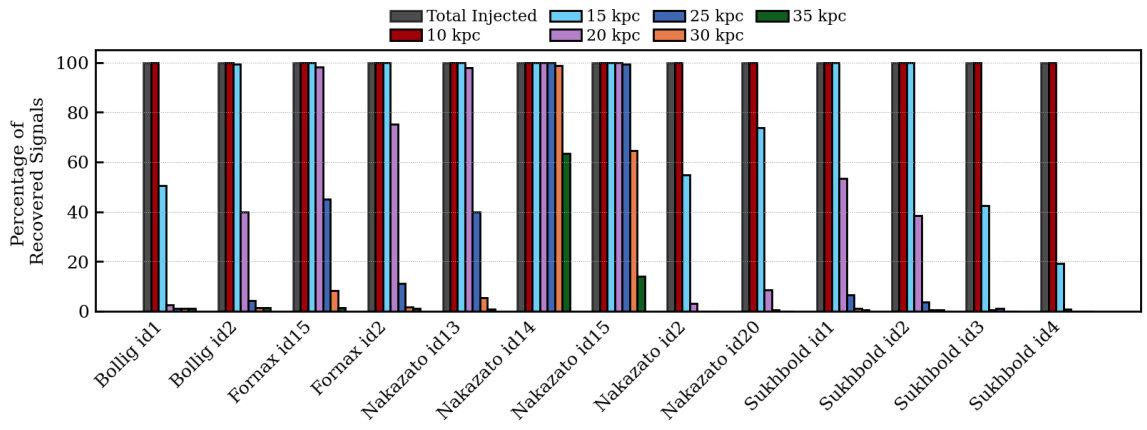

Figure 6.3: The percentage of recovered supernova signals for each model considered in this study. Different bars represents the signal at different distances. All simulated models are identified when they are at 10 kpc and the number of identification decreases with increasing distance.

The performance of the trigger algorithm is influenced by three key components: signal selection, trigger parameters (such as prominence), and the tuning of the CNN classification model. Variations in signal selection can alter the shape of the final rate evolution curves, potentially causing the current trigger algorithm to perform suboptimally. The trigger parameters, optimized for supernova signals at 10 kpc, were calibrated to detect signal amplitudes and widths characteristic of such a distance. Adjusting these parameters could improve performance for signals at other distances. Lastly, the CNN classification model, trained to identify true signal–peaks based on 10 kpc simulations, may not generalize

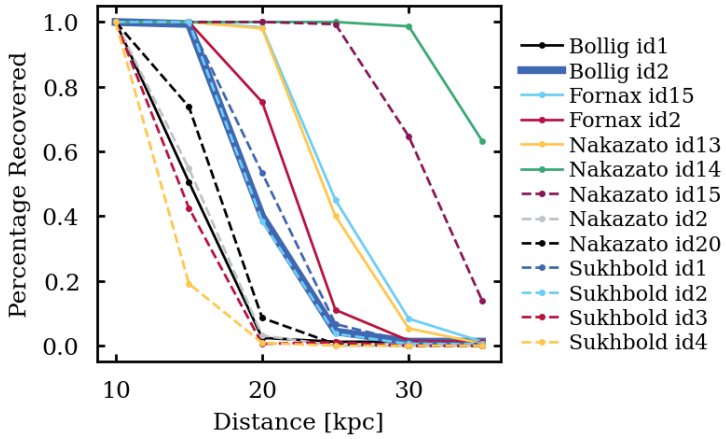

Figure 6.4: The percentage of triggered supernova for a given model as a function of distance. Most of the Nakazato black hole models (id13, id4, and id15) can be identified even at greater distances due to their unique signatures where the signal rate increase is sharp.

effectively to signals from greater distances. A more flexible signal selection, paired with adapted trigger parameters and a CNN model designed to handle a broader range of distances, could significantly enhance trigger performance for supernovae at larger distances.

6.2.3 Further Remarks

Effects of Higher Drift Field: In XENONnT, the electric drift and extraction fields are generated using three parallel electrodes: the cathode, gate, and anode. Initially, the drift field was set at approximately 100 V/cm but was later reduced to 23 V/cm due to a short circuit [93]. The strength of the electric field plays a crucial role in predicting both the light and charge yields in nuclear recoil interactions, as illustrated in Figure 6.5 using NEST simulations. Furthermore, the drift field affects the likelihood of electrons surviving as they travel through the LXe medium, due to potential capture by electronegative impurities dispersed in the LXe. Because of the high xenon purity of XENONnT, electron lifetimes (> 10 ms) are longer than the maximum drift time of 2.2 ms [93], minimizing electron losses during detection.

Although the lower drift field has had minimal impact on electron survival rates in XENONnT, higher drift fields could improve charge yield, potentially leading to improved signal detection. Furthermore, a higher drift field also improves the NR/ER discrimination, especially for low–energy nuclear recoils. Understanding these effects is helpful for future DM experiments, which may operate at different field strengths. Although this analysis focuses on the current setup, simulating the detector response under higher drift fields could provide valuable insights for future upgrades and the potential replacement of electrodes in XENONnT.

Improved Signal Selection: The exact supernova model is unknown, and the analysis should be robust against variations in models. The identification of a potential supernova signature in the XENONnT background data starts with signal selection. The current study uses the complete SR1 dataset in conjunction with 2194 simulated CCSN burst signals

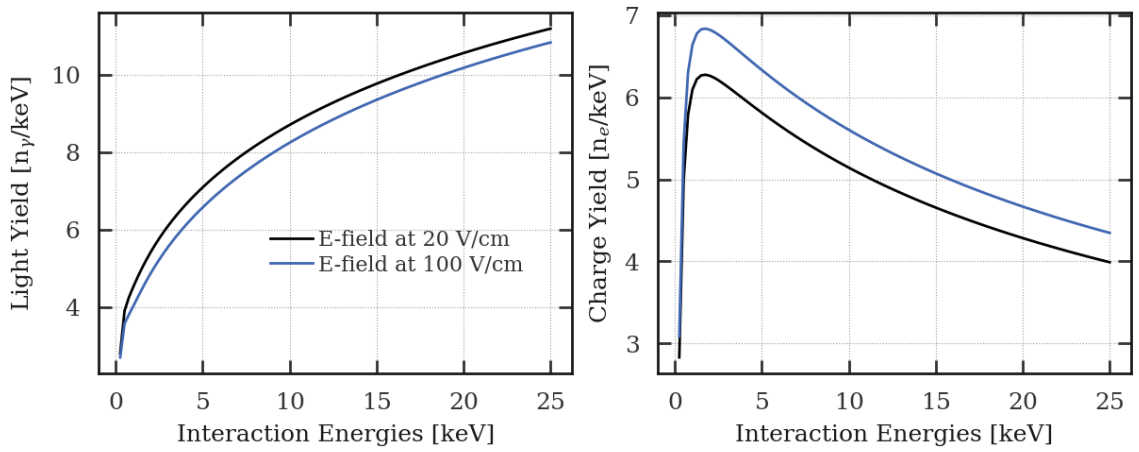

Figure 6.5: The expected light and charge yield under different electric fields as predicted by NEST. The number of quanta shown as a function of recoil energy.

derived from 13 progenitors across 4 supernova models, all at a distance of 10 kpc. While the signal selection is optimized on the salted data with Bollig 2016 model (27 $M_\odot$), the subsequent optimization of supernova burst detection, incorporates the full set of simulated data to ensure the analysis remains robust against variations in models. The goal is to isolate supernova signatures more efficiently by employing both conventional methods and decision tree-based techniques.

Although recoil energy spectra across models are similar, the variability in time profiles suggests that future improvements should consider minimizing model-specific time dependencies. Similarly, adjusting the time-proximity for more distant supernovae, which is expected to result in fewer interactions, could further enhance the selection process.

Improvements to the CNN-based identification system would reduce reliance on strict background suppression, provided that the remaining background signals remain distinguishable from supernova light curves. The training sample can be expanded to improve the discrimination power of the model.

Supernova Search During Calibration: An important extension to this study could involve analyzing the detector's response during calibration runs. As long as the rate evolution of the background and supernova signal can be modeled, an additional trigger can be built. Although these calibration periods less than 10% of the primary data-taking phases, conducting a separate supernova search during these runs could increase the live time of the experiment for the supernova search. This approach would reduce the likelihood of missing a potential supernova event by extending the periods in which signals are actively monitored.

Additionally, the existing framework could be modified to prioritize signals from the veto detectors during calibration runs. In this case, if a supernova-like signal is detected by the veto detectors, the system could immediately trigger a warning to SNEWS, even if no coincident signal is observed in the TPC. This would maximize the detector's utility across all operational phases.

6.3 Supernova Physics with Next Generation DM Experiments

The sensitivity of the XENONnT experiment and similar size DM experiments to neutrinos originating from CCSN has been demonstrated in the previous chapters. The experiment has a high chance of detecting the next CCSN event within the Milky Way. Furthermore, the current analysis framework employed by XENONnT for the detection of galactic supernovae was optimized for a supernova at 10 kpc. Nonetheless, the same framework can still identify supernova signatures from most models (including the benchmark Bollig model with 27 $M_{\odot}$ progenitor) more than 40% of the time up to 20 kpc, and in cases of failed supernovae more than 60% of the time up to 30 kpc.

Next–generation dual–phase noble liquid DM experiments are projected to use a target mass of about 50 tons [74]. These experiments are expected to observe supernova signals with statistics similar to those of existing neutrino observatories during a supernova event [47].

Although prompt detection of the forthcoming galactic CCSN is essential for initiating early alerts and enabling multimessenger observations, it simply marks the beginning of a deeper scientific exploration. Upon identification of neutrino interactions emanating from the supernova, an extensive investigation can ensue. As discussed in [47], the improved statistical data from next–generation DM detectors will enable detailed reconstruction of the neutrino light curve, allowing for a more precise evaluation of theoretical models and the rejection of those that are inconsistent with observed supernova behavior. Also, as additional data become available, such as distance estimates from electromagnetic observatories and flavor-specific data from other neutrino observatories that are otherwise degenerate, more detailed investigations can be conducted. For example, distance information can assist in resolving ambiguities among different progenitor models, whereas single-flavor data, when combined with the total energy emitted in neutrinos, can be employed to constrain neutrino properties.

6.3.1 Expected Number of Interactions

Assuming that the forthcoming DM experiments will employ 50 tons of liquid xenon in its active volume, it is possible to compute the expected number of interactions. For the same Bollig benchmark model with 27 $M_{\odot}$ progenitor, the differential interaction rates for each flavor, along with the total numbers, are presented in Figure 6.6 as functions of time (left) and recoil energy (right). Furthermore, in Figure 6.7, the integrated total number of interactions of the supernova burst is shown for each progenitor considered in this study.

The number of detected events will be influenced by the detection efficiencies and background suppression in future experiments. However, for a supernova at 10 kpc, it is generally expected that around a thousand interactions will be observed. This is comparable to the detection numbers projected by current neutrino observatories (a comparison of estimated total interactions can be found in Table 4 of [118]).

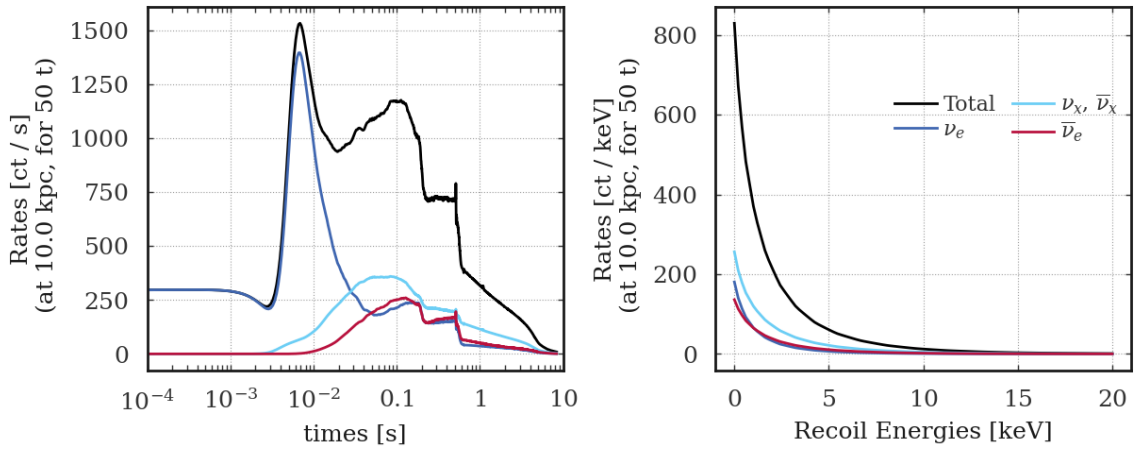

Figure 6.6: The differential interaction rates of the Bollig model with a 27 $M_\odot$ progenitor as a function of time (*left*), and as a function of recoil energy (*right*), shown for each neutrino flavor and the total. The rates are scaled for 50 tons of target mass and a supernova progenitor at 10 kpc.

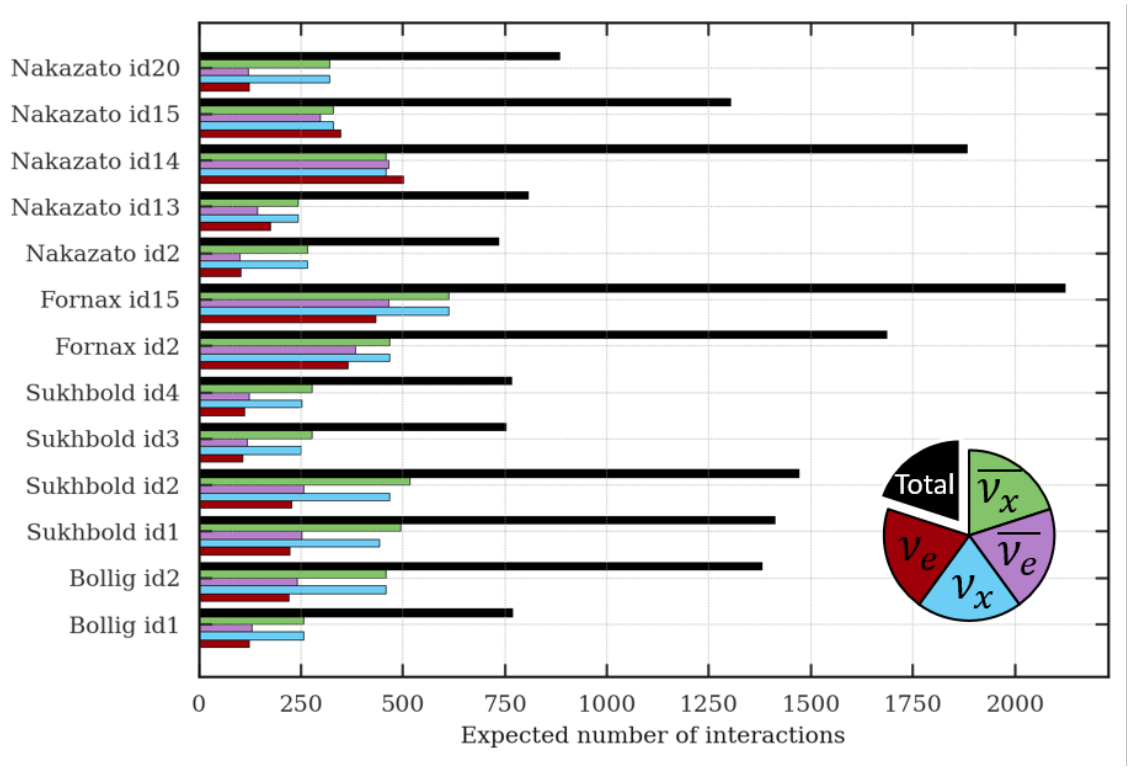

Figure 6.7: Total expected neutrino interactions from a supernova at 10 kpc for each model considered in this study for a next generation 50 ton target. The individual total counts for each neutrino flavor are also shown. For the identifier id, see Table 4.1.

The models with the highest and lowest expected number of interactions were scaled for different distances and shown in Figure 6.8, indicating that at least 100 interactions are expected within the Milky Way and a non-zero number of interactions are expected from a supernova in the large (~ 50 kpc) and small (~ 60 kpc) magellanic clouds.

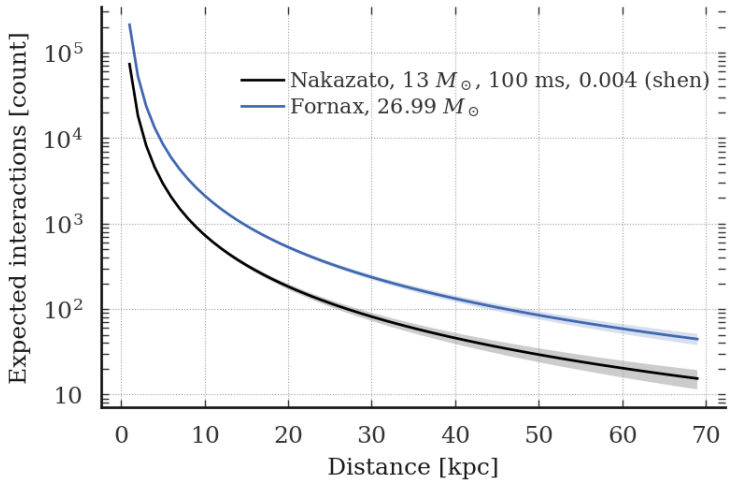

Figure 6.8: Expected number of interactions for a next-generation DM experiment with 50 tons of liquid xenon. The models predicting the highest (Fornax model with 26.99 $M_\odot$ progenitor, Fornax id15) and lowest (Nakazato model with 13 $M_\odot$ progenitor, Nakazato id2) expected counts are shown as a function of distance.

6.3.2 Likelihood Ratio Tests for Model Discrimination

With sufficient statistical data, it becomes possible to evaluate the evolution of an observed signal compared to theoretical models. In the following, a simplified analysis is presented. To exclude a model based on the observed data, the expected count distribution in different time bins was computed and scaled. The total observed count, N, was fixed, and the models were scaled to match the total predicted counts at different distances.

This analysis was performed using binned data in two approaches: linear and logarithmic binning. Linear binning ensures uniform intervals (here, chosen as 100 ms), which are suited for simpler a comparison, while logarithmic binning is ideal for capturing the finer details in the signal shape, especially during early phases. Figure 6.9 illustrates the expected counts for two Bollig models with progenitor masses of 11 $M_\odot$ and 27 $M_\odot$, using both binning methods. To assess the ability to discriminate between these models, toy Monte Carlo simulations were generated. For each simulation, 100,000 events were sampled for each bin from a Poisson distribution, where the mean of each distribution was given by the underlying model. In this example, the model with a 27 $M_\odot$ progenitor was used to sample the data n_i, and expected counts λ_i were computed for both 11 $M_\odot$ and 27 $M_\odot$ progenitors. The likelihood that each sampled data point belongs to a particular model, L, was calculated by multiplying the Poisson probability P of sampled counts in each bin, n_i, with the expected mean λ_i for that bin,

$$L = \prod_{i=1}^{N} P_i(n_i|\lambda_i). \tag{6.5}$$

For computational efficiency, the negative log-likelihood is subsequently calculated,

$$-\ln(L) = -\left(\sum_{i=1}^{N} \ln\left(P_i(n_i|\lambda_i)\right)\right). \tag{6.6}$$

An example of the sampled data is shown in Figure 6.9, along with the Poisson probabilities for each data point based on the expectations of both models, displayed in the subpanels below each graph. The log-likelihood value for this particular sample is also provided in the same figure.

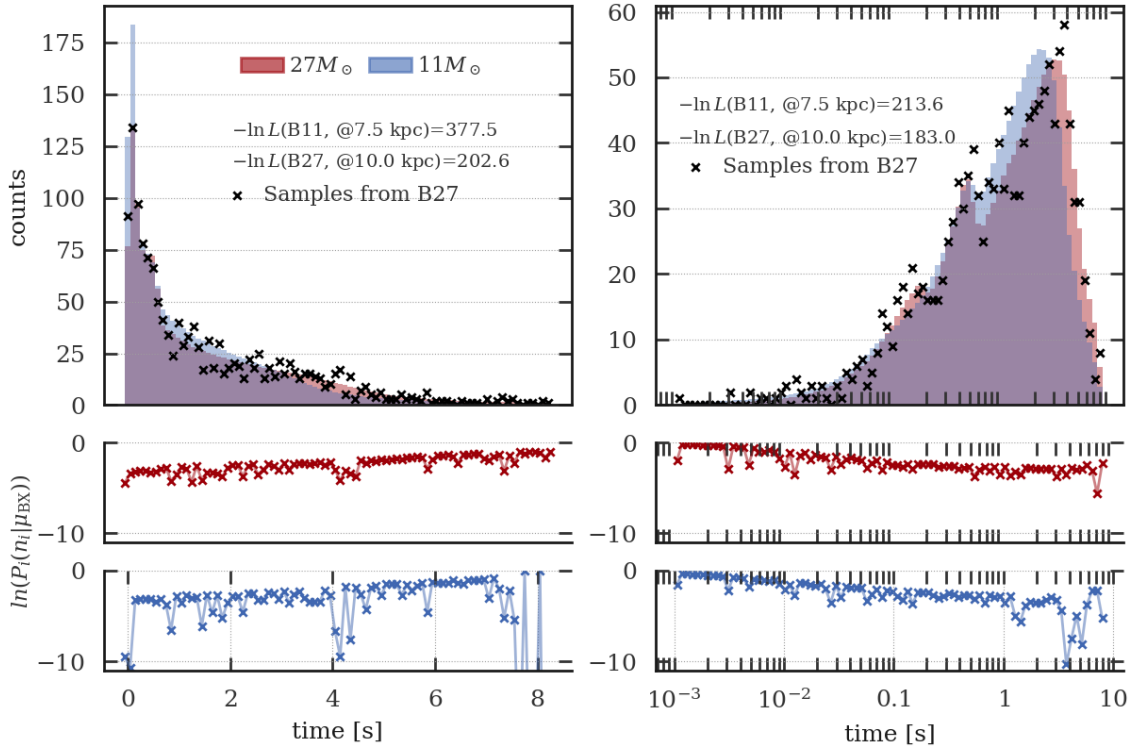

Figure 6.9: Linearly (*left*) and logarithmically (*right*) binned counts over time for both progenitors are shown as histograms. Sampled data from the Bollig model with a 27 $M_\odot$ progenitor is included in the plots. The logarithm of the Poisson probabilities for each bin are displayed below each graph: red for the model with a 27 $M_\odot$ progenitor, and blue for the model with a 11 $M_\odot$ progenitor.

It is important to note that the log–likelihood value alone is not sufficient to reject a model; a comparison of the log-likelihood distributions of the two models, which reveals their similarities or differences, is also needed. As shown in Figure 6.10, the distribution of log–likelihood values for each Monte Carlo sample is plotted, from which the statistical *power* of the test (its ability to distinguish between models) can be derived. Power of a test is the probability of correctly rejecting a false null hypothesis. A test with high power is more likely to distinguish between competing models based on observed data [152]. In this case, the observed log-likelihood values for the two models show significant differences, indicating that the test has a high discrimination power.

To evaluate the test power, a *critical value* was chosen based on the null hypothesis distribution (e.g., Bollig model with a 27 $M_\odot$ progenitor). This critical value was chosen to match a

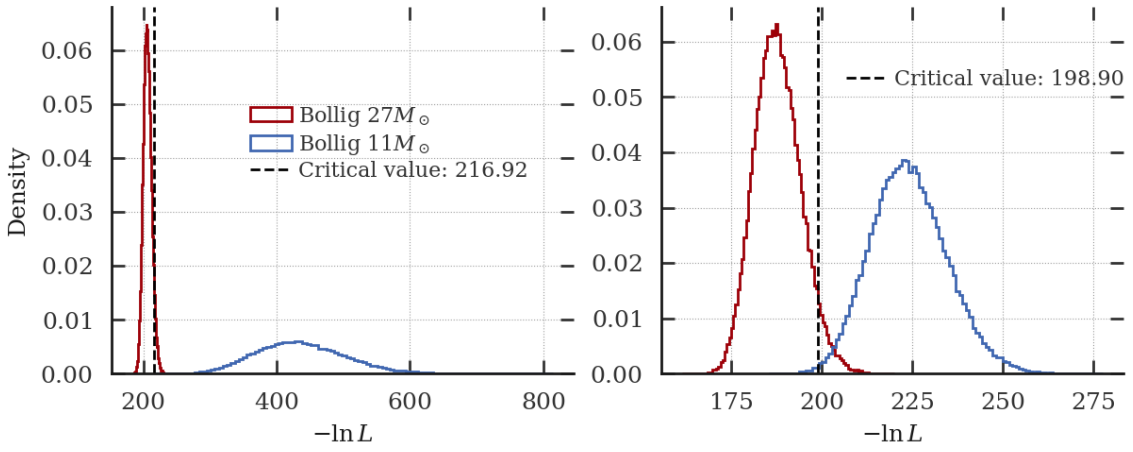

Figure 6.10: Distribution of negative log-likelihood values from toy Monte Carlo samples under different model assumptions. The red distribution corresponds to the Bollig model with a 27 $M_\odot$, while the blue corresponds to the Bollig model with a 11 $M_\odot$. The left panel shows results for the linearly binned analysis, and the right panel shows results for the logarithmic binning. In all cases, the sampled data is based on the Bollig 27 $M_\odot$ model.

predefined significance level, typically $\alpha = 0.05$, corresponding to 5% false positive rate, such that

$$-\ln L_{\text{crit}} = \text{percentile}(-\ln L_{\text{null}}, 100 \times (1-\alpha)). \tag{6.7}$$

Once this critical value is determined, the statistical power β, can be calculated as the probability that the log-likelihood value for the alternative hypothesis (e.g., Bollig model with a 11 $M_\odot$ progenitor) exceeds the critical value.

$$\text{Power} = \beta = P(-\ln L_{\text{alt}} > -\ln L_{\text{crit}}). \tag{6.8}$$

In the given example, the log-likelihood test against the Bollig model with 11 $M_\odot$ progenitor has a rejection power of 100% for the linearly binned analyses and a power of 99.6% for the logarithmically binned analyses, meaning that any observation of a log-likelihood value above the given critical value can allow to reject the alternative hypothesis. Here, the slightly better power of the linearly binned analysis can be attributed to the high counts in the first bins where for the alternative model, Poisson probabilities are lower.

In Figure 6.11 log-likelihood distributions of several other models are shown. For this sampled dataset (from Bollig model with 27 $M_\odot$ progenitor, id2, shown in dashed orange distribution), the log-likelihood test can reject all models with a 100% power except for the following three models: Bollig model with 11 $M_\odot$ progenitor (id1, green distribution), Sukhbold models with 9.6 $M_\odot$ (id3, light blue distribution) and 27 $M_\odot$ (id1, dark blue distribution) progenitors. The rejection power for these models is listed in Table 6.1 for both linear and logarithmically binned cases.

It is important to emphasize that this is a simplified analysis based on several assumptions, including the precise knowledge of the signal's onset. In this analysis, the models were

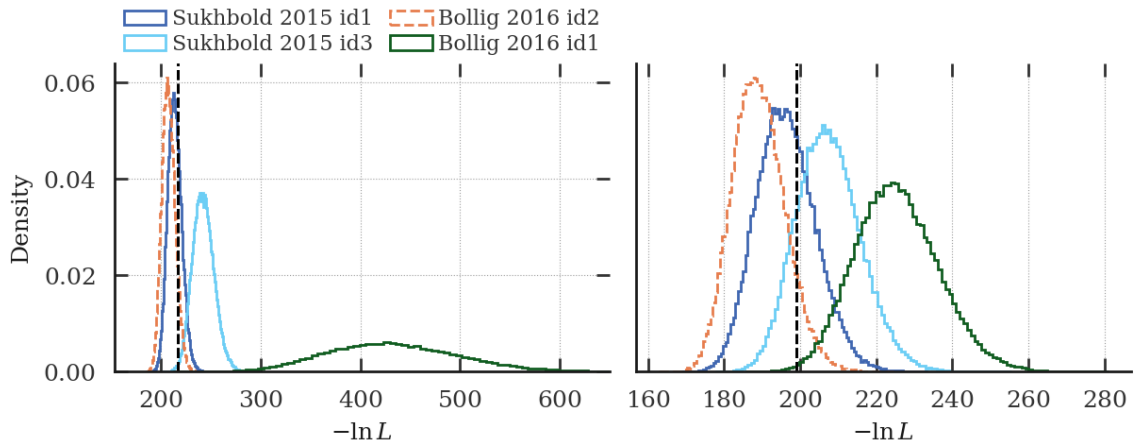

Figure 6.11: Distribution of negative log-likelihood values from toy Monte Carlo samples based on different model expectations. Bollig models (*id1* and *id2*) correspond to progenitors with masses of 11 $M_\odot$ and 27 $M_\odot$, respectively. Sukhbold models (*id1* and *id3*) represent progenitors of 27 $M_\odot$ and 9.6 $M_\odot$. The left panel shows results for the linearly binned analysis, while the right panel shows the logarithmic binning. The sampled data is always based on the Bollig 27 $M_\odot$ model.

	Bollig id1	Sukhbold id1	Sukhbold id3
Linear scale (%)	99.9%	21.0%	99.0%
Logarithmic scale (%)	99.3%	26.0%	77.0%

Table 6.1: Rejection power of the log-likelihood test statistic against different models. When tested with samples from Bollig model with a 27 $M_\odot$ progenitor, all models except these three can be rejected with 100% power.

aligned by matching the peak signal rate and binning data from the first event. However, in reality, accurately identifying the first event is important for a binned analysis to be reliable. Unbinned strategies, where the initial time is treated as a free parameter, could be explored in future studies. Additionally, this analysis neglects any potential background contamination, assuming only Poisson fluctuations in the observed counts. The effects of detection efficiencies and signal survival probabilities were similarly ignored. Since the recoil energy spectra are assumed to be similar across models, the reductions due to detection efficiencies and signal selections can be treated as uniform scaling in distance. As such, once the total count N is observed, this reduction can be treated as a distance scaling factor. Furthermore, following the supernova event, electromagnetic observatories are expected to provide a distance estimate within hours, which will enhance the precision of log-likelihood tests and facilitate easier model discrimination.

7

Summary and Final Remarks

Supernovae are some of the most fascinating phenomena in the universe, responsible for the synthesis of elements crucial to life on Earth. Despite significant progress in the last decades, including the SN1987A, the explosion mechanism of supernovae remains poorly understood. DM detectors, designed primarily to search for rare interactions, are also sensitive to neutrinos emitted during these cosmic events through CEνNS.

We have developed methods to effectively select supernova signals against background noise. Our XGB selection technique, utilizing advanced feature engineering, achieves a signal survival rate of 78.79%, while reducing background acceptance to just 0.08%. Given the short duration of a supernova burst, clustering of surviving peaks is instrumental in identifying the signal. Peak-finding algorithms have been optimized to detect all considered supernova models at 10 kpc.

To address false positives, we implemented analysis frameworks for multiple data levels, including a layer of Convolutional Neural Networks to reject spurious signals. Additionally, our statistical analysis shows that coincidence detection with outer veto detectors significantly increases the significance of observed signals, reinforcing the confidence in a true supernova detection.

Our study demonstrates that XENONnT's TPC data is highly likely to detect supernova neutrino signals under various assumptions, including different supernova models and distances. The analysis framework can identify the benchmark Bollig model with 100% accuracy up to 15 kpc, and has a non-zero detection probability extending to the edge of the Milky Way. For supernova events within 10 kpc, all models are identified with 100% accuracy, and the detection probability increases for closer distances, ensuring that any supernova within this range will be detected with high confidence.

The SNEWS is undergoing a major upgrade, introducing new tools such as the SNEWS Publishing Tools and the SNEWS Coincidence System, both developed in this work. These improvements enhance coordination between observatories and experiments. XENONnT, as the first DM experiment integrated into the SNEWS network, is fully operational; receiving alerts, sending heartbeats, and soon automating the transmission of observation messages to the network.

Looking ahead, there are exciting prospects for future work. Next-generation DM detectors with larger target masses will provide higher statistics during a supernova explosion, enabling the reconstruction of supernova light curves and the ability to rule out theoretical models. Since the CEνNS channel provides flavor-blind information, combining these data with observations from large neutrino observatories such as IceCube could offer new insights into

neutrino properties. Distance estimates from electromagnetic observatories will further refine model assumptions and facilitate deeper investigation into supernova physics.

In 1987, only 24 neutrino events were detected from a supernova with three observatories. Since then, remarkable advancements have been made, and current experiments are well positioned to capture valuable data from the next galactic supernova. This time, the multimessenger approach will allow us to explore stellar evolution, nuclear processes, and neutrino physics with unprecedented precision.

8

Bibliography

[1] R. Oerter. *The theory of almost everything: The standard model, the unsung triumph of modern physics.*: Penguin Publishing Group (2006)ISBN: isbn: 978-0452287860.

[2] M. K. Gaillard, P. D. Grannis, and F. J. Sciulli: The standard model of particle physics. In *Rev. Mod. Phys.* (1999), volume 71:S96–S111.
doi: 10.1103/RevModPhys.71.S96.

[3] T. W. Kibble: The Standard Model of Particle Physics. In *European Review* (2015), volume 23(1).
doi: 10.1017/S1062798714000520.

[4] CERN. *OPEN-PHO-CHART-2015-001-1: Particle Physics Chart*. 2015.
url: `https://cds.cern.ch/images/OPEN-PHO-CHART-2015-001-1/`, Accessed: 2024-12-01.

[5] S. Horiuchi and J. P. Kneller: What can be learned from a future supernova neutrino detection? In *Journal of Physics G: Nuclear and Particle Physics* (2018), volume 45(4):043002.
doi: 10.1088/1361-6471/aaa90a.

[6] I. Iben and A. Renzini: Asymptotic Giant Branch Evolution and Beyond. In *Annual Review of Astronomy and Astrophysics* (1983), volume 21:271–342.
doi: 10.1146/annurev.aa.21.090183.001415.

[7] E. O'Connor and C. D. Ott: Black Hole Formation in Failing Core-Collapse Supernovae. In *The Astrophysical Journal* (2011), volume 730(2):70.
doi: 10.1088/0004-637X/730/2/70.

[8] S. J. Smartt: Progenitors of Core-Collapse Supernovae. In *Annual Review of Astronomy and Astrophysics* (2009), volume 47:63–106.
doi: 10.1146/annurev-astro-082708-101737.

[9] A. Burrows, D. Radice, and D. Vartanyan: Three-dimensional Supernova Explosion Simulations and Progenitor Dependence. In *Monthly Notices of the Royal Astronomical Society* (2020), volume 491(3):2715–2735.
doi: 10.3847/2041-8213/aaa967.

[10] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann: How Massive Single Stars End Their Life. In *The Astrophysical Journal* (2003), volume 591(1):288–300.
doi: 10.1086/375341.

[11] S. E. Woosley and H.-T. Janka, "The Physics of Core-Collapse Supernovae", *Nature Physics*, vol. 1, 3, Nature Publishing Group, 2005, 147–154,
doi: 10.1038/nphys172.

[12] NASA Goddard Space Flight Center. *Lifecycle of Stars*. 2009.
url: `https://imagine.gsfc.nasa.gov/science/objects/images/stars_lifecycle_full.jpg`, Accessed: 2024-12-01.

[13] H.-T. Janka, "Neutrino Emission from Supernovae", *Handbook of Supernovae*, Springer, 2017, 1575–1604,
doi: 10.1007/978-3-319-20794-0_4-1.

[14] A. Mirizzi, I. Tamborra, H.-T. Janka, N. Saviano, K. Scholberg, R. Bollig, L. Hüdepohl, and S. Chakraborty: Supernova Neutrinos: Production, Oscillations and Detection. In *Rivista del Nuovo Cimento* (2016), volume 39(1-2):1–112.
doi: 10.1393/ncr/i2016-10120-8.

[15] S. Wanajo, B. Müller, H.-T. Janka, and A. Heger: Nucleosynthesis in the Innermost Ejecta of Neutrino-driven Supernova Explosions in Two Dimensions. In *The Astrophysical Journal* (2018), volume 852(1):40.
doi: 10.3847/1538-4357/aa9d97.

[16] Northern Arizona University Meteorite Laboratory. *The Origin of the Chemical Elements*. 2007.
url: `https://en.wikipedia.org/wiki/Nucleosynthesis#/media/File:Nucleosynthesis_periodic_table.svg`, Available on Wikimedia Commons and sourced from NASA.

[17] K. Nakamura, S. Horiuchi, M. Tanaka, K. Hayama, T. Takiwaki, and K. Kotake: Multimessenger signals of long-term core-collapse supernova simulations: synergetic observation strategies. In *Monthly Notices of the Royal Astronomical Society* (2016), volume 461(3):3296–3313.
doi: 10.1093/mnras/stw1453.

[18] J. Vink. *Supernova 1604, Kepler's Supernova, and its Remnant.*: Springer (2017).

[19] K. Hirata et al.: Observation of a Neutrino Burst from the Supernova SN 1987A. In *Physical Review Letters* (1987), volume 58(14):1490–1493.
doi: 10.1103/PhysRevLett.58.1490.

[20] R. M. Bionta, G. Blewitt, C. B. Bratton, D. Caspere, and A. Ciocio: Observation of a Neutrino Burst in Coincidence with Supernova SN 1987A in the Large

Magellanic Cloud. In *Physical Review Letters* (1987), volume 58(14):1494–1496.
DOI: 10.1103/PhysRevLett.58.1494.

[21] E. N. Alexeyev, L. N. Alexeyeva, I. V. Krivosheina, and V. I. Volchenko: Detection of the Neutrino Signal from SN 1987A in the Baksan Underground Scintillation Telescope. In *Physics Letters B* (1988), volume 205(2-3):209–214.
DOI: 10.1134/S1063773708110030.

[22] M. Aglietta et al.: On the Event Observed in the Mont Blanc Underground Neutrino Observatory During the Occurrence of Supernova 1987A. In *Europhysics Letters* (1987), volume 3(12):1315.
DOI: 10.1209/0295-5075/3/12/011.

[23] A. Burrows and J. M. Lattimer: Neutrinos from SN 1987A. In (1987), volume 318:L63.
DOI: 10.1086/184938.

[24] H. A. Bethe: Supernova Mechanisms. In *Reviews of Modern Physics* (1990), volume 62(4):801–866.
DOI: 10.1103/RevModPhys.62.801.

[25] D. H. Clark and F. R. Stephenson, "The Historical Supernovae", *Supernovae: A Survey of Current Research*, ed. by M. J. Rees and R. J. Stoneham, Dordrecht: Springer Netherlands, 1982, 355–370, ISBN: 978-94-009-7876-8,
DOI: 10.1007/978-94-009-7876-8_22.

[26] F. R. Stephenson and D. A. Green: Historical supernovae and their remnants. In *International Series in Astronomy and Astrophysics* (2002), volume 5.
DOI: 10.1093/acprof:oso/9780198507666.001.0001.

[27] K. Brecher, R. A. Fesen, S. P. Maran, and J. C. Brandt: Ancient records and the Crab Nebula supernova. In *The Observatory* (1983), volume 103:106–113.
DOI: 103(1054), 106-113.

[28] S. P. Reynolds, K. J. Borkowski, U. Hwang, J. P. Hughes, C. Badenes, J. M. Laming, and J. M. Blondin: A Deep Chandra Observation of Kepler's Supernova Remnant: A Type Ia Event with Circumstellar Interaction. In *The Astrophysical Journal* (2007), volume 668(2):L135.
DOI: 10.1086/522830.

[29] S. M. Adams, C. S. Kochanek, J. F. Beacom, M. R. Vagins, and K. Z. Stanek: Observing the Next Galactic Supernova. In *The Astrophysical Journal* (2013), volume 778(2):164.
DOI: 10.1088/0004-637x/778/2/164.

[30] K. Rozwadowska, F. Vissani, and E. Cappellaro: On the rate of core collapse supernovae in the milky way. In *New Astronomy* (2021), volume 83:101498.
doi: 10.1016/j.newast.2020.101498.

[31] IceCube Collaboration: Search for Galactic Core-collapse Supernovae in a Decade of Data Taken with the IceCube Neutrino Observatory. In *The Astrophysical Journal* (2024), volume 961(1):84.
doi: 10.3847/1538-4357/ad07d1.

[32] A. A. Hakobyan et al.: Supernovae and their host galaxies – III. The impact of bars and bulges on the radial distribution of supernovae in disc galaxies. In *Monthly Notices of the Royal Astronomical Society* (2015), volume 456(3):2848–2860.
doi: 10.1093/mnras/stv2853.

[33] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller: Theory of core-collapse supernovae. In *Physics Reports* (2007), volume 442(1):38–74.
doi: 10.1016/j.physrep.2007.02.002.

[34] R. Laha, J. F. Beacom, and S. K. Agarwalla. *New Power to Measure Supernova ν_e with Large Liquid Scintillator Detectors*. 2014.
doi: 10.48550/arXiv.1412.8425.

[35] K. Collaboration and T. S.-K. Collaboration: Combined Pre-supernova Alert System with KamLAND and Super-Kamiokande. In *The Astrophysical Journal* (2024), volume 973(2):140.
doi: 10.3847/1538-4357/ad5fee.

[36] A. Burrows and J. M. Lattimer: Neutrino Trapping and Accretion during Gravitational Collapse. In *The Astrophysical Journal* (1986), volume 307:178–196.
doi: 10.3847/1538-4357/aaa296.

[37] Y.-Z. Qian and S. E. Woosley: Nucleosynthesis in Neutrino-driven Winds: II. The r-Process. In *The Astrophysical Journal* (1996), volume 471:331–351.
doi: 10.1086/304181.

[38] S. E. Woosley, J. R. Wilson, G. J. Mathews, R. D. Hoffman, and B. S. Meyer: The r-Process and Neutrino-heated Supernova Ejecta. In *The Astrophysical Journal* (1994), volume 433:229–246.
doi: 10.1086/174638.

[39] A. Bellerive, J. R. Klein, A. B. McDonald, A. J. Noble, and A. W. P. Poon: The Sudbury Neutrino Observatory. In *Nucl. Phys. B* (2016), volume 908:30–51.
doi: 10.1016/j.nuclphysb.2016.04.035.

[40] K. Scholberg: Supernova Neutrino Detection. In *Annual Review of Nuclear and Particle Science* (2012), volume 62:81–103.
doi: 10.1146/annurev-nucl-102711-095006.

[41] F. Tortorici, V. Bellini, C. Petta, and C. Sutera: The ICARUS Experiment. In *Nuclear and Particle Physics Proceedings* (2019), volume 306-308:154–162. doi: 10.1016/j.nuclphysbps.2019.07.022.

[42] B. Abi et al.: Volume I. Introduction to DUNE. In *Journal of Instrumentation* (2020), volume 15(08):T08008. doi: 10.1088/1748-0221/15/08/T08008.

[43] E. Aprile et al.: Observation of Two-Neutrino Double Electron Capture in ^{124}Xe with XENON1T. In *Nature* (2020), volume 568:532–535. doi: 10.1038/s41586-019-1124-4.

[44] D. S. Akerib et al.: The LUX-ZEPLIN (LZ) Experiment. In *Nucl. Instrum. Meth. A* (2020), volume 953:163047. doi: 10.1016/j.nima.2019.163047.

[45] K. Scholberg: Coherent Elastic Neutrino-Nucleus Scattering. In *Journal of Physics: Conference Series* (2020), volume 1468(1):012126. doi: 10.1088/1742-6596/1468/1/012126.

[46] XENON Collaboration: First Indication of Solar ^{8}B Neutrinos via Coherent Elastic Neutrino-Nucleus Scattering with XENONnT. In *Phys. Rev. Lett.* (2024), volume 133:191002. doi: 10.1103/PhysRevLett.133.191002.

[47] R. F. Lang, C. McCabe, S. Reichard, M. Selvi, and I. Tamborra: Supernova neutrino physics with xenon dark matter detectors: A timely perspective. In *Physical Review D* (2016), volume 94(10). doi: 10.1103/physrevd.94.103009.

[48] M. Lisanti, "Lectures on Dark Matter Physics", *New Frontiers in Fields and Strings*, World Scientific, 2017, chap. Chapter 7, 399–446, doi: 10.1142/9789813149441_0007, eprint: `https://www.worldscientific.com/doi/pdf/10.1142/9789813149441_0007`.

[49] G. Bertone, D. Hooper, and J. Silk: Particle dark matter: evidence, candidates and constraints. In *Physics Reports* (2005), volume 405(5):279–390. doi: 10.1016/j.physrep.2004.08.031.

[50] G. Bertone. *Particle Dark Matter: Observations, Models and Searches.*: Cambridge University Press (2010)ISBN: isbn: 978-0521763684.

[51] Planck Collaboration: Planck2018 results: VI. Cosmological parameters. In *Astronomy & Astrophysics* (2020), volume 641:A6. doi: 10.1051/0004-6361/201833910.

[52] G. Börner. *The Early Universe. Facts and Fiction*. Berlin, Germany: Springer (2003)ISBN: isbn: 978-3-540-44197-7, 978-3-642-07915-3, 978-3-662-05267-9.

[53] M. Cirelli, A. Strumia, and J. Zupan. *Dark Matter*. 2024.
doi: 10.48550/arXiv.2406.01705.

[54] J. Schaye et al.: The EAGLE project: simulating the evolution and assembly of galaxies and their environments. In *Monthly Notices of the Royal Astronomical Society* (2014), volume 446(1):521–554.
doi: 10.1093/mnras/stu2058.

[55] T. Sawala et al.: The APOSTLE simulations: solutions to the Local Group's cosmic puzzles. In *Monthly Notices of the Royal Astronomical Society* (2016), volume 457(2):1931–1943.
doi: 10.1093/mnras/stw145.

[56] W. Tucker et al.: 1E 0657–56: A Contender for the Hottest Known Cluster of Galaxies. In *The Astrophysical Journal* (1998), volume 496(1):L5.
doi: 10.1086/311234.

[57] V. C. Rubin, J. Ford W. K., and N. Thonnard: Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa -> Sc. In (1978), volume 225:L107–L111.
doi: 10.1086/182804.

[58] E. Corbelli and P. Salucci: The extended rotation curve and the dark matter halo of M33. In (2000), volume 311(2):441–447.
doi: 10.1046/j.1365-8711.2000.03075.x.

[59] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky: A Direct Empirical Proof of the Existence of Dark Matter. In *The Astrophysical Journal* (2006), volume 648(2):L109–L113.
doi: 10.1086/508162.

[60] R. Massey, T. Kitching, and J. Richard: The dark matter of gravitational lensing. In *Reports on Progress in Physics* (2010), volume 73(8):086901.
doi: 10.1088/0034-4885/73/8/086901.

[61] R. Durrer. *The Cosmic Microwave Background*. Cambridge: Cambridge University Press (2008)ISBN: isbn: 978-0-511-81720-5.

[62] W. Hu and S. Dodelson: Cosmic Microwave Background Anisotropies. In *Annual Review of Astronomy and Astrophysics* (2002), volume 40(1):171–216.
doi: 10.1146/annurev.astro.40.060401.093926.

[63] V. Springel, C. S. Frenk, and S. D. M. White: The large-scale structure of the Universe. In *Nature* (2006), volume 440(7088):1137–1144.
doi: 10.1038/nature04805.

[64] H. E. Haber and L. S. Haskins, "Supersymmetric Theory and Models", *Anticipating the Next Discoveries in Particle Physics*, WORLD SCIENTIFIC, May 2018, 355–499,
doi: 10.1142/9789813233348_0006.

[65] J. L. Feng: Dark Matter Candidates from Particle Physics and Methods of Detection. In *Annual Review of Astronomy and Astrophysics* (2010), volume 48(1):495–545.
doi: 10.1146/annurev-astro-082708-101659.

[66] R. Peccei and H. Quinn: CP Conservation in the Presence of Pseudoparticles. In *Physical Review Letters - PHYS REV LETT* (1977), volume 38:1440–1443.
doi: 10.1103/PhysRevLett.38.1440.

[67] S. Dodelson and L. M. Widrow: Sterile neutrinos as dark matter. In *Physical Review Letters* (1994), volume 72(1):17–20.
doi: 10.1103/physrevlett.72.17.

[68] B. Carr, F. Kühnel, and M. Sandstad: Primordial black holes as dark matter. In *Physical Review D* (2016), volume 94(8).
doi: 10.1103/physrevd.94.083504.

[69] A. M. Green: Primordial black holes as a dark matter candidate - a brief overview. In *Nucl. Phys. B* (2024), volume 1003:116494.
doi: 10.1016/j.nuclphysb.2024.116494.

[70] E. Aprile et al.: The XENONnT dark matter experiment. In *Eur. Phys. J. C* (2024), volume 84(8):784.
doi: 10.1140/epjc/s10052-024-12982-5.

[71] X. Cao et al.: PandaX: A Liquid Xenon Dark Matter Experiment at CJPL. In *Sci. China Phys. Mech. Astron.* (2014), volume 57:1476–1494.
doi: 10.1007/s11433-014-5521-2.

[72] PandaX Collaboration: Dark matter direct search sensitivity of the PandaX-4T experiment. In *Sci. China Phys. Mech. Astron.* (2019), volume 62(3):31011.
doi: 10.1007/s11433-018-9259-0.

[73] J. Aalbers et al.: DARWIN: towards the ultimate dark matter detector. In *Journal of Cosmology and Astroparticle Physics* (2016), volume 2016(11):017.
doi: 10.1088/1475-7516/2016/11/017.

[74] L. Baudis: DARWIN/XLZD: A future xenon observatory for dark matter and other rare interactions. In *Nuclear Physics B* (2024), volume 1003:116473.
doi: 10.1016/j.nuclphysb.2024.116473.

[75] J. Aalbers et al.: The XLZD Design Book: Towards the Next-Generation Liquid Xenon Observatory for Dark Matter and Neutrino Physics. In *(forthcoming)* (2024), volume.
doi: 10.48550/arXiv.2410.17137.

[76] J. Angevaare. *First WIMP Results of XENONnT and Its Signal Reconstruction.* 2023.
url: https://hdl.handle.net/11245.1/140ec608-e872-4715-8417-3b7b2660e529, PhD thesis, University of Amsterdam, Faculty of Science (FNWI), Institute of Physics (IoP). Supervisors: M.P. Decowski, A.P. Colijn. ISBN: 9789464198096, 179 pages.

[77] SuperCDMS Collaboration. *A Strategy for Low-Mass Dark Matter Searches with Cryogenic Detectors in the SuperCDMS SNOLAB Facility.* 2023.
doi: 10.48550/arXiv.2203.08463.

[78] CRESST Collaboration: Detector development for the CRESST experiment. In *Journal of Low Temperature Physics* (2024), volume.
doi: 10.1007/s10909-024-03154-6.

[79] Y. Hochberg, Y. Zhao, and K. M. Zurek: Superconducting Detectors for Superlight Dark Matter. In *Phys. Rev. Lett.* (2016), volume 116:011301.
doi: 10.1103/PhysRevLett.116.011301.

[80] E. Aprile and T. Doke: Liquid xenon detectors for particle physics and astrophysics. In *Rev. Mod. Phys.* (2010), volume 82:2053–2097.
doi: 10.1103/RevModPhys.82.2053.

[81] E. Aprile et al.: The XENON100 dark matter experiment. In *Astroparticle Physics* (2012), volume 35(9):573–590.
doi: 10.1016/j.astropartphys.2012.01.003.

[82] E. Aprile et al.: The XENON1T Dark Matter Experiment. In *Eur. Phys. J. C* (2017), volume 77(12):881.
doi: 10.1140/epjc/s10052-017-5326-3.

[83] M. Schumann: Direct detection of WIMP dark matter: concepts and status. In *Journal of Physics G: Nuclear and Particle Physics* (2019), volume 46(10):103003.
doi: 10.1088/1361-6471/ab2ea5.

[84] XENON Collaboration: Energy resolution and linearity of XENON1T in the MeV energy range. In *Eur. Phys. J. C* (2020), volume 80(8):785.
doi: 10.1140/epjc/s10052-020-8284-0.

[85] XENON Collaboration: Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T. In *Nature* (2019), volume 568(7753):532–535.
doi: 10.1038/s41586-019-1124-4.

[86] EXO Collaboration: Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector. In *Phys. Rev. C* (2014), volume 89:015502.
doi: 10.1103/PhysRevC.89.015502.

[87] XENON Collaboration: Double-weak decays of ^{124}Xe and ^{136}Xe in the XENON1T and XENONnT experiments. In *Phys. Rev. C* (2022), volume 106:024328.
doi: 10.1103/PhysRevC.106.024328.

[88] K. J. R. Rosman and P. D. P. Taylor: Isotopic compositions of the elements 1997 (Technical Report). In *Pure and Applied Chemistry* (1998), volume 70(1):217–235.
doi: doi:10.1351/pac199870010217.

[89] XENON Collaboration: Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment. In *Journal of Cosmology and Astroparticle Physics* (2020), volume 2020(11):031.
doi: 10.1088/1475-7516/2020/11/031.

[90] K. Abe et al.: XMASS detector. In *Nucl. Instrum. Meth. A* (2013), volume 716:78–85.
doi: 10.1016/j.nima.2013.03.059.

[91] M. Auger et al.: The EXO-200 detector, part I: Detector design and construction. In *JINST* (2012), volume 7:P05010.
doi: 10.1088/1748-0221/7/05/P05010.

[92] E. Aprile, A. E. Bolotnikov, A. L. Bolozdynya, and T. Doke. *Noble Gas Detectors.*: Wiley (2006)ISBN: isbn: 9783527610020.

[93] XENON Collaboration: First Dark Matter Search with Nuclear Recoils from the XENONnT Experiment. In *Physical Review Letters* (2023), volume 131(4).
doi: 10.1103/physrevlett.131.041003.

[94] J. A. Formaggio and G. P. Zeller: From eV to EeV: Neutrino Cross Sections Across Energy Scales. In *Reviews of Modern Physics* (2012), volume 84(3):1307–1341.
doi: 10.1103/RevModPhys.84.1307.

[95] D. Z. Freedman, D. N. Schramm, and D. L. Tubbs: The Weak Neutral Current and its Effects in Stellar Collapse. In *Annual Review of Nuclear Science* (1977), volume 27:167–207.
doi: 10.1146/annurev.ns.27.120177.001123.

[96] A. Drukier and L. Stodolsky: Principles and Applications of a Neutral Current Detector for Neutrino Physics and Astronomy. In *Physical Review D* (1984), volume 30(11):2295–2309.
doi: 10.1103/PhysRevD.30.2295.

[97] P. Bhattacharjee, A. Bandyopadhyay, S. Chakraborty, S. Ghosh, K. Kar, and S. Saha: Inelastic charged-current interactions of supernova neutrinos in two-

phase liquid xenon dark matter detectors. In *Phys. Rev. D* (2022), volume 106:043029. DOI: 10.1103/PhysRevD.106.043029.

[98] R. H. Helm: Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei. In *Physical Review* (1956), volume 104(5):1466–1475. DOI: 10.1103/PhysRev.104.1466.

[99] L. Vietze, P. Klos, J. Menéndez, W. Haxton, and A. Schwenk: Nuclear structure aspects of spin-independent WIMP scattering off xenon. In *Physical Review D* (2015), volume 91(4). DOI: 10.1103/physrevd.91.043520.

[100] XENON Collaboration: Design and performance of the XENON10 dark matter experiment. In *Astroparticle Physics* (2011), volume 34(9):679–698. DOI: 10.1016/j.astropartphys.2011.01.006.

[101] XENON Collaboration: Physics reach of the XENON1T dark matter experiment. In *Journal of Cosmology and Astroparticle Physics* (2016), volume 2016(04):027. DOI: 10.1088/1475-7516/2016/04/027.

[102] XENON Collaboration: The triggerless data acquisition system of the XENONnT experiment. In *Journal of Instrumentation* (2023), volume 18(07):P07054. DOI: 10.1088/1748-0221/18/07/P07054.

[103] J. Aalbers and et al. *AxFoundation/strax: Streaming analysis for xenon experiments.* 2021. DOI: 10.5281/zenodo.5576034.

[104] XENON Collaboration. *XENONnT/straxen: Streaming analysis for XENON(nT).* 2022. DOI: 10.5281/zenodo.6102391.

[105] XENON Collaboration. *XENONnT Analysis: Signal Reconstruction, Calibration and Event Selection.* 2024. DOI: 10.48550/arXiv.2409.08778.

[106] XENON Collaboration: Emission of single and few electrons in XENON1T and limits on light dark matter. In *Phys. Rev. D* (2022), volume 106:022001. DOI: 10.1103/PhysRevD.106.022001.

[107] XENON Collaboration: Observation and applications of single-electron charge signals in the XENON100 experiment. In *Journal of Physics G: Nuclear and Particle Physics* (2014), volume 41(3):035201. DOI: 10.1088/0954-3899/41/3/035201.

[108] XENON Collaboration: Search for New Physics in Electronic Recoil Data from XENONnT. In *Phys. Rev. Lett.* (2022), volume 129:161805.
doi: 10.1103/PhysRevLett.129.161805.

[109] LUX Collaboration: Radiogenic and muon-induced backgrounds in the LUX dark matter detector. In *Astroparticle Physics* (2015), volume 62:33–46.
doi: 10.1016/j.astropartphys.2014.07.009.

[110] A. M. Daniel Layos Garcia. *Preliminary Work on analysis of muon and neutron veto for supernova detection.* 2024, Preliminary findings, [2024].

[111] XENON Collaboration. *XENONnT WIMP Search: Signal & Background Modeling and Statistical Inference.* 2024.
doi: 10.48550/arXiv.2406.13638.

[112] C. A. J. O'Hare: New Definition of the Neutrino Floor for Direct Dark Matter Searches. In *Phys. Rev. Lett.* (2021), volume 127:251802.
doi: 10.1103/PhysRevLett.127.251802.

[113] M. Kara et al.: The SNEWS2.0 alert software for the coincident detection of neutrinos from core-collapse supernovae. In *Journal of Instrumentation* (2024), volume 19(10):P10017.
doi: 10.1088/1748-0221/19/10/P10017.

[114] SNEWS Collaboration. *SNEWS 2.0 Website.*
url: `https://snews2.org/`.

[115] K. Rozwadowska, F. Vissani, and E. Cappellaro: On the rate of core collapse supernovae in the milky way. In *New Astronomy* (2021), volume 83:101498.
doi: 10.1016/j.newast.2020.101498.

[116] P. Antonioli et al.: SNEWS: the SuperNova Early Warning System. In *New Journal of Physics* (2004), volume 6:114–114.
doi: 10.1088/1367-2630/6/1/114.

[117] SNEWS Collaboration. *The SuperNova Early Warning System.* 2020.
url: `https://snews.bnl.gov`.

[118] S. A. Kharusi et al.: SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy. In *New Journal of Physics* (2021), volume 23(3):031201.
doi: 10.1088/1367-2630/abde33.

[119] F. Abe et al.: Multi-messenger Observations of a Binary Neutron Star Merger. In *Nature* (2021), volume 551:85–88.
doi: 10.1038/nature24471.

[120] P. Baerwald, L. Hüdepohl, B. Müller, and W. Winter: Neutrino Supernova Signals and Their Correlation with Electromagnetic and Gravitational-Wave Emissions. In *Astroparticle Physics* (2012), volume 35(10):641–659.
doi: 10.1016/j.astropartphys.2012.02.005.

[121] A. Abramowski et al.: Multi-wavelength Observations of Supernova Remnants: Gravitational Waves, Electromagnetic Signals, and Particle Acceleration. In *Astronomy & Astrophysics* (2011), volume 531:A72.
doi: 10.1051/0004-6361/201016084.

[122] Y. Fukuda et al.: The Super-Kamiokande detector. In *Nucl. Instrum. Meth. A* (2003), volume 501:418–462.
doi: 10.1016/S0168-9002(03)00425-X.

[123] K. Abe et al.: Hyper-Kamiokande Design Report. In *(forthcoming)* (2018), volume.
doi: 10.48550/arXiv.1805.04163.

[124] I. A. Belolaptikov et al.: The Baikal underwater neutrino telescope: Design, performance and first results. In *Astropart. Phys.* (1997), volume 7:263–282.
doi: 10.1016/S0927-6505(97)00022-4.

[125] M. G. Aartsen et al.: The IceCube Neutrino Observatory: Instrumentation and Online Systems. In *JINST* (2017), volume 12(03):P03012.
doi: 10.1088/1748-0221/12/03/P03012.

[126] A. Margiotta: The KM3NeT deep-sea neutrino telescope. In *Nucl. Instrum. Meth. A* (2014), volume 766:83–87.
doi: 10.1016/j.nima.2014.05.090.

[127] R. Acciarri et al.: Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 1: The LBNF and DUNE Projects. In *arXiv e-prints* (2016), volume:arXiv:1601.05471.
doi: 10.48550/arXiv.1601.05471.

[128] S. Amerio et al.: Design, construction and tests of the ICARUS T600 detector. In *Nucl. Instrum. Meth. A* (2004), volume 527:329–410.
doi: 10.1016/j.nima.2004.02.044.

[129] G. Alimonti et al.: The Borexino detector at the Laboratori Nazionali del Gran Sasso. In *Nucl. Instrum. Meth. A* (2009), volume 600:568–593.
doi: 10.1016/j.nima.2008.11.076.

[130] F. An et al.: Neutrino Physics with JUNO. In *J. Phys. G* (2016), volume 43(3):030401.
doi: 10.1088/0954-3899/43/3/030401.

[131] A. A. of Variable Star Observers. *AAVSO: Neutrino Astronomy Collaboration.*
url: `https://www.aavso.org`.

[132] N. Sravan, D. Milisavljevic, J. M. Reynolds, G. Lentner, and M. Linvill: Real-time, Value-driven Data Augmentation in the Era of LSST. In *The Astrophysical Journal* (2020), volume 893(2):127.
doi: 10.3847/1538-4357/ab8128.

[133] S. Healy et al.: Red supergiant candidates for multimessenger monitoring of the next Galactic supernova. In *Monthly Notices of the Royal Astronomical Society* (2024), volume 529(4):3630–3650.
doi: 10.1093/mnras/stae738.

[134] B. P. Abbott et al.: Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. In *Living Reviews in Relativity* (2020), volume 23(1).
doi: 10.1007/s41114-020-00026-9.

[135] F. Azfar et al. *Data-driven core collapse supernova multilateration with first neutrino events*. 2024.
doi: 10.48550/arXiv.2410.11984.

[136] A. Brazier, "SCIMMA: Collaboration and information transfer cyberinfrastructure for Multi-Messenger Astrophysics", *American Astronomical Society Meeting Abstracts*, vol. 237, American Astronomical Society Meeting Abstracts, January 2021, 146.03, 146.03.

[137] A. Brazier et al., "SCiMMA Hopskotch: performant, robust, and secure messaging cyberinfrastucture for Multi-Messenger Astrophysics", *American Astronomical Society Meeting #240*, vol. 240, American Astronomical Society Meeting Abstracts, June 2022, 348.10, 348.10.

[138] M. J. Sax. *Apache Kafka*. Cham: Springer International Publishing (2018)ISBN: isbn: 978-3-319-63962-8.

[139] B. T. Tsang, D. Vartanyan, and A. Burrows: Applications of Machine Learning to Predicting Core-collapse Supernova Explosion Outcomes. In *The Astrophysical Journal Letters* (2022), volume.
doi: 10.3847/2041-8213/ac8f4b.

[140] H. Nagakura, A. Burrows, D. Radice, and D. Vartanyan: Towards an understanding of the resolution dependence of Core-Collapse Supernova simulations. In *Monthly Notices of the Royal Astronomical Society* (2019), volume.
doi: 10.1093/mnras/stz2730.

[141] E. O'Connor and S. Couch: Exploring Fundamentally Three-dimensional Phenomena in High-fidelity Simulations of Core-collapse Supernovae. In *The Astrophysical Journal* (2018), volume.
doi: 10.3847/1538-4357/aadcf7.

[142] A. Burrows, D. Vartanyan, J. Dolence, M. A. Skinner, and D. Radice: Crucial Physical Dependencies of the Core-Collapse Supernova Mechanism. In *Space Science Reviews* (2016), volume.
doi: 10.1007/s11214-017-0450-9.

[143] M. T. Keil, G. G. Raffelt, and H.-T. Janka: Monte Carlo Study of Supernova Neutrino Spectra Formation. In *The Astrophysical Journal* (2003), volume 590(2):971–991.
doi: 10.1086/375130.

[144] G. G. Raffelt. *Stars as laboratories for fundamental physics: The astrophysics of neutrinos, axions, and other weakly interacting particles.*: University of Chicago Press (1996)ISBN: isbn: 978-0-226-70272-8.

[145] R. Bollig, H.-T. Janka, A. Lohs, G. Martínez-Pinedo, C. J. Horowitz, and T. Melson: Muon Creation in Supernova Matter Facilitates Neutrino-Driven Explosions. In *Phys. Rev. Lett.* (2017), volume 119:242702.
doi: 10.1103/PhysRevLett.119.242702.

[146] K. Nakazato, K. Sumiyoshi, and H. Suzuki: Supernova Neutrino Light Curves and Spectra for Various Progenitor Stars: From Core Collapse to Proto-neutron Star Cooling. In *The Astrophysical Journal Supplement Series* (2013), volume 205(2):2.
doi: 10.1088/0067-0049/205/2/2.

[147] B. Pang et al. *Detecting Neutrinos from Supernova Bursts in PandaX-4T*. 2024.
doi: 10.1088/1674-1137/ad3efe.

[148] T. Sukhbold, T. Ertl, S. E. Woosley, J. M. Brown, and H.-T. Janka: Core-collapse Supernovae from 9 to 120 Solar Masses Based on Neutrino-powered Explosions. In *The Astrophysical Journal* (2016), volume 821(1):38.
doi: 10.3847/0004-637X/821/1/38.

[149] J. Brodsky, C. Tunnell, mszydagis, jbalajth, V. Velan, and J. Huang. *NEST-Collaboration/nest: Geant4 Integration Fixes and Updates*. 2019.
doi: 10.5281/zenodo.2535713.

[150] H. S. Eissing et al. *XENONnT/fuse: 1.4.2*. 2024.
doi: 10.5281/zenodo.13940786.

[151] M. Kara and R. Peres. *XENONnT/multimessenger : snax*. 2024.
url: `https://github.com/XENONnT/multimessenger/`.

[152] G. Cowan, K. Cranmer, E. Gross, and O. Vitells: Asymptotic formulae for likelihood-based tests of new physics. In *The European Physical Journal C* (2011), volume 71(2).
doi: 10.1140/epjc/s10052-011-1554-0.

[153] A. Sheshukov, A. Vishneva, and A. Habig: Combined detection of supernova neutrino signals. In *Journal of Cosmology and Astroparticle Physics* (2021), volume 2021(12):053.
doi: 10.1088/1475-7516/2021/12/053.

[154] A. Sheshukov. *Sheshuk/sn_stat: v0.3.3*. 2021.
doi: 10.5281/zenodo.5478391.

List of Abbreviations

AC Accidental Coincidence

AFT Area Fraction Top

API Application Programming Interface

AUC Area Under the Curve

BAO Baryonic Acoustic Oscillations

BNL Brooklyn National Labaroty

CC Charged Current

CCSN Core Collapse Supernova

CEνNS Coherent Elastic Neutrino Nucleus Scattering

CLI Command Line Interface

CMB Cosmic Microwave Background

CNN Convolutional Neural Networks

CS Coincidence System

DAQ Data Acquisition

DM Dark Matter

EOS Equations of State

ER Electronic Recoil

FAR False Alarm Rate

fuse Framework for Unified Simulation of Events

GCN Gamma–Ray Coordinates Network

IBD Inverse Beta Decay

INFN Istituto Nazionale di Fisica Nucleare

JINST Journal of Instrumentation

LAr/GAr Liquid/Gaseous argon

LNGS Laboratori Nazionali del Gran Sasso

LRT Likelihood Ratio Test

LXe/GXe Liquid/Gaseous xenon

LZ LUX-ZEPPLIN

MMA Multi-Messenger Astronomy

NC Neutral Current

NEST Noble Element Simulation Technique

NR Nuclear Recoil

PE Photoelectron

PMT Photomultiplier Tube

PNS Proto-Neutron Star

PR Precision Recall

PT Publishing Tools

QCD Quantum Chromodynamics

ReLU Rectified Linear Unit

ROI Region of Interest

RSG Red Super Giant

S1 Prompt Scintillation Signal

S2 Proportional Scintillation Signal

SCiMMA Scalable Cyberinfrastructure to support Multi–Messenger Astrophysics

snax Supernova analysis in xenon

SNEWS Supernova Early Warning System

SR Science Run

TPC Time Projection Chamber

UDP User Datagram Protocol

VOEvents Virtual Observatory Events

VUV Vacuum Ultraviolet

WIMP Weakly Interacting Massive Particle

XGB eXtreme Gradient Boost