INSTITUT POLYTECHNIQUE DE PARIS

NNT : 2023IPPAX002

Thèse de doctorat

Hunting for the Diffuse Supernova Neutrino Background with Super-Kamiokande: latest searches and prospects for the Gd era

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°626 École doctorale de l'Institut Polytechnique de Paris (EDIPP)
Spécialité de doctorat : Physique

Thèse présentée et soutenue à Palaiseau, le 17 janvier 2023, par

ALBERTO GIAMPAOLO

Composition du Jury :

Yves Sirois Directeur de recherche, LLR – École Polytechnique	Président
Sara Bolognesi Directrice de recherche, CEA Paris-Saclay	Rapporteur
Jean-Sebastien Real Directeur de recherche, LPSC – Université Grenoble Alpes	Rapporteur
Maria Cristina Volpe Directrice de recherche, APC – Université Paris Cité	Examinateur
Pascal Paganini Directeur de recherche, LLR – École Polytechnique	Directeur de thèse

Acknowledgements

The work in this thesis couldn't have been possible without the support of the wonderful people who accompanied me and lead me through these intense few years, inside the lab as well as outside of it.

First and foremost, my supervisor, Pascal Paganini. Research can be a winding maze, but one that we were able to navigate together. I've come to appreciate our relationship not just on a professional level, but on a personal one, too. Like a good marathon runner, you always pushed me to go the extra mile, cheering me on not as a student, but as a fellow researcher. I feel fortunate to have had you as my mentor. A heartfelt thank-you to my colleague, Sonia El Hedri, who patiently guided me in the first years of my thesis. As one of the pillars of the DSNB analysis at SK, many of the ideas in this work are borne of your suggestions, of our conversations, of our late-night discussions. I feel that you've not only formed me as a researcher, but you've also given me the confidence to stand by my own ideas and forge my own path.

I am also grateful to everyone at LLR who contributed to making the lab the warm, friendly working environment it is. Special thanks to Thomas Müller, head of the LLR Super-K group, who taught me much about neutrinos, but also about the subtleties of Pozzuoli Italian. A sincere thank-you as well to Margherita Avanzini, Olivier Drapier, Benjamin Quilain, and Emilie Maurice, who supported me the entire way with their experienced advice. The other PhD candidates and post-docs at the lab I worked and interacted with also contributed to a vibrant work atmosphere. Thank you to Alice Coffani, a dear friend I shared an office with for much of my thesis, and Felipe Garcia, colleague by day, climbing buddy by night. Thank you to the office mates I had the pleasure to work and hang out with daily in the latter half of my thesis: Andrew "Pigargue" Santos (thanks for bringing back a piece of America I didn't know I missed) and Antoine Beauchêne (my favorite almost-Breton). Thank you also to Andrés Muñoz, Jafaar Chakrani, Christine Quach, Lena Osu, Laura Bernard, Rudolph Rogly, and Viêt Nguyen. I've been lucky to count you among my friends during my thesis (and beyond!). I would also be remiss not to mention the miracle workers at LLR's IT department: Michel Lastes, Frederic Magniette, Michael Mellin, and Andrea Sartirana, who saved my work from technological ruin on more than one occasion.

Many thanks to everyone who helped me from the Super-Kamiokande group, advising me either from across the Channel, from the other side of the Ocean, or the other side of the globe. First, the SK collaborators who worked with me on the DSNB analyses on the SK-IV and SK-VI datasets, namely Yosuke Ashida, Masayuki Harada, Masataka Shinoki, Soniya Samani, Hiroyuki Sekiya, Ryosuke Akutsu, and Linyan Wan. The counsel and guidance from the SK leadership has also been invaluable: thank you to Micheal Smy and Yasuo Takeuchi (leading the low-energy

analysis group), Mark Vagins (the mastermind behind the Sk-Gd project), and for the advice provided by Yusuke Koshio, Yasuhiro Nakajima, and the Super-Kamiokande spokesperson, Masayuki Nakahata.

I would like to express my sincere gratitude to the members of the jury: Sara Bolognesi, Jean-Sebastien Real, Maria Cristina Volpe, and Yves Sirois, for your precious feedback and insightful discussions.

Heartfelt thanks also to all my friends from outside the lab that accompanied me. Thanks to Marie, David, Giulia, Sjoerd, Guillaume, Kaustuv, Rens, and Sophie, for keeping in contact all these years. Thanks also to Luci for our many hikes around Île de France, Megan for an excellent DnD campaign, and to Jonas, fellow university classmate, colleague, housemate, and friend, who shared with me the highs and lows of the research life.

Finally, I am deeply grateful to my family. In particular, to my parents, my brother Michele, my aunt Pina, and my cousin Giulia, for your unwavering support. It wouldn't have been possible without you!

Abstract

The Super-Kamiokande experiment, located in Japan, allows for the study of neutrinos from a variety of sources. One of the interests of the experiment is the study of astrophysical neutrinos, and in particular of neutrinos produced during core-collapse supernovae, marking the death of massive stars. In 1987, Super-Kamiokande's predecessor experiment detected a dozen neutrinos originating from a close-by supernova in the Large Magellanic Cloud. The 1987 supernova remains the only instance of supernova neutrino detection to date. Indeed, supernovas that are near enough for neutrino detection are a rare occurrence: roughly three supernovae per century are expected in the Milky Way. In addition to direct supernova detection, one of the main challenges for the experiment is the first detection of the integrated flux from all past supernovae in the observable universe. Detection of this flux, also known as Diffuse Supernova Neutrino Background (DSNB), would allow for the improvement of our understanding of supernovae, as well as the refinement of current cosmological models, providing hints about the history of star formation, about nucleosynthesis, and about the evolution of the universe.

In this thesis, an analysis searching for the DSNB using Super-Kamiokande is presented. In this water Cherenkov detector, the main detection channel for supernova neutrinos is the inverse beta decay, in which an electron antineutrino interacts with a proton, producing a positron and a neutron. Particular attention is given to the background reduction obtained with the detection of the neutron in coincidence with the positron. The analysis of the data from the SK-IV period (between September 2008 and June 2018) allowed for the determination of upper limits on the DSNB flux that are currently the world's tightest, excluding several optimistic models. The second component of this thesis consists in extending the analysis to the most recent data, after doping the detector water with Gadolinium sulfate in 2020 (also known as the SK-Gd phase), significantly strengthening the signal of a neutron in the detector. This work shows that the sensitivity of an SK-Gd search, after just three years of runtime (i.e. in mid-2023), should be statistically equivalent to that obtained with SK-IV, paving the way for discovery of the DSNB in the medium term. This thesis also offers an early look at new SK data with Gd, showing an improved sensitivity to the DSNB, and an increased significance of the observed excess.

Résumé en français

Aujourd'hui, l'un des axes de recherche actuels de l'astrophysique est l'étude des supernovae, ces explosions cataclysmiques qui provoquent l'effondrement des étoiles massives. Ces événements sont riches en informations et nous fournissent des indices sur la dynamique d'évolution des étoiles, sur leur composition élémentaire et sur le taux de formation et de mort stellaire dans l'univers visible. L'observation optique des supernovae a une longue histoire, avec des enregistrements d'événements de supernova remontant à l'antiquité. Mais l'observation optique est loin d'être la fin de l'histoire. En effet, les supernovae sont des accélérateurs de particules à haute énergie propres à la nature, éjectant toutes sortes de particules élémentaires au cours de leurs explosions. Les neutrinos qui en sortent sont particulièrement intéressants à étudier, car ils constituent la grande majorité de l'énergie émise et, contrairement aux photons, ils traversent directement la majeure partie de la matière dense de la supernova, nous renseignant ainsi sur les couches les plus profondes de l'étoile dans ses derniers instants.

Le domaine de l'observation des neutrinos de supernova s'ouvrit avec la première (et jusqu'à présent la seule) observation directe d'une supernova avec des détecteurs de neutrinos en 1987. Depuis, l'apparition répétée d'une supernova suffisamment proche pour être à la portée des capacités actuelles de détection des neutrinos est un événement très attendu. Mais nous ne devrons pas nécessairement attendre une supernova proche pour détecter d'autres neutrinos de supernova. En effet, le flux de toutes les supernovae passées dans l'univers observable, une fois additionné, devrait être suffisamment élevé pour être détecté par les expériences de neutrinos modernes. Ce type d'observation serait très différent de l'observation directe d'une seule supernova. Outre le décalage vers le rouge induit par le flux provenant de sources lointaines, nous aurions affaire à un flux faible, isotrope et constant, d'où son nom de fond diffus de neutrinos de supernova (Diffuse Supernova Neutrino Background en anglais, ou DSNB). Dans le cadre de cette thèse, nous nous intéressons à la possibilité de détection par l'expérience Super-Kamiokande (SK), qui est actuellement le plus grand détecteur Cherenkov à eau du monde. Depuis ses débuts, Super-Kamiokande a été à la pointe de l'astronomie des neutrinos.

Déjà avant cette thèse, SK a placé les limites supérieures les plus strictes sur le flux DSNB, dans la gamme d'énergie accessible au détecteur. L'objectif de cette thèse est donc de faire progresser les efforts de détection du DSNB avec SK, à la lumière de nouveaux échantillons de données et de nouvelles capacités du détecteur. Tout d'abord, une introduction des concepts physiques pertinents et du détecteur est donnée dans les chapitres 1 et 2. La stratégie de réduction des données pour une recherche utilisant l'ensemble de données SK-IV est ensuite présentée au chapitre 3. En utilisant le canal d'interaction de la désintégration bêta inverse, les neutrinos

DSNB ont été recherchés dans le volume fiduciel d'eau de 22,5 ktonnes du détecteur en utilisant la période de données SK-IV. Dans le cadre de cette analyse, la stratégie d'identification du positron IBD et, en particulier, du signal de neutron de coïncidence est détaillée. La détection des neutrons, possible pour la première fois dans SK à partir de SK-IV, est une caractéristique marquante de cette recherche. Ce travail utilise un modèle BDT pour maximiser notre capacité à détecter le faible signal de capture de neutrons de 2,2 MeV sur l'hydrogène. L'efficacité de détection des neutrons obtenue pour les captures de neutrons dans l'eau pure de SK-IV se situe entre 20 et 30 %, pour une réduction du bruit de fond accidentel de 3 à 4 ordres de grandeur.

Le traitement statistique de l'échantillon complet de données SK-IV est ensuite discuté au chapitre 4. Un traitement statistique spécifique, s'appuyant sur un ajustement spectral des données, a été développé pour tirer parti du signal neutronique, tout en étant capable de produire une mesure combinée utilisant l'ensemble des données SK jusqu'à SK-IV. Une série de modèles théoriques modernes de flux DSNB ont été testés avec une exposition totale de $22,5 \times 5823$ kton-jour. Un excès d'environ $4,6^{+2,5}_{-2,3}$ événements IBD/an (selon le modèle DSNB) a été trouvé dans l'ensemble de données SK-IV. En combinant cette observation avec les ajustements spectraux sur les périodes antérieures de SK, on obtient des limites supérieures sur le flux DSNB pour E_ν>17,3 MeV d'environ $2,6\ \nu_e/\mathrm{cm}^2/\mathrm{an}$, à un niveau de confience de 90%. De plus, la sensibilité de l'ajustement combiné s'avére comparable à plusieurs modèles de flux modernes et réalistes, ce qui permet d'envisager avec optimisme l'avenir du signal DSNB. Cette recherche représente la première analyse de SK prenant en compte l'ensemble des données de SK-IV et combinant son observation avec toutes les périodes précédentes de SK.

Après avoir discuté les résultats de la recherche et leurs implications sur les modèles modernes de DSNB, nous passons à l'examen des perspectives du DSNB dans l'avenir de l'expérience. Plus précisément, les perspectives du DSNB en ce qui concerne le détecteur SK nouvellement amélioré, qui, ayant été dopé au sulfate de gadolinium en 2020, voit sa sensibilité au DSNB augmenter de façon significative. Le mécanisme et l'étendue de ces gains sont étudiés au chapitre 5. Avec l'avènement de l'ère du Gd dans la prise de données SK, le signal de capture neutronique s'améliore en fonction de la concentration du gadolinium dans l'eau. L'experience a commencé par une concentration de 0,01% dans SK-VI et a atteint une concentration de 0,03% dans SK-VII, ce qui actuellement en cours. SK est donc prêt à augmenter de manière significative sa sensibilité aux signaux neutroniques, et en particulier au DSNB. L'analyse SK-IV prépare effectivement le terrain pour les analyses sensibles SK-Gd, en fournissant un schéma directeur pour la stratégie des recherches à venir. Le potentiel de l'analyse à la lumière du dopage au Gd a été directement évalué à l'aide de nouvelles études de sensibilité. L'amélioration de la capacité de SK à détecter les captures de neutrons est testée, en observant l'efficacité du signal pour les captures de neutrons sur le Gd, par rapport aux captures sur l'hydrogène, pour des coupures réalistes d'identification des neutrons.

La manière dont ce gain d'efficacité s'applique aux différentes époques du SK est étudiée, en projetant les performances de l'analyse sur l'avenir de l'expérience. Il s'agit de la première étude évaluant la sensibilité du SK-Gd au DSNB dans le cadre d'une analyse réaliste, tenant pleinement compte de l'impact des incertitudes systématiques. Les résultats montrent que SK peut atteindre la sensibilité combinée de l'analyse spectrale SK-IV en seulement trois ans de fonctionnement, à la mi-2023. À la fin du projet SK-Gd, la sensibilité d'une analyse SK combinée couvrira probablement une grande partie de l'espace des paramètres, balayant les modèles de flux les plus optimistes. Des progrès réalistes dans notre traitement des incertitudes systématiques

permettront probablement d'améliorer encore plus ces perspectives.

Enfin, le chapitre 6 présente une première analyse du jeu de données SK-VI, obtenu avec le détecteur nouvellement amélioré. Bien que la statistique soit encore limitée, une sensibilité accrue a été démontrée, et la significance de l'excès précédemment observé a été renforcée. Cette thèse espère donc servir de tramplin pour de futurs efforts et études sur la voie d'une première détection du DSNB.

Dans le cadre de SK, les pistes à explorer sont encore nombreuses. L'amélioration des incertitudes systématiques clés dans l'analyse est discuté. En particulier, la réduction de l'incertitude systématique associée aux fonds atmosphériques NCQE améliorera considérablement notre sensibilité à la gamme d'énergie inférieure du flux DSNB et nous permettra de tirer pleinement parti de l'efficacité de détection des neutrons plus élevée du SK-Gd. Une autre incertitude statistique importante à maîtriser sera celle de la multiplicité des neutrons des fonds atmosphériques, qui peut limiter notre compréhension de la façon dont ces fonds répondent aux coupures d'identification des neutrons. Avec SK-Gd, la structure et la stratégie de l'analyse DSNB pourraient être revues afin d'exploiter davantage le signal neutronique. Les performances des coupures de spallation s'amélioreront probablement en même temps que la visibilité des neutrons secondaires associés au dépôt d'énergie du muon de spallation. SK-Gd ouvre également la voie à un traitement spécifique des fonds atmosphériques pour tenter de discriminer directement les neutrons secondaires provenant de ces fonds, au lieu de discriminer uniquement les neutrons de type IBD par rapport aux coïncidences accidentelles. La simulation de spallation, nouvellement disponible pour SK, améliorera aussi considérablement notre capacité de comprendre, de modéliser et de traiter la spallation, et sera essentielle pour réaliser le plein potentiel de l'expérience. Enfin, des algorithmes plus sophistiqués tels que les modèles de Deep Learning pourraient contribuer à l'améliorations des performances de détection des neutrons, à une meilleure discrimination des fonds de type NCQE, et à une reconstruction combinée du positron et du neutron dans l'événement IBD complet. Nombre de ces recherches sont déjà en cours au sein de la collaboration. Les perspectives pour l'avenir de la recherche du DSNB à Super-Kamiokande sont donc très prometteuses. Cette thèse est donc l'une des premières briques d'un avenir riche en possibilités.

Contents

Introduction

E quando miro in cielo arder le stelle;
Dico fra me pensando:
A che tante facelle?
Che fa l'aria infinita, e quel profondo
Infinito seren? Che vuol dir questa
Solitudine immensa? Ed io che sono?

Giacomo Leopardi

The mysteries of the night sky have inspired generations of artists, thinkers, poets and scientists, for as long as we've been able to admire it. To this day, despite awesome leaps in our understanding of cosmology, of astrophysics, of stellar formation, evolution, and death, much is still unknown about the laws governing the behavior of the stars. One of the thrusts of current investigations is the study of supernovae, the cataclysmic explosions bringing about the collapse of massive stars. These events are rich with information, providing us with clues about a star's evolution dynamics, its elemental composition, and about the rate of stellar formation and death in the visible universe. The optical observation of supernovae has a long history, with records of supernova events being observed dating back to antiquity. But optical observation is far from the end of the story. Supernovae are in fact Nature's high-energy particle accelerators, ejecting all manners of elementary particles in their outburst of energy. Of particular interest are neutrinos, which make up the vast majority of the energy emission, and—unlike photons—travel straight through most of the dense supernova matter, carrying with them information about the deepest layers of the star in its final moments.

The field of supernova neutrino observation opened with the first (and, so far, only) direct observation of a supernova with neutrino detectors in 1987. Ever since, the repeated occurrence of a supernova close enough to be in reach of current neutrino detection capabilities is a much-awaited event. But we won't necessarily have to wait for a close-by supernova to detect more supernova neutrinos. In fact, the flux from all past supernovae in the observable universe, when summed up, should be high enough to be detectable in modern neutrino experiments. This type of observation would differ in important ways from the direct observation of a single supernova. In addition to the redshift incurred by the flux from far-away sources, we would deal with a faint, isotropic, and constant flux, motivating its name as the *Diffuse Supernova Neutrino Background* (DSNB). For the purpose of this thesis, we are interested in the possibility of detection with the Super-Kamiokande (SK) experiment, currently the world's largest water Cherenkov detector. Since the beginning, Super-Kamiokande has been at the forefront of neutrino astronomy.

Already before this thesis work, SK placed the tightest upper limits on the DSNB flux, in the energy range accessible to the detector. The goal of this thesis is then to advance the DSNB detection efforts with SK, in light of new data samples and new detector capabilities. First, an introduction of the relevant physical concepts and of the detector are given in Chapters 1 and 2. The data reduction strategy

for a search using the SK-IV dataset is then presented Chapter 3. Here, particular attention is given to the detection of neutrons at SK, a new possibility with SK-IV. Then, the statistical treatment of the full SK-IV data sample is discussed in Chapter 4. This search represents the first SK analysis considering the full SK-IV dataset and combining its observation with all previous SK periods. After discussing the results of the search and their implications on modern DSNB models, we move to considering the outlook of the DSNB in the future of the experiment. Specifically, the prospects of the DSNB with respects to the newly upgraded SK detector, which, having been doped with Gadolinium sulfate in 2020, will see significant gains in its sensitivity to the DSNB. The mechanism for and extent of these gains are explored in Chapter 5. Finally, Chapter 6 gives an early analysis of the SK-VI dataset, obtained from the newly-upgraded detector. This work will then give the reader a picture of the current status of the DSNB at SK, as well as its direction and future outlook.

1
Neutrino physics

The observation of supernovae with neutrinos is a game of contradictions. One must look underground to gaze at the stars. One must search for the tiny to find the immense. To sufficiently introduce the theory relevant to this thesis, we must therefore explore the two fields of physics that concern themselves with the two opposite extremes of our universe: namely, particle physics and astrophysics. We can start our exploration from the extremely small world of elementary particles, of which the neutrino is an extremely light inhabitant. Indeed, out of the massive elementary particles known today, neutrinos are by far the lightest. First postulated in 1930, they allowed Wolfgang Pauli to explain how the β-decay in nuclear physics could agree with the conservation of energy and momentum. As we now know, the neutrino is abundantly present on Earth and across the universe. Despite this abundance, its weak interaction with matter and its dainty mass make it a particularly elusive particle to detect. The fact that we can detect a neutrino at all is owed in no small part to the great strides in theoretical neutrino physics that have been made since Pauli's intuition. Our current understanding of neutrinos—and of particle physics as a whole—is founded upon the Standard Model (SM). To give the reader a broad overview of the theoretical backing of this thesis, this is where we'll start. The SM is a *gauge theory*, built around the idea that certain local transformations leave the dynamics of our universe unchanged. In the framework of a gauge theory, one can theorize the mathematical *symmetries* that are inherent in nature, and identify the corresponding *conservation laws*. This kind of theory now forms the foundation of modern theoretical physics.

1.1 Neutrinos in the Standard Model

The beginnings of the SM can be traced back to the development of quantum electrodynamics (QED) in the 1920s by Paul Dirac, describing the dynamics of electrically-charged particles and their interaction with the photon. It wouldn't be until the 1940s that foundational papers by Feynman, Schwinger, Tomonaga and Dyson [1, 2, 3, 4] would flesh out a full formulation of QED as a gauge theory. In QED, the electron and other electrically-charged particles are mathematically described as fields, invariant under symmetry transformations belonging the U(1) group. The electromagnetic force between electrically-charged particles is then mediated by the photon field, with the consequent conservation of electrical charge. The idea of explaining fundamental forces with fundamental symmetry transformations was then applied by Yang and Mills in 1954 [5] to kickstart the quantum description of the strong force between nucleons. Later, the quantum theory of the strong force—quantum chromodynamics (QCD)—would be more precisely understood as acting between *quarks*. This time, the quark fields obey an SU(3) symmetry, and the *gluon* field is introduced to mediate the strong force, with conservation of color charge. The third fundamental force to

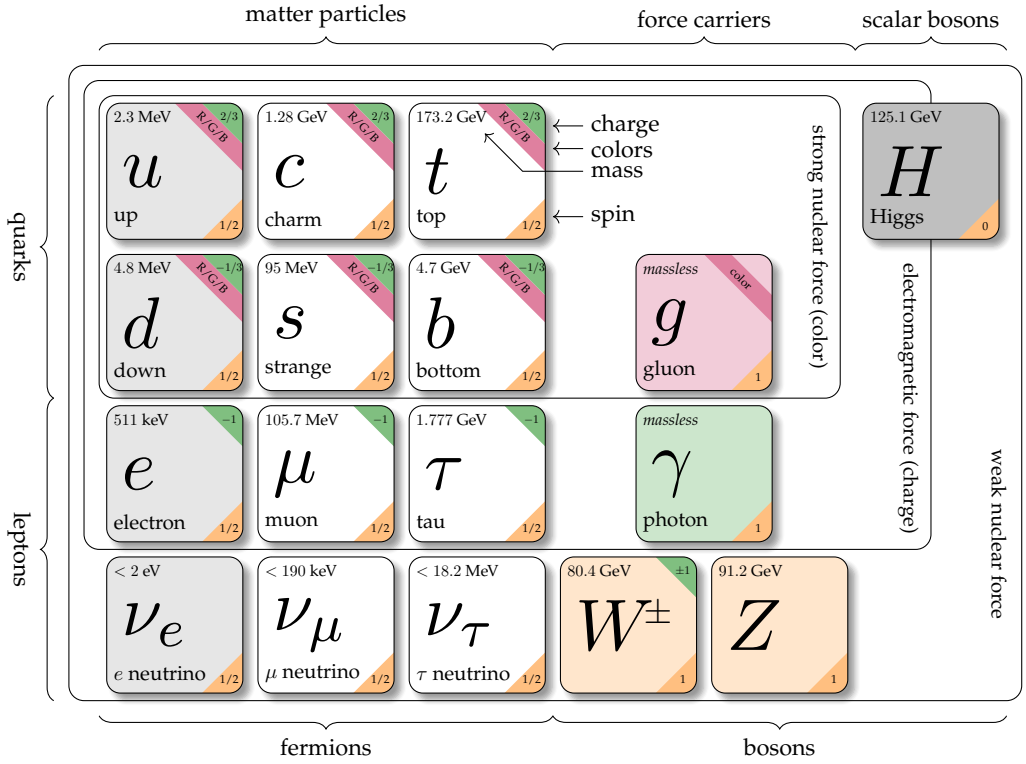

FIGURE 1.1: The particles of the Standard Model, with their respective mass, electric charge, color charge, and spin. Each column of matter particles represents one of the three generations, with the first generation (shaded) containing the lightest matter. Since the light quarks and charged leptons are the most stable, they are also the most common particles found in ordinary matter, along with the electron neutrino, which participates in many interactions with the electron. Matter particles are also grouped together according to which interaction they experience, with neutrinos being subject only to the weak force. For each matter particle, the SM also contains a corresponding antiparticle, with equal mass and spin but opposite electric and color charge. Illustration adapted from Ref. [8].

receive the gauge treatment was the weak force, introduced to explain the β-decay in nuclear physics. First theorized by Fermi in 1933, the weak interaction is the only one directly relating the neutrino with other particles of the SM. It would be thanks to Glashow, Salam and Weinberg [6, 7] that the weak force would be unified with the electromagnetic force in a single theory, founded on an SU(2)×U(1) symmetry. Three particle fields mediate the weak force: the W^+, W^-, and Z bosons. The final component and linchpin of the SM in its current form is the Higgs boson. Crucially, the presence of the Higgs boson and the associated Higgs mechanism brings about the generation of non-zero masses for the W and Z bosons. Before its introduction, physicists could not fully reconcile the prediction of massless bosons with the precise measurement of the W and Z masses at CERN. The existence of the Higgs field would be confirmed in 2012 by CERN [9, 10] to much international excitement. For all its successes, it must be noted that the Standard Model is far from a complete quantum theory of nature: it does not explain the gravitational force in any capacity, and its particle content only covers ∼5% of the energy composition of the universe (i.e., it does not offer an explanation for dark matter nor dark energy). Nonetheless, the SM represents a triumph of modern particle physics, being able to make precise, observable predictions that have so far found excellent agreement with experimental projects.

The full SM is then based on a SU(3)×SU(2)×U(1) symmetry group. Its particle content is summarized in Fig. 1.1. The existence of three matter particle generations—identical to each other in all respects save their mass—has been established by empirical observation [11]. The finding that neutrinos carry mass, running contrary to earlier expectations, is the latest modification of our understanding of the model. In total, the SM has 26 free parameters (including all nonzero masses) that are not predicted by the theory and must be fixed by experimental measurement. In the following introduction, a brief overview of SM aspects relevant to neutrino physics is given. Specifically, how the neutrino interacts with other matter, the phenomenon of neutrino oscillation, and how the latter is modified in the presence of matter. In the next section, a sketch of the main neutrino source relevant to this work—indeed, its *raison d'être*—supernova neutrinos.

The electroweak interaction

The interactions governed by the electroweak sector of the SM can be grouped into charged current (CC) interactions mediated by the $W^{\pm}$, and neutral current (NC) interactions mediated by the γ and Z. Hence the interaction electroweak Lagrangian can be written as [12]:

$$\mathscr{L}_I = \mathscr{L}_{CC}^{(W)} + \mathscr{L}_{NC}^{(Z)} + \mathscr{L}_{NC}^{(\gamma)} \tag{1.1}$$

A peculiarity of the weak interaction is that it treats left-chiral particles differently from right-chiral particles. Therefore, explicit interaction terms must specify the chirality of the particles. Focusing on the leptons (e_L, e_R, and the corresponding ν_L), we can expand the interaction Lagrangian in terms of the W, Z, and A (photon) fields as:

$$\mathscr{L}_{CC}^{(W)} = -\frac{g}{\sqrt{2}}\{\overline{\nu_{eL}}\not{W} e_L + \overline{e_L}\not{W}^{\dagger}\nu_{eL}\} \tag{1.2}$$

$$\mathscr{L}_{NC}^{(Z)} = -\frac{g}{2\cos\theta_{\mathrm{W}}}\{\overline{\nu_{eL}}\not{Z}\nu_{eL} - (1-2\sin^2\theta_{\mathrm{W}})\overline{e_L}\not{Z}e_L + 2\sin^2\theta_{\mathrm{W}}\overline{e_R}\not{Z}e_R\} \tag{1.3}$$

$$\mathscr{L}_{NC}^{(\gamma)} = g\sin\theta_{\mathrm{W}}\overline{e}\not{A}e \tag{1.4}$$

where g, θ_{W} are free parameters, the *Feynman slash* operator acts on each field as $\not{A} = \gamma^{\mu}A_{\mu}$ (with γ^{μ} being the *Dirac γ matrices*) and the dagger (†) denotes the Hermitian conjugate. These interaction terms generate vertices of the following type:

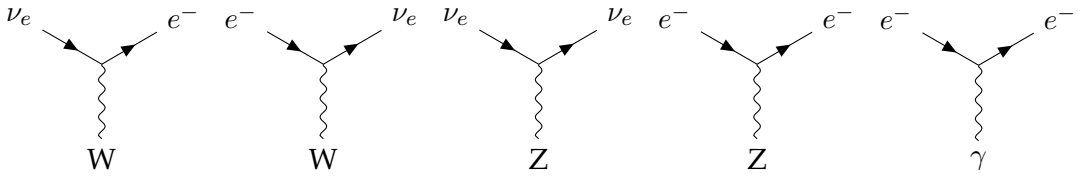

as well as the corresponding antiparticle vertices, obtained with the replacement $\nu_e \rightarrow \overline{\nu}_e$ and $e^- \rightarrow e^+$ and inverting the fermion direction. Similar terms allow quarks to scatter through photon and weak neutral currents. The interaction with the W induces a change in quark flavor, preserving, to first order, the quark generation (i.e. favoring $u \leftrightarrow d$, $t \leftrightarrow b$, $c \leftrightarrow s$ transitions). These interactions are essential in nuclear physics, as they allow flavor-changing interactions between leptons and nucleons (themselves bound states of quarks). As we will see, particularly relevant for this work is the *inverse beta decay*, (IBD) wherein an electron antineutrino can interact with a proton (a uud bound state), yielding in the final state a positron and a neutron (a ddu bound state):

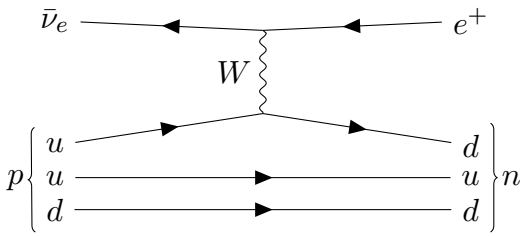

(1.5)

This interaction is a canonical detection channel for neutrino experiments, owing to a large cross-section and an abundance of protons in a typical detector.

Neutrino mixing and oscillations

As it turns out, the neutrino eigenstates associated with interaction (ν_e, ν_μ, ν_τ), exhibiting well-defined flavor, are not the same as those associated with propagation in vacuum (ν_1, ν_2, ν_3), exhibiting well-defined masses (m_1, m_2, m_3). The two are related by the vacuum mixing matrix U:

$$|\nu_\alpha\rangle = \sum_k U^\star_{\alpha k}|\nu_k\rangle \qquad (\alpha = e, \mu, \tau; k = 1, 2, 3) \tag{1.6}$$

A neutrino in the flavor eigenstate $|\nu_\alpha\rangle$ at $t = 0$ therefore turns into a superposition of flavor eigenstates at $t > 0$ according to

$$|\nu_\alpha(t)\rangle = \sum_\beta \left(\sum_k U^\star_{\alpha k} e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle \tag{1.7}$$

We can adopt an ultra-relativistic neutrino approximation, wherein $E_k \approx E + \frac{m_k^2}{2E}$ and $t \approx L$, where L is the travel distance at time t. In this case, the probability of *oscillation* (i.e., of transition between different neutrino flavor states) in vacuum is given by

$$P_{\alpha\to\beta}(L, E) \equiv |\langle\nu_\beta|\nu_\alpha(t)\rangle|^2 = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \exp\left(-i\frac{\Delta m^2_{kj} L}{2E}\right) \tag{1.8}$$

Here, $\Delta m^2_{kj} \equiv m^2_k - m^2_j$. Therefore, the very existence of neutrino oscillations, first observed by the Super-Kamiokande collaboration in 1998 [13], implies a non-zero mass difference between mass eigenstates, proving the existence of massive neutrinos. The full vacuum mixing matrix, also known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, describing the mixing of all three flavor states, can be parametrized by three angles and a phase (θ_1, θ_2, θ_3, δ) and is associated with three separate mass squared differences (Δm^2_{21}, Δm^2_{23}, Δm^2_{13}, of which only two are independent parameters). In practice, a useful approximation is the two-neutrino mixing case, in which only two neutrino mass eigenstates are considered as taking part in oscillation. In this case, the effective mixing matrix $U_{2\nu}$ can be parametrized by a single angle θ, corresponding to a single squared mass difference Δm^2:

$$U_{2\nu} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{1.9}$$

$$P_{\alpha\to\beta}(L, E) = \sin^2 2\theta \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \tag{1.10}$$

This approximation is particularly useful since many neutrino experiments and analyses are not sensitive enough to measure the effects of the full three-neutrino oscillation (or, more precisely, are most sensitive to one type of oscillation). In a typical oscillation experiment, the amplitude of the oscillation therefore informs us about the parameters of the mixing matrix that we are sensitive to. Meanwhile, the phase of the oscillation provides us with the absolute value of the squared mass difference.

Indeed, such measurements do not directly measure the neutrino masses themselves, nor their relative size, since Eq. 1.10 is also independent of the sign of Δm^2. In experiments studying neutrinos from the sun, where the θ_{12} mixing dominates, the absolute mass squared difference has been determined to be $|\Delta m^2_{21}| \approx 7.5 \times 10^{-5}\text{eV}^2$ [14]. In experiments studying neutrinos produced in the atmosphere θ_{23} dominates, with the finding $|\Delta m^2_{23}| \approx 2.4 \times 10^{-3}\text{eV}^2$ [14]. Currently, no precise measurement of the neutrino masses themselves is available, though the most recent limit on the absolute neutrino mass scale was placed by the KATRIN collaboration at < 0.8 eV, at a 90% confidence level in 2022 [15]. Because of the ambiguity surrounding the values of the neutrino masses, two general mass configurations are possible. Either $m_1 \approx m_2 \lesssim m_3$, in which case we talk of a *normal hierarchy*, or $m_3 \lesssim m_1 \approx m_2$, where we talk of an *inverted hierarchy*.

Oscillations in matter

The behavior of neutrinos in matter is altered by the coherent elastic scattering of the neutrinos on the stable constituents of the matter itself (i.e., on electrons and nucleons). In particular, neutrino oscillation is affected by the CC scattering process $\nu_e + e^- \to e^- + \nu_e$. Scattering processes due to NC interactions can be shown to have no impact on oscillation probability in the three-flavor paradigm. The CC interaction modifies the Hamiltonian of neutrino propagation in vacuum $\mathscr{H}_0$, with the effective potential V_{CC} yielding the effective matter Hamiltonian $\mathscr{H}_{\text{eff}}$ for the $\nu_\alpha \to \nu_\beta$ transition amplitude $\psi_{\alpha\beta}$:

$$\mathscr{H}_{\text{eff}}\psi_{\alpha\beta} = \left(\mathscr{H}_0 + V_{CC}\delta_{\alpha e}\delta_{\beta e}\right)\psi_{\alpha\beta} \tag{1.11}$$

$$V_{CC} = \sqrt{2}G_F N_e. \tag{1.12}$$

Here, G_F is the Fermi coupling constant, and N_e is the electron number density in the medium. Shifting our focus on the two-neutrino case for simplicity, where the mixing is between ν_e, ν_μ and ν_1, ν_2, we write

$$\mathscr{H}_{\text{eff}} = \left(\frac{\Delta m^2}{4E}\right)\begin{pmatrix}\cos 2\theta & \sin 2\theta \\ -\sin 2\theta & \cos 2\theta\end{pmatrix} + \begin{pmatrix}V_{CC} & 0 \\ 0 & 0\end{pmatrix}. \tag{1.13}$$

To ease our discussion, we can apply a constant shift to the Hamiltonian, leaving the dynamics of our system unchanged:

$$\mathscr{H}_{\text{eff}} \to \mathscr{H}_{\text{eff}} - \begin{pmatrix}V_{CC}/2 & 0 \\ 0 & V_{CC}/2\end{pmatrix} \tag{1.14}$$

$$\mathscr{H}_{\text{eff}} = \left(\frac{\Delta m^2}{4E}\right)\begin{pmatrix}\cos 2\theta + A_{CC} & \sin 2\theta \\ -\sin 2\theta & \cos 2\theta - A_{CC}\end{pmatrix} \tag{1.15}$$

where $A_{CC} = \frac{2EV_{CC}}{\Delta m^2}$. We can now define a new basis of states diagonal in $\mathscr{H}_{\text{eff}}$. This basis is then associated with the matter parameters θ_M and Δm^2_M:

$$\mathscr{H}_{\text{eff}} = \left(\frac{\Delta m^2_{\text{M}}}{4E}\right)\begin{pmatrix}\cos 2\theta_{\text{M}} & \sin 2\theta_{\text{M}} \\ -\sin 2\theta_{\text{M}} & \cos 2\theta_{\text{M}}\end{pmatrix} \tag{1.16}$$

$$\Delta m^2_{\text{M}} = \Delta m^2 \sqrt{(\cos 2\theta - A_{CC})^2 + (\sin 2\theta)^2} \tag{1.17}$$

$$\tan 2\theta_{\text{M}} = \frac{\tan 2\theta}{1 - \frac{A_{CC}}{\cos 2\theta}} \tag{1.18}$$

The effect of matter interactions on neutrino oscillations can be therefore interpreted as a modification of the mixing angle and the neutrino masses. For a material with constant electron density, $\frac{\mathrm{d}\theta}{\mathrm{d}x} = 0$, Equation 1.18 implies the existence of a resonance, in which maximal flavor mixing is reached with $\theta_{\text{M}} = \frac{\pi}{4}$, for the neutron density

$$N_e^{\text{res}} = \frac{\Delta m^2 \cos 2\theta}{2\sqrt{2} E G_F}. \tag{1.19}$$

This resonant mixing, named the *MSW effect* after Mikhev, Smirnov and Wolfenstein [16, 17] is important in explaining the propagation of neutrinos in dense matter such as inside stellar cores such as the Sun's, and in the inner regions of supernovae. In particular, such a resonance can induce maximal mixing even if the vacuum mixing angle is small, significantly altering predictions of neutrino fluxes from stellar and supernova sources. Observation of the flavor conversion resulting from the MSW effect allows us to access the sign of Δm^2, assuming that the quadrant of the relevant mixing angle is known.

1.2 Supernova neutrinos

We now switch gears, turn our eyes to the sky and expand the scale of our discussion to the rich and complex world of stellar evolution. We focus in particular on the neutrino source of primary importance for this thesis: supernovae. Supernovae (SNe) are dramatic explosions, typically brighter than their host galaxy, that mark the death of massive stars. In an outburst of energy of the order of 10^{58} MeV, several solar masses are violently expelled, possibly leaving behind a neutron star or a black hole as remnant. In this section, an introduction to our current understanding of supernova dynamics is presented, with particular attention to the production and propagation of neutrinos. After this introduction, the observational target of this work is discussed, namely, the Diffuse Supernova Neutrino Background.

Supernova explosion dynamics

When discussing SNe, one will frequently come across the typical SN classification used in observational astronomy, based on the SN's optical spectrum (which is related to the elemental content of the ejected material). Accordingly, we talk of a *Type I* SN if the optical spectrum does not display any Hydrogen absorption lines; otherwise, we talk of a *Type II* SN. Type I SNe are further classified into *Type Ia* (showing a strong ionized silicon line), Type Ib (no strong silicon line, but a strong non-ionized helium line), and *Type Ic* (no strong lines associated with either ionized silicon or non-ionized helium). Type II SNe can be further classified according to peculiarities in their spectrum. In particular, we find *Type IIn* SNe (with unusually narrow hydrogen lines) and *Type IIb* (whose spectrum changes over time from a typical Type II spectrum to a helium-rich one, strongly resembling a Type Ib spectrum).

For the purposes of this text, we can also draw a distinction between two broad categories, according to the mechanism triggering the explosion: thermonuclear and core-collapse SNe. The thermonuclear SN, corresponding to a type Ia SN, is the result of runaway fusion reactions inside a white dwarf (i.e., a compact star supported by the degeneracy pressure of its electrons). When a white dwarf is part a binary system together with another star, it can gradually accrete mass spilling from its companion. Before being able to reach its Chandrasekhar mass (when the gravitational pressure would overcome the electron degeneracy pressure), the increasing mass causes the density and temperature of the dwarf to increase, inducing carbon fusion reactions. Since, unlike main-sequence stars, white dwarfs are not supported by fusion pressure, the increasing fusion rate does not cause the star to expand and cool down. We talk then of *runaway* fusion reactions, whose energy release is eventually large enough to unbind the star with a violent shockwave. These kinds of SNe are however not significant sources of neutrinos and are therefore not the focus of neutrino physicists.

Rather, we are interested in the core-collapse SN (corresponding to a SN of Types Ib, Ic, or II), which is responsible for the generation of an enormous flux of all types of neutrinos. Unlike the thermonuclear SN, it avoids runaway fusion and reaches the Chandrasekhar limit of the progenitor star. The core-collapse SN is then initiated by the gravitational collapse of a star with a mass between 8-9 $M_\odot$ and 40-60 $M_\odot$ ($M_\odot \approx 2 \times 10^{30}$ kg is the solar mass). In the ensuing explosion, neutrinos carry away as much as 99% of the gravitational binding energy, producing $\sim 10^{58}$ neutrinos with average energy of ~ 10 MeV. Unlike thermonuclear SNe, core-collapse SNe leave behind a compact remnant, in the form of a neutron star or black hole. In addition, since such SNe can be initiated in stars with a wide variety of masses, metal contents, and star envelope structures, the visible effects of the explosion are also widely variable.

In the early stages of the SN, the progenitor star's inner core is composed of mostly iron, which, as the element with the highest binding energy, cannot be used in further fusion reactions. This reduces the star's ability to support its own gravitational pressure. Still, the star resists collapse thanks to the electron degeneracy pressure. In the *capture phase*, as the core temperature rises, electrons in core get captured by nuclei and free protons,

$$e^- + \mathcal{N}(\mathrm{Z}, \mathrm{A}) \to \mathcal{N}(\mathrm{Z}-1, \mathrm{A}) + \nu_e \tag{1.20}$$

$$e^- + p \to n + \nu_e \tag{1.21}$$

In this process most of the electron energy is carried away by the electron neutrino, for whom at this stage the core is still transparent. In addition, photodissociation processes, in which the iron is broken down by energetic photons,

$$\gamma + {}^{56}\mathrm{Fe} \to 13\alpha + 4n, \tag{1.22}$$

further reduce the energy of the core. As the energy available in the core drops, together with the number of electrons, so does the electron pressure, and with it, the Chandrasekhar limit. When the latter falls below the mass of the iron core, collapse is inevitable.

As the core shrinks and becomes denser, it eventually becomes opaque even to neutrinos, as their mean-free-path becomes smaller than the core size. At this stage, as neutrinos are produced they are effectively trapped inside the core. Eventually, the inner core reaches nuclear density, at which point the collapse is decisively halted by the degeneracy pressure of the neutrons: a proto-neutron star is born, roughly

10 km in radius. At its surface, a powerful shockwave is formed by the infalling material suddenly rebounding outwards. The shock propagates outwards, traveling through the inner and outer regions of the core. As further infalling material from the outer core crosses the shockwave front, it slows down and accretes on the inner proto-neutron star, stabilizing it. Photodissociation processes leave free nucleons in the shock's wake, inducing copious production of electron neutrinos through electron capture. While the shockwave front is in the densest regions of the star, these neutrinos remain trapped and accumulate behind the front. As the shock progresses outwards, eventually a low enough medium density is reached from which the neutrinos can finally escape. This causes a sudden, brief electron neutrino burst, sometimes called the *neutronization burst*, at the beginning of the supernova event. Despite its relatively high luminosity, this burst is very short-lived, and as such does not represent the main source of neutrino emission.

In the dense SN matter, a variety of processes are responsible for producing neutrinos of all types. Other than electron neutrinos being produced via electron capture as in Eq. 1.21, electron antineutrinos are similarly produced by

$$e^+ + n \to p + \overline{\nu}_e \tag{1.23}$$

In addition, a host of neutral-current processes can produce neutrinos of all flavors:

$$e^- + e^+ \to \nu + \overline{\nu} \tag{1.24}$$

$$\gamma \to \nu + \overline{\nu} \tag{1.25}$$

$$\gamma + e^\pm \to e^\pm \nu + \overline{\nu} \tag{1.26}$$

$$e^\pm + N \to e^\pm + N + \nu + \overline{\nu} \tag{1.27}$$

$$N + N \to N + N + \nu + \overline{\nu}, \tag{1.28}$$

where N in Eq. 1.27 and 1.28 are nucleons taking part in *bremsstrahlung* processes. The boundary surface at which neutrinos are no longer trapped and can freely escape the core is called the *neutrinosphere*. Owing to the different ways in which neutrino species interact with the SN matter, three different neutrinospheres can be defined. Electron neutrinos, able to interact with neutrons via CC interactions, have the hardest time leaving the core and their neutrinosphere is the largest, producing a flux that is lower in energy compared to the other neutrino types. Electron antineutrinos can similarly interact with protons, which are however less abundant than neutrons, resulting in a smaller neutrinospheres. All other neutrino types (namely ν_μ, ν_τ, $\overline{\nu}_\mu$, and $\overline{\nu}_\tau$), only able to interact with the medium via flavor-independent neutral current interactions, are associated with the deepest neutrinosphere and the hottest flux.

Neutrino emission is currently understood to be vital to the achievement of the characteristic SN explosion. The consensus among astrophysicists, supported by recent simulation work, is that the rebound shockwave described above is, on its own, not enough to generate the expected explosion. As it propagates through the stellar medium, photodissociation saps the shock of its energy, eventually stalling it completely. In fact, if enough matter continues to accrete onto the proto-neutron star, a black hole can be formed without any explosion at all (in this case, we talk of a *failed supernova*). To arrive at the explosion, then, the shock needs to be revived in some way. Currently, the accepted mechanism for this revival is the energy deposition onto the shock by the enormous flux of neutrinos produced in the previous stages of the SN. By the time the shock is stalled, the shockwave front will sit well outside even the largest neutrinosphere, and is therefore able to absorb energy from all emitted

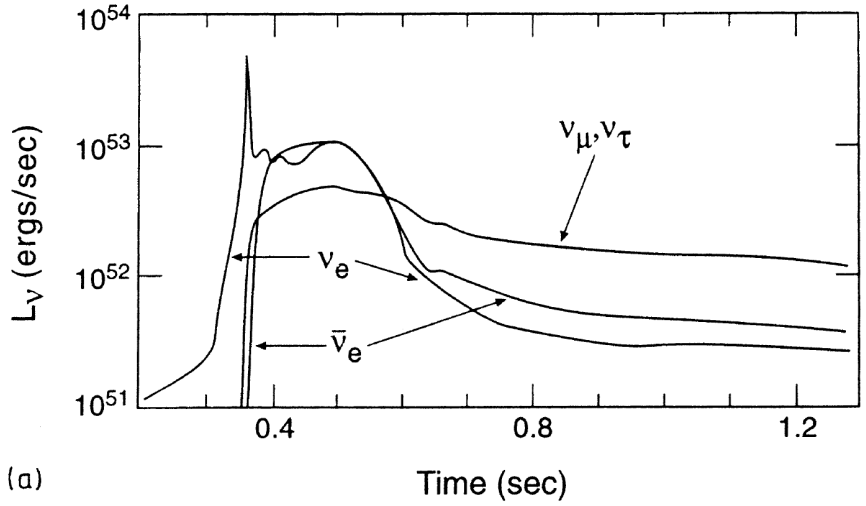

FIGURE 1.2: Time evolution of neutrino luminosity in the early phases of a SN with progenitor mass $M \approx 25M_\odot$. The prompt emission of electron neutrinos during the neutronization burst is visible as a sharp peak in luminosity. The following luminosity "hump" representing emission of neutrinos of all flavors, makes up the majority of the neutrino flux. Figure from Ref. [18]

neutrinos. Meanwhile, as the shock is stalled, more infalling matter has time to pass through the front and accrete, further heating the core and enhancing neutrino production. After about 0.5 s, enough energy is deposited to restart to shock and carry it through all the outer layers of the star, and a SN is generated.

1.2.1 The Diffuse Supernova Neutrino Background

The direct detection of supernovas with neutrinos is a rare event. Despite the enormous neutrino flux, small interaction cross-sections currently only allow for detection of SNe in our own galaxy or in its immediate vicinity, at an expected rate of just a few SNe per century. Future detectors such as Hyper-Kamiokande will broaden the detection horizon, but "only" by roughly one order of magnitude. Neutrino experiments are, for this reason, constantly on the lookout for such events. Optical detection of SNe has a much broader range, but provides different—and complementary—information about SNe. In particular, photons take longer to travel through the stellar core, and are only able to escape it further away from the center than neutrinos, so they carry more information about the outer regions of the SN. In addition, unlike neutrinos, photons interact with the interstellar medium on their way to Earth, further distorting their spectrum. Meanwhile, neutrinos escape the SN from much deeper and arrive to us completely undisturbed, providing a much cleaner probe of the SN core and clearer insight into intrinsic properties of SNe.

Rather than attempting direct SN neutrino detection, the focus of this work is the road to the detection of the Diffuse Supernova Neutrino Background (DSNB). In short, the DSNB is made up by the flux of neutrinos and antineutrinos from all causally-reachable core-collapse supernovae. The DSNB leverages the fact that, on a cosmic scale, SNe are no rare event at all: roughly one SN is expected to occur every second in the visible universe. As its name suggests, for practical observations on Earth the DSNB will appear as an isotropic and time-independent signal. Indeed, one of the appeals of the DSNB is that it allows for a *guaranteed steady source* of supernova

neutrinos. Since the first description of the DSNB, several flux models have been put forward over the years. In general, a DSNB model relies on the convolution of three ingredients: a model of the average SN neutrino energy spectrum $\mathrm{d}N_\nu/\mathrm{d}E_\nu$, an expected SN rate R_{SN}, and the expected correction due to oscillation effects in the SN matter. This can be formulated as a line-of-sight integral over the redshift parameter z:

$$\Phi_{\mathrm{DSNB}} \propto R_{\mathrm{SN}} \otimes \frac{\mathrm{d}N_\nu}{\mathrm{d}E_\nu} \otimes P_{\mathrm{oscillation}} \tag{1.29}$$

$$\frac{\mathrm{d}\Phi_{\mathrm{DSNB}}}{\mathrm{d}E_\nu} = \frac{c}{H_0}\int_0^{z_{\mathrm{max}}} R_{\mathrm{SN}}(z)\frac{\mathrm{d}N_\nu(E_\nu(1+z))}{\mathrm{d}(E_\nu(1+z))}\frac{\mathrm{d}z}{\sqrt{\Omega_m(1+z)^3+\Omega_\Lambda}}. \tag{1.30}$$

In Eq. 1.30, E_ν is the neutrino energy detected on Earth, while $E_\nu(1+z)$ is the neutrino energy at the time of emission. Effects due to oscillation are incorporated in the spectrum. With z_{max} we denote the redshift at the furthest visible gravitational collapse. Finally, the cosmological parameters H_0, Ω_m, and Ω_Λ are, respectively, the Hubble constant, the matter density and the dark energy density, describing the expansion of the universe. This formulation of the flux, as illustrated in Eq. 1.29, also lays bare some of the physics motivations for observing, or placing limits on, the DSNB flux. The neutrino spectrum at production is, as mentioned above, an important probe of supernova physics, depending not only on explosion dynamics, but also on black-hole-forming failed supernova scenarios. The supernova rate is an implicit measure of the stellar formation rate, while the inclusion of oscillation effects in matter probes properties inherent to the neutrinos themselves, such as the mixing angles and the neutrino mass hierarchy.

Cosmic supernova rate

The lifetime of a star is insignificant in the cosmic timescale. For this reason, the stellar death rate can be considered effectively equal to the stellar formation rate R_{SF}. The latter has been measured to good precision by studying the optical emission of massive stars. A DSNB observation would provide a further, *independent*, measurement of R_{SF}. As a function of z, stellar formation is observed to be $\sim$10 times at $z = 1$ than it is today, $z = 0$, and remains flat until at least $z \sim$ 4-5. Stellar formation is usually given as a rate density, in units of $M_\odot \mathrm{Mpc}^{-3}\mathrm{yr}^{-1}$. To obtain the core-collapse SN rate R_{SN}, we divide by the mean mass and rescale by the probability of undergoing gravitational collapse, assuming core collapse occurs for progenitors with masses in the range ΔM_{CC} (a typical range is $8M_\odot < M_{CC} < 50M_\odot$):

$$R_{\mathrm{SN}}(z) = R_{\mathrm{SF}}(z)\frac{\int_{\Delta M_{CC}}\psi(M)\mathrm{d}M}{\int_{0.1}^{100} M\psi(M)\mathrm{d}M}. \tag{1.31}$$

Here, $\psi(M)$ is the initial mass function (IMF), describing the mass distribution of main-sequence stars at birth. The traditional choice of IMF is the Saltpeter IMF, where $\psi \propto M^{-2.35}$.

Core-collapse SN neutrino emission

As mentioned above, due to differences in the medium opacity for a given neutrino species, the neutrino spectrum can be effectively decomposed into three different thermal spectra. A typical characterization of such thermal spectra is a Fermi-Dirac

distribution for emission at effective temperature T:

$$\frac{\mathrm{d}N}{\mathrm{d}E} = \frac{120}{7\pi^4}\frac{E_\nu^{\mathrm{tot}}E^2}{T^4}\frac{1}{e^{E/T}+1}, \tag{1.32}$$

where E_ν^{tot} is the total energy liberated for the neutrino species considered. The most recent simulations converge around the following values for the average neutrino temperatures and energies:

$$\langle T_{\nu_e}\rangle \approx \text{3-5 MeV} \qquad \langle T_{\overline{\nu}_e}\rangle \approx \text{4-6 MeV} \qquad \langle T_{\nu_{\mu,\tau}}\rangle \approx \text{4-7 MeV} \tag{1.33}$$

$$\langle E_{\nu_e}\rangle \approx \text{10-15 MeV} \qquad \langle E_{\overline{\nu}_e}\rangle \approx \text{13-19 MeV} \qquad \langle E_{\nu_{\mu,\tau}}\rangle \approx \text{13-22 MeV} \tag{1.34}$$

This simplified emission model already captures significant information about our understanding of the core collapse. Nevertheless, modern descriptions usually modify this spectrum by including shape parametrizations, for example with the inclusion of a *pinching parameter* (see Fig. 1.3), i.e. sharpening the Fermi-Dirac peak with the addition of an effective chemical potential:

$$\frac{\mathrm{d}N}{\mathrm{d}E} = \frac{(1+\alpha)^{(1+\alpha)}}{\Gamma(1+\alpha)}\frac{E_\nu^{\mathrm{tot}}E^\alpha}{\langle E\rangle^{2+\alpha}}\exp\left[-(1+\alpha)\frac{E}{\langle E\rangle}\right], \tag{1.35}$$

where $\langle E\rangle$ is the mean neutrino energy and α is the pinching parameter. In addition, we now know that neutrino emission has a substantial dependence on the mass of the progenitor star. This effect can be studied through simulation, and is usually added to the DSNB model by including a weighting function dependent on the IMF:

$$\frac{\mathrm{d}N}{\mathrm{d}E} = \sum_i \frac{\int_{\Delta M_i}\psi(M)\mathrm{d}M}{\int_{\Delta M_{CC}}\psi(M)\mathrm{d}M}\frac{\mathrm{d}N_i}{\mathrm{d}E}, \tag{1.36}$$

where the mass of progenitor i, with core-collapse neutrino emission spectrum $\mathrm{d}N_i/\mathrm{d}E_\nu$, varies in the range ΔM_i.

The picture is further modified by considering the rarer possibility of failed SNe. If the collapse terminates with creation of a black hole, the neutrino emission spectrum is significantly altered. Due to the significant shrinking of the core prior to the final collapse into the black hole, neutrino emission will be slightly more luminous and significantly more energetic, with average energies of 20-24 MeV for all flavors. Upon collapse into a black hole, emissions are brusquely interrupted, limiting the total flux from these events. Contributions to the overall shape of the DSNB spectrum are therefore significant. The DSNB can then be a handle to probe the *black hole fraction*, describing how frequently core collapses end in black holes.

Oscillation effects in matter

Due to the highly dense matter in which SN neutrinos are produced, matter effects on mixing are significant. In particular, the high density allows MSW resonances to occur well within the propagation path of the neutrinos within the SN matter, meaning that strong conversion between flavors is possible. For instance, a neutrino produced as a $\overline{\nu}_e$ will coincide almost exactly with the lightest mass eigenstate ($\overline{\nu}_1$, assuming normal hierarchy). If the matter density varies slowly enough (i.e, in the *adiabatic* case), the neutrino will stay in the same mass eigenstate as it propagates and exits the SN, and as it arrives to terrestrial detectors. The resulting flux measurement then depends on the flavor mixing in this mass eigenstate. We can then write the

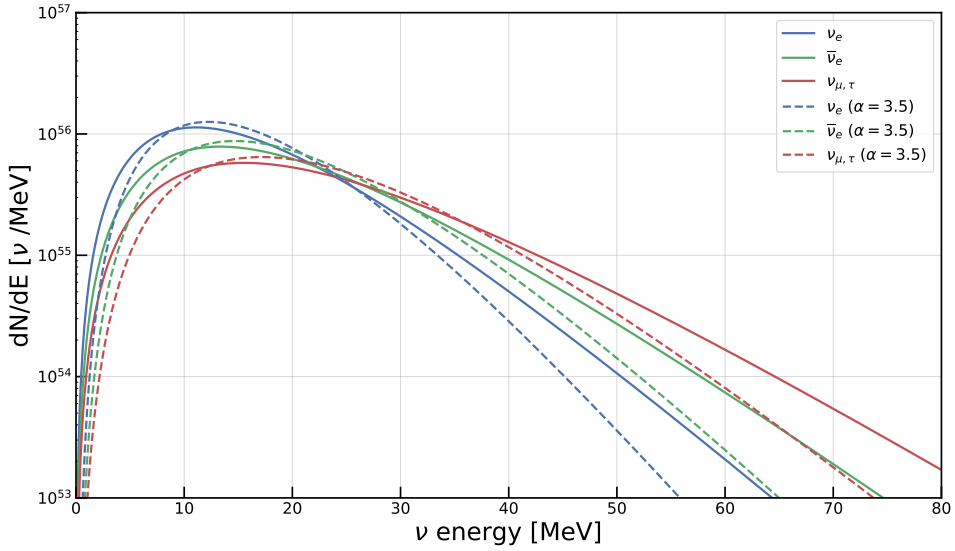

FIGURE 1.3: Comparison of standard (solid lines) and pinched (dashed lines) Fermi-Dirac emission spectra for neutrino species emitted in a core-collapse SN with total energy $E_\nu^{\text{tot}} = 2 \times 10^{59}$ MeV. Neutrino temperatures of 5 MeV, 6 MeV, and 7 MeV assumed for ν_e, $\overline{\nu}_e$, and $\nu_{\mu,\tau}$, respectively. A pinching parameter $\alpha = 3.5$ is chosen. Total energy is assumed to be equally divided among neutrino species, so each plotted spectrum represents $1/6$ of the SN energy.

measured $\overline{\nu}_e$ flux, $\phi_{\overline{\nu}_e}$, as a function of the neutrino fluxes at production, $\phi^0_{\overline{\nu}_x}$:

$$\phi_{\overline{\nu}_e} = |U_{e1}|^2 \phi^0_{\overline{\nu}_e} + (1 - |U_{e1}|^2)\phi^0_{\overline{\nu}_{\mu,\tau}} \quad \text{(NH)} \tag{1.37}$$

$$\approx \cos^2\theta_{12}\phi^0_{\overline{\nu}_e} + \sin^2\theta_{12}\phi^0_{\overline{\nu}_{\mu,\tau}} \tag{1.38}$$

In the inverted hierarchy hypothesis, $\overline{\nu}_e$ coincides with $\overline{\nu}_3$ inside the SN, and hence

$$\phi_{\overline{\nu}_e} = |U_{e3}|^2 \phi^0_{\overline{\nu}_e} + (1 - |U_{e3}|^2)\phi^0_{\overline{\nu}_{\mu,\tau}} \quad \text{(IH)} \tag{1.39}$$

$$= \sin^2\theta_{13}\phi^0_{\overline{\nu}_e} + \cos^2\theta_{13}\phi^0_{\overline{\nu}_{\mu,\tau}} \tag{1.40}$$

$$\approx \phi^0_{\overline{\nu}_{\mu,\tau}}, \tag{1.41}$$

where the approximations come from the fact that θ_{13} is small enough ($\theta_{13} = 8.6°$) that, to first order, $\sin^2\theta_{13} \approx 0$ and $\cos^2\theta_{13} \approx 1$. Since the muon and tau neutrino SN flux is more energetic, the result of flavor conversions inside the SN envelope is a significant hardening of the spectrum. The inverted hierarchy case, presenting near-total flavor conversion and more energetic spectrum, could therefore be differentiated from the normal hierarchy case. Additional mixing effects appear if the density of the star does not change adiabatically along the travel path of the neutrino. In particular, this would allow for a non-zero probability for the neutrino to switch from one mass eigenstate to another near a resonance. Additional effects can further be induced by the effective matter potential due to forward ν-ν scattering, which would modify the Hamiltonian in Eq. 1.11 and hasn't been discussed in this text.

These are some of the elements currently used to develop theoretical DSNB descriptions. This introduction hopefully highlighted the physical richness of the DSNB signal, and the value of its eventual measurement and characterization to the neutrino physics, astrophysics, and cosmology communities. DSNB flux predictions for a variety of models, operating under a range of different assumptions and parameters,

are plotted of In Fig. 1.4.

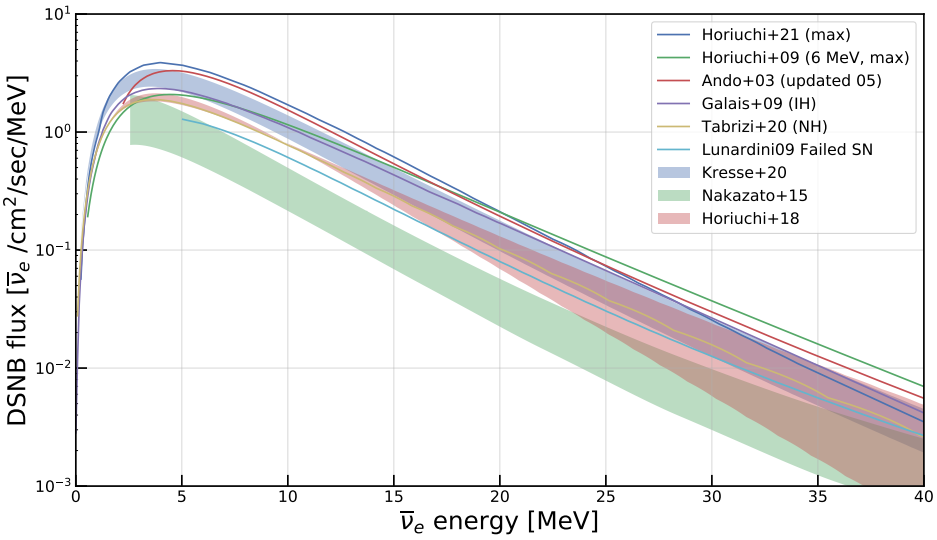

FIGURE 1.4: DSNB $\bar{\nu}_e$ flux predictions from various theoretical models [19, 20, 21, 22, 23, 24, 25, 26, 27]. Refer to the corresponding publications for detailed descriptions of each model. The large variety of models is reflected in significant variation in flux shape and absolute scale. Some numerical estimates are not available for the entire energy range.

1.2.2 Detection landscape

SN1987A

So far, the only detected supernova neutrinos come from SN1987A. This supernova was observed at 07:35 UT, on February 24th, 1987, simultaneously by three neutrino experiments: Kamiokande, in Japan, detecting 12 antineutrinos [28]; IMB, in the U.S., with 8 antineutrinos [29]; and Baksan, in Russia, with 5 antineutrinos [30]. This detection occurred several hours before confirmation came from optical observatories, and effectively marked the beginning of neutrino astronomy. Neutrino detection was at all possible due to the SN's relative proximity: while not in our own galaxy, it originated in the Large Magellanic Cloud, a satellite galaxy of the Milky Way roughly 50 kpc away from the Solar System. As such, it is the best-studied SN, even outside of neutrino physics, and is the only one to be visible to the naked eye in recent history (the previous recording of such a clearly-visible SN dates back to the Kepler SN in 1604 [31]).

A worldwide total of 25 observed neutrinos may seem small, but from these precious data points, astrophysicists have been able to infer a remarkable amount. For one, the neutrino signal was showed to be in impressive agreement with accepted SN neutrino emission models, which had, until then, no observational backing. In particular, the study in Ref. [32] showed that the neutrino emission was consistent with a supernova in which 99% of the binding energy is radiated away as neutrinos through the delayed—and not the prompt—shock mechanism. The neutrino observation was enough to decisively favor the delayed explosion scenario, found to be roughly 100 times more likely. Two timescales were identified, associated with the neutrino emission in two stages: $\tau_{\text{prompt}} \approx 4$ s and $\tau_{\text{delayed}} \approx 0.7$ s. The neutrino energy itself, ≈ 15 MeV was well within expectations for $\overline{\nu}_e$ emissions (e.g. Eq. 1.34). In addition, SN1987A afforded us an independent bound on the ν_e mass scale. A

model-independent mass limit can be extracted [33] with the assumption that the variance in the neutrino time-of-flight ΔT, due to the variance in emission energy ΔE, is smaller than the intrinsic and observed burst duration, ΔT_0 and ΔT_{obs}, respectively:

$$\Delta T < \Delta T_0 \leq \Delta T_{\text{obs}}. \tag{1.42}$$

For highly relativistic particles such as neutrinos, emitted from a source at distance D with average velocity v and average energy E,

$$v = \frac{p}{E} \simeq 1 - \frac{m_{\nu_e}^2}{2E^2} \tag{1.43}$$

$$\Delta T = \Delta \left(\frac{D}{v} \right) \simeq \frac{m_{\nu_e}^2}{E^2} D \frac{\Delta E}{E} \tag{1.44}$$

yielding $m_{\nu_e} \lesssim 30$ eV. Model-dependent limits can bring this figure down, notably by making assumptions about the intrinsic duration of the burst. The model-dependent study in Ref. [32] placed a limit at $m_{\nu_e} < 5.7$ eV to 95 % CL for a delayed-shock SN model.

Current and future supernova burst detectors

Since 1987, the vibrant field of neutrino astronomy has produced a new generation of detectors, or *neutrino telescopes*, able to characterize future nearby SN bursts with greater precision, using a variety of detection technologies.

Two main classes of such detectors are currently in use: *water Cherenkov detectors* and *scintillator detectors*. The former look for neutrinos by detecting *Cherenkov radiation* (explored in more detail in the following chapter) associated with their interaction inside a volume of water. Super-Kamiokande (the experiment of relevance to this thesis) belongs, like its predecessor, Kamiokande, to this category, sporting a 50-kton volume of water. Other detectors, running at the time of writing, in this category are IceCube (using a cubic kilometer of Antarctic ice as detection volume), and Baikal (operating under the homonymous Russian lake). In the following years, these will be joined by Hyper-Kamiokande (with 1000 ktons of water) and KM3NeT (detecting interactions under the Mediterranean at three different European sites). Scintillator detectors, in contrast, utilize scintillating material as their detection medium, looking for the scintillation light associated with charged particles traversing their volume. Among these, Baksan, KamLAND (in the same underground facility as Super-Kamiokande), Borexino and LVD (both located under the Gran Sasso mountain in Italy), NOνA (at Fermilab, in the U.S.) and Daya Bay (in southern China) are the main scintillator detectors, operational at the time of writing, with supernova burst detection capabilities. In the years to come, the JUNO and SNO+ experiments will join their ranks. A further range of technologies is employed by additional experiments: HALO (detecting neutrino interactions in a volume of lead), DUNE (still under construction, using a Liquid Argon Time Projection Chamber), as well as Xenon and ArDM (detecting both ionization and scintillation signals).

The combined sensitivity from the currently-operational experiments affords the detection of thousands supernova neutrinos in the case of a burst in or near the Milky Way, an impressive number when compared to the 25 neutrinos observed in 1987, allowing for equally impressive gains in our ability to characterize the SN in detail. A number of these experiments (namely, Super-Kamiokande, KamLAND, IceCube, Borexino, LVD, Daya Bay, and HALO) are part of the SuperNova Early Warning System (SNEWS) [34], set up to promptly and swiftly alert astronomers

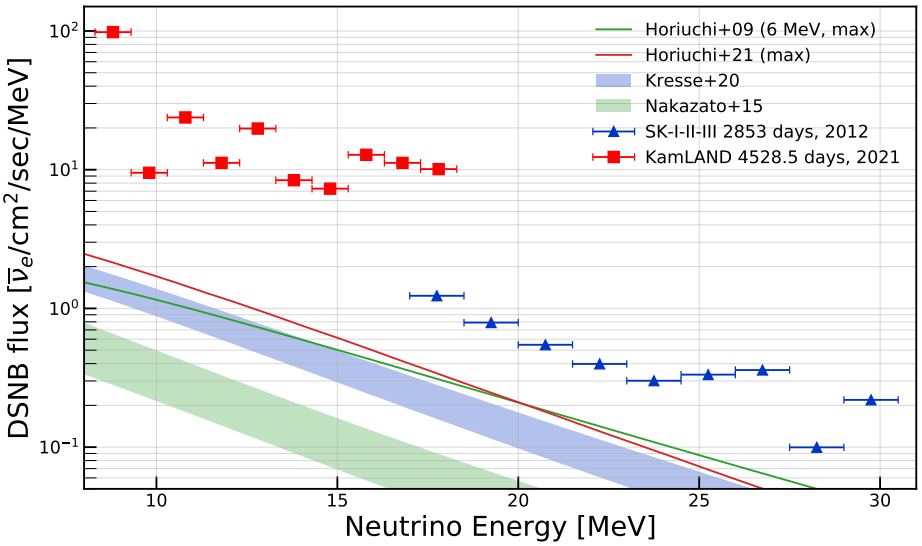

FIGURE 1.5: The best 90% CL upper DSNB flux limits before the SK-IV data-taking phase. The KamLAND upper limits (red squares) are the best ones at the low end of the DSNB spectrum, while SK (here including data-taking runs I, II, and III) places better limits at the high end (blue triangles). A selection of DSNB models, representing the range of modern flux predictions, is also shown.

worldwide in the case a positive detection. Supernova neutrinos, reaching Earth several hours before any electromagnetic radiation can be detected, can alert optical observatories of the upcoming burst. Further, a combined angular sensitivity of roughly 5° would allow astronomers to point their telescopes to the relevant section of the sky and optimize their observation. Next-generation experiments such as Hyper-Kamiokande, ever pushing the boundaries, will allow detection in the order of 10^5 neutrinos from a single burst, and expand the horizon of neutrino detection out to neighboring galaxies such as Andromeda.

Past and future limits on the DSNB

As of yet, the DSNB signal has proved too faint to be observed by any neutrino telescope. A few experiments are nonetheless sensitive enough to have placed upper limits on the DSNB flux, excluding the most optimistic estimates. At the low end ($E_\nu < 17.3$ MeV) of the DSNB energy spectrum, and before considering the SK-IV period, KamLAND placed the tightest limits in 2022 [35]. Using their 1-kton scintillator detector, with a total exposure of 6.72 kton-years, they searched for DSNB events for neutrino energies between 8.3 and 30.8 MeV. The analysis observed 18 candidate events, but no significant excess, placing model-independent upper flux limits between 7.3 and 98.1 $\overline{\nu}_e\text{cm}^{-2}\text{s}^{-1}\text{MeV}^{-1}$ at 90% confidence level (CL). For energies $E_\nu > 17.3$ MeV, the best limits have been placed by SK. Before the analysis of data from the SK-IV data-taking phase (detailed in this work), the best limits were placed using data from the SK-I, II, and III phases [36]. Such an analysis, leveraging 177 kton-years of exposure, placed model-independent limits between 1.1 and 0.1 $\overline{\nu}_e\text{cm}^{-2}\text{s}^{-1}\text{MeV}^{-1}$ at 90% CL, for neutrino energies in 17-30.5 MeV. The KamLAND and SK-I-II-III limits are displayed in Figure 1.5 This latter result is tantalizingly

close to several modern DSNB models, demonstrating in practice the feasibility of detection.

One of the main objectives of the work presented in this text is indeed to update the SK upper flux limits on the DSNB, using the latest SK data. The total exposure available to SK-IV, roughly 180 kton-years, is comparable to the previous SK search, and can therefore provide an equally sensitive measurement, before considering the improvements in the SK data inherent to this phase (see Section 2.2). The near future will also bring exciting developments. For one, SK itself has already undergone an important transformation in late 2020, its water having been doped with Gadolinium (Gd), dramatically enhancing DSNB sensitivity (see Section 2.3). In addition, JUNO and HK—both having the DSNB as a major physics target—will start their data-taking within the decade. Improvements to our sensitivity to the DSNB, already near the upper edge of many predictions, are therefore accelerating at this very moment. Simply put, the DSNB is one of the most likely discoveries in neutrino physics to expect in the years to come—and a much awaited one! Being first described nearly 40 years ago, the DSNB is finally in striking distance.

2
The Super-Kamiokande detector

The work in this thesis relies on the Super-Kamiokande (SK) experiment, located inside the historic Kamioka mine in central Japan. Shielded from the radiation above by mount Ikeno, SK holds inside its 40-meter-high steel tank 50,000 tons of water, detecting the neutrino interactions therein by looking for Cherenkov radiation. It is, at the time of writing, the world's largest water Cherenkov detector.

Super-K's predecessor, Kamiokande, was envisioned in 1982 as a means to detect proton decay [37]. Indeed, the name *Kamiokande*, originally standing for *Kamioka Nucleon Decay Experiment* betrays the project's initial purpose. Proton decay searches became popular starting in the mid-70's, when Georgi, Quinn and Glashow [38, 39] estimated the proton lifetime as $\tau_p = 6 \times 10^{31}$ years, based on the hypothesized unification of the weak, electromagnetic, and strong interaction under the single transformation gauge group SU(5). This prediction sparked excitement in the experimental world: with 1000 tons of water, corresponding roughly to 6×10^{32} nucleons, one could try their hand at the first observation of proton decay, and the first hints for the unification of the SM forces. With its ~ 2000 tons of water, Kamiokande was one of the forerunners in the race for proton decay. Alas, the aforementioned proton lifetime was eventually ruled out experimentally, and no proton decay has been detected to date. Though falling short of the awaited discovery, Kamiokande proved to be an excellent tool for neutrino detection. In its original proposal, atmospheric and supernova burst neutrinos were mentioned as secondary physics goals. In 1987, the historic observation of neutrinos from supernova SN1987A [28] affirmed the detector's low-energy capabilities, and neutrino astronomy started taking up more space in the collective consciousness of the experimental community.

Plans for a larger detector would already begin soon after the start of Kamiokande itself. Initially proposed as an extension of the proton decay project, Super-Kamiokande began being considered more and more as a neutrino detector, as the quality of Kamiokande's neutrino measurements—both supernova and solar neutrinos [40]—became apparent. By the time SK started taking data in 1996, the priorities of the collaboration had shifted to include neutrino astronomy, supernova neutrinos, solar neutrinos, and atmospheric neutrinos. Since then, SK contributed to a great deal of measurements and discoveries in neutrino physics, including the historic 1998 discovery of neutrino oscillations [13]. Today, the SK collaboration comprises ~ 200 physicists and engineers from around the world [41]. In this chapter, an overview of the detector is given, describing its physical design and its data systems. Finally, the detector's current state and its future are discussed: namely, the recent SK-Gadolinium upgrade, of great importance to DSNB searches.

2.1 Detector design

Cherenkov radiation

Super-Kamiokande is a *water Cherenkov detector*, meaning that it detects the interaction between particles inside its volume of water by looking for the associated Cherenkov radiation. Water is a convenient medium for Cherenkov detectors due to its easy availability and good transparency. Cherenkov light is emitted when a charged particle propagates in a dielectric medium with a speed higher than the speed of light in that medium. This effect is analogous to the *sound boom* observed when an aircraft moves at speeds higher that the speed of sound, emitting sound waves that accumulate in a cone pointing in the direction of motion. Under normal conditions, electrically-charged particles moving at constant speed will polarize a surrounding dielectric without emission of electromagnetic radiation. However, in a medium in which the speed of light, c is lower than the speed of light in vacuum, c_0, matter can move with speed $v > c$. In this case, the resulting polarization wave fronts will cross and constructively interfere with each other, with consequent light emission. The number of photons N emitted per unit length x and wavelength λ is given by [42]

$$\frac{\mathrm{d}N}{\mathrm{d}x\mathrm{d}\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2}\left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right) \tag{2.1}$$

where α is the fine-structure constant, Z is the atomic number of the material, $n = c/c_0$ is the refractive index of the material, and $\beta = v/c_0$. Lower-wavelength photons are thus produced in higher numbers, giving Cherenkov light its characteristic blue glow. Note that, since $n(\lambda) \to 1$ as $\lambda \to 0$, Eq. 2.1 is not divergent. Photons propagate along the surface of the cone defined by the angle θ_C, measured from the direction of motion:

$$\cos\theta_C = \frac{1}{\beta n} \tag{2.2}$$

Neutrinos, themselves being neutral, can then be detected when their interaction is associated with a charged particle propagating at superluminal speed in the medium.

Physical characteristics

The location of Super-Kamiokande is determined by the need to shield it from cosmic ray muons, which, due to their capacity to penetrate to high depths, would represent an unacceptably-high background. SK's positioning inside the Kamioka mine, under Mt. Ikeno, affords it about 1 km of rock cover, corresponding to 2700 m-water-equivalent. At this depth, a muon flux of 6×10^{-8} cm^{-2} s^{-1} sr^{-1} is observed, roughly four orders of magnitude smaller than the surface flux.

Cherenkov radiation is detected within 50 ktons of water contained in a cylindrical stainless-steel tank, 39 m in diameter and 42 m in height. This volume is further divided into two concentric regions: an outer detector (OD) and an inner detector (ID). Separating the two is a 55-cm-thick cylindrical structure, optically isolating the ID and OD from each other. Its inner surface, defining the ID volume, measures 33.8 m in diameter and 36.8 m in height. A schematic view of the experiment is given in Figure 2.1. The ID represents the main detection volume, while the OD is used to detect and veto incoming cosmic muons. Detection of light occurs with an array of photomultipliers (PMTs) lining the surface of the ID (i.e. facing inwards) and the

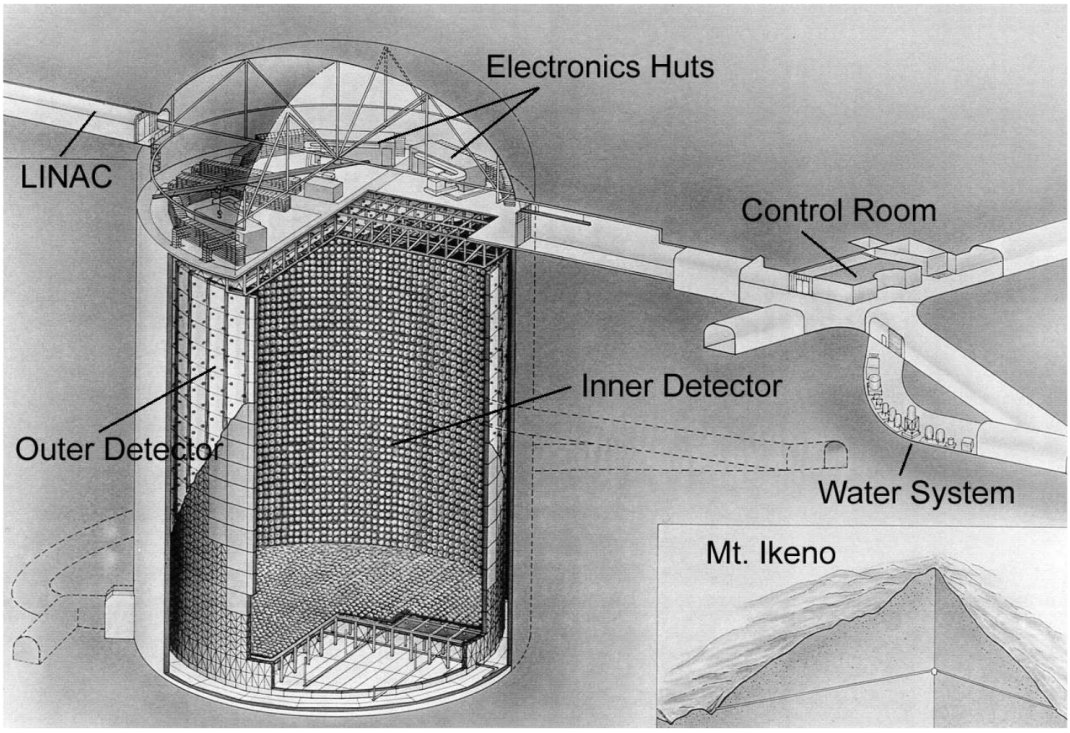

FIGURE 2.1: The SK detector. Illustration from Ref. [42]. In addition to the PMT-lined outer and inner detectors, the mine access tunnels are shown. On top of the tank, the electronic huts contain the hardware necessary for processing and storing data from the PMTs (more information in Section 2.2). The water system ensures the stability and quality of the SK water. The LINAC (LINear ACcelerator) is used to accelerate electrons into the tank volume, and serves as one of methods for the calibration of the PMTs.

inner surface of the OD (facing outwards). The inward-facing array consists of 11,129 hemispherical PMTs, each 50 cm in diameter, covering 40% of the ID surface. The outward-facing array consists of 1885 hemispherical PMTs, each 20 cm in diameter.

The SK PMTs consist of an evacuated glass bulb with a photosensitive surface, tightly connected through a seal with a series of dynodes and an anode, leading to a single cable for each PMT. Figure 2.3 shows a sketch of an inner detector PMT. When photons strike the photosensitive area of the PMT, electrons are emitted through the photoelectric effect, which are then directed to the dynode chain by an electric field. As each of the electrons hit the first dynode, several electrons are emitted through secondary emission, effectively multiplying the electron flux. This effect repeats, resulting in a cascade of electrons which, by the time it reaches the anode, is easily detected by the electronics. This mechanism allows for highly sensitive detection of low-energy photons. In practice, there are several sources of inefficiency inherent to the PMTs. For one, the photosensitive material (a bialkali photocathode composed of Antimony, Potassium and Carbon) has a peak quantum efficiency of 21% for wavelengths between 360-400 nm, as shown in Figure 2.3. In addition, the collection efficiency of the first dynode is $>$70%.

Water and air purification

Even small impurities in SK's water and air can be a threat to the quality of its data. Specifically, contamination by radioactive species (most importantly, Rn) will add bothersome backgrounds, while a deterioration in the transparency of the water

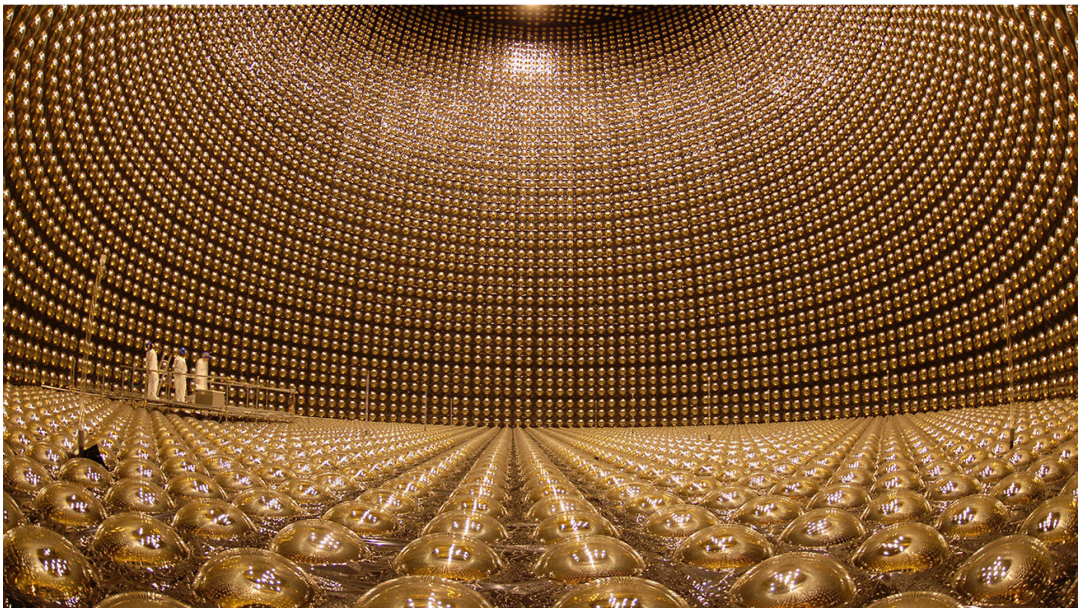

FIGURE 2.2: The interior of SK, during a maintenance period, emptied of its water. The three scientists standing on the bottom surface suggest the scale of the experiment. Such a view of the empty detector, revealing its golden-yellow PMT-lined surface in its full glory is a rare sight!

will decrease the light yield of the detector. Therefore, an extensive water and air purification system is employed to keep contamination levels low and water transparency high.

The water filtration system, pictured in Figure 2.4, consists in a closed loop, through which SK's 50 ktons of water are continuously processed at a rate of roughly 30 tons per hour. The SK water starts its journey from the top of the tank, herded into the cycle by a pump. A first filtration is achieved through a 1μm mesh filter, removing dust and other particles that could be Rn sources and contribute to low transparency. The temperature of the water is regulated down to 12.9C by two heat exchangers at different points in the cycle. Lowering the temperature reduces PMT noise and bacterial growth (itself reducing transparency). A cartridge polisher is then used to remove heavy ions. Any leftover bacteria are killed off by a UV sterilizer. A reverse osmosis system further removes particulates. The vacuum degasifier module removes gases that are dissolved in the water. Radon is reduced by 96% at this stage, which is helped by the introduction of radon-less air beforehand. Other than radon gas, the vacuum degasifier notably removes Oxygen gas, to discourage bacterial growth. An ultrafilter, able to remove particles down to 10 nm, is then used for finer filtration. Finally, a membrane degasifier provides additional removal of dissolved gases, and the now-filtered water returns to the tank from its bottom end. The efficacy of this system is impressive, rendering the water *ultra-pure*, with a resistivity near the theoretical maximum of 18.2 MΩ-cm.

Air purification is aimed at reducing Rn levels. A dedicated system purifies the air in the gap between the water surface and the top of the tank, consisting of several filters, compressors, and dryers, controlling the temperature and pressure of the filtered air. A slight over-pressure is induced to keep the Rn-rich air from outside the tank from leaching back in. To reduce contamination in the dome area above the tank, fresh air is continuously pumped in from outside the mine.

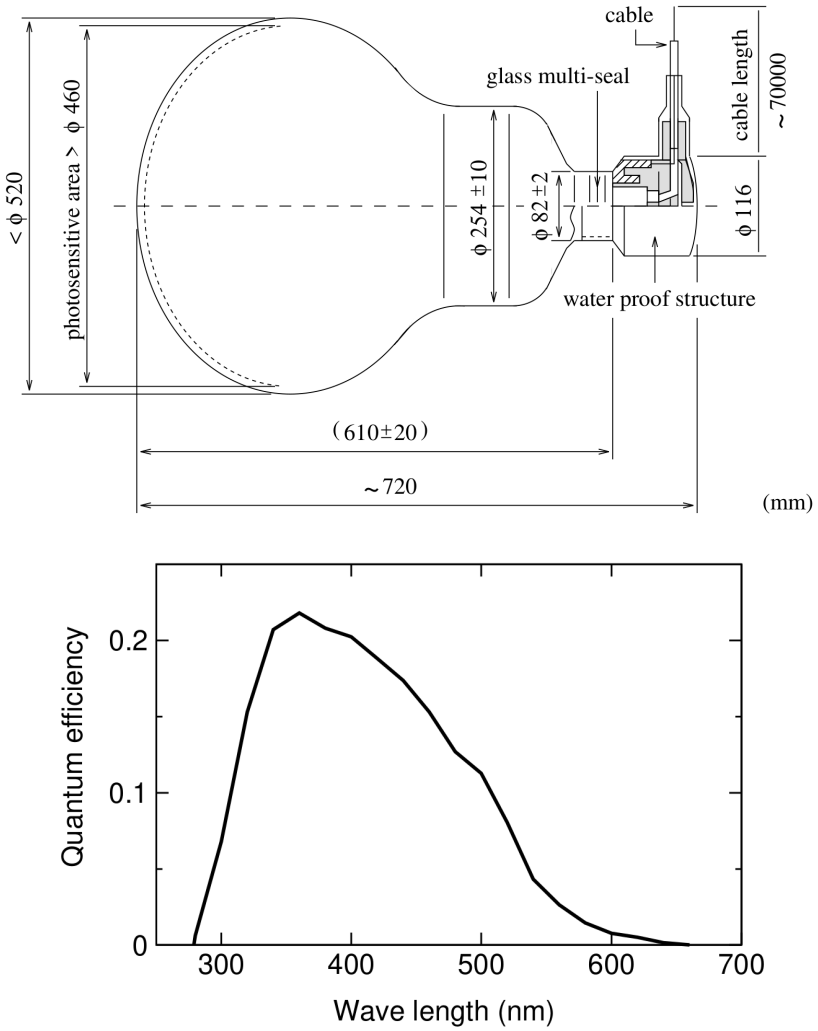

FIGURE 2.3: Top: schematic view of a 50-cm PMT located in the inner detector. Bottom: quantum efficiency of the photocathodes in the SK PMTs, as a function of wavelength. From Ref. [42].

Detector performance

The detector design described above effectively defines the kind of events we are able to detect with SK. Firstly, it defines the accessible energy range. Over the years, the SK collaboration has continuously strived to push down the lower energy threshold of the detector. The latter is limited by our ability to suppress low-energy backgrounds, which is in large part achieved through the aforementioned purification methods. Currently, the SK energy threshold stands at $\sim$ 3.5 MeV. At the low end of the accessible energy we find neutrinos emitted by nearby nuclear reactors, with energies ranging from 3.5 to 10 MeV, as well as neutrinos produced in fusion reactions in the Sun, from 3.5 to 15 MeV. As discussed in the previous chapter, neutrinos from SNe—either from a nearby burst or from the DSNB—will arrive with an energy < 50 MeV. Searches for proton decay look for neutrinos with energies between 100 MeV and 1 GeV. Finally, atmospheric neutrinos, generated in the particle shower from cosmic rays interacting with the atmosphere, occupy the entire accessible energy range to varying degrees, up to energies of $\mathcal{O}(100)$ GeV.

The energy resolution ranges from 14.4% (at 10 MeV) for low-energy events,

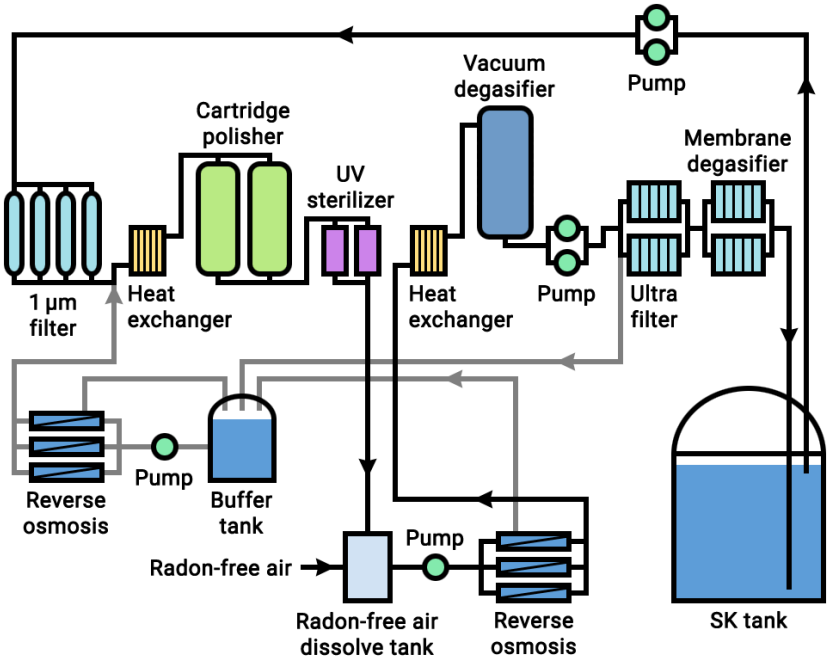

FIGURE 2.4: Diagram showing the stages of the water purification cycle of SK. Illustration from Ref. [42].

to $(1.7 + \sqrt{E(\mathrm{GeV})}\%)$ for high-energy atmospheric neutrinos. The time resolution is determined by the intrinsic time resolution of the PMTs, or 2.8 ns for a single photoelectron. Angular information for single events is not always available. An example of an interaction with accessible directionality is the forward scattering $\nu_e + e \to \nu_e + e$, which can be used to identify solar neutrinos with $\sim 20^\circ$ resolution at 10 MeV, limited by a low PMT activation at these energies, as well as by the multiple scatterings undergone by the electron. For GeV-scale interactions with muons as the recoil particle, $\nu_\mu + X \to \mu + X'$ the angular resolution can be as good as 2° for neutrinos arriving at the bottom of the detector.

2.2 Data acquisition and electronics

When a PMT in SK detects a photoelectron, two quantities are recorded: the detection time and the charge deposited in the PMT. Combining these detections allows us to produce images of the Cherenkov rings associated with neutrino interactions. Some of these "photographs," as it were, can be quite striking—beautiful, even–particularly for higher-energy events (see Fig. 2.5). For the purposes of high-performance Cherenkov imaging, Super-K's detection energy range, spanning six orders of magnitude, requires both a wide dynamic range and high sensitivity to single-photoelectron charge reading. In addition, good time resolution is essential to allow reasonable reconstruction of interaction vertices. Finally, the processing speed of the electronics processing the incoming signal needs to be high enough to handle high signal throughput, particularly in the case of fast events such as nearby SN bursts.

To better fulfill these needs, the electronics system at SK underwent a major upgrade in 2008 [43, 44]. The details that follow refer to these new electronics. This new system, dubbed QBEE (QTC-based Electronics with Ethernet) can be broken down into two main processing steps. The first step, using charge-to-time converters

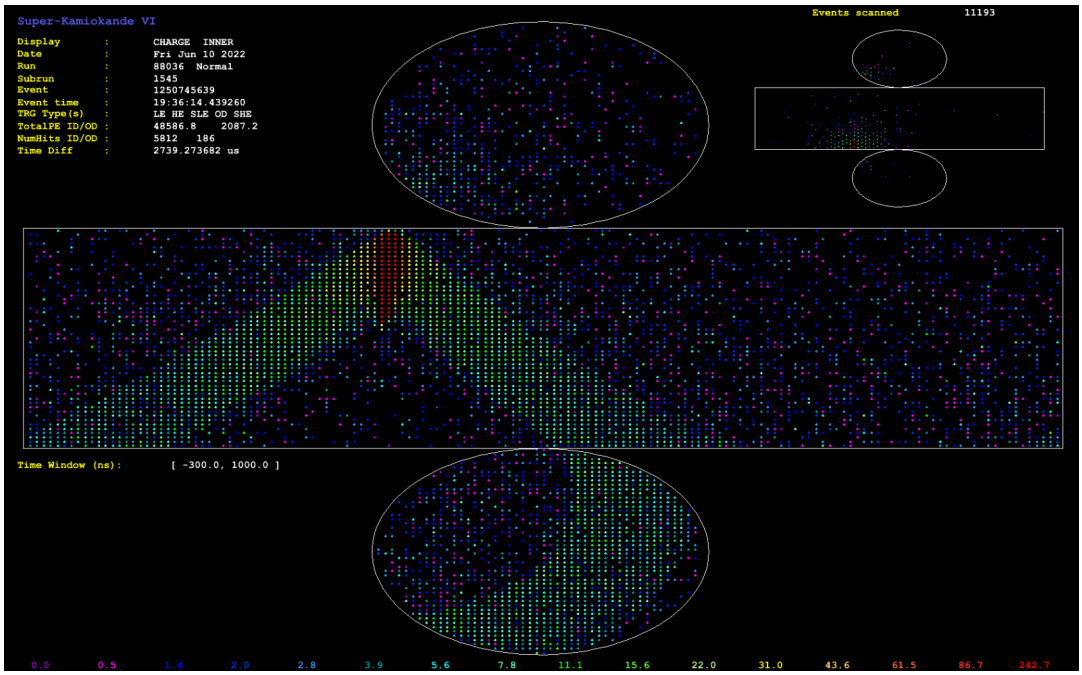

FIGURE 2.5: Example of a muon event at SK from June 2022. The cylindrical detector surface is unraveled, separating the two endcaps from the barrel. Each lit-up pixel corresponds to a PMT hit, with warmer colors indicating a higher total charge deposition. The Outer Detector (upper right) is lit up, indicating a muon event. The Inner Detector displays the typical Cherenkov cone structure of an SK event, here striking the detector from the top, and traversing the volume near the side. Cosmic muons such as this one, are highly energetic and suffer little deflection thanks to their large mass, leading to copious PMT activity and a well-defined cone.

(QTC), converts the charge signal from the PMT to a square waveform, encoding the hit time and the total charge deposited as the start time and width of the waveform. In the second step, a multi-hit time-to-digital converter (TDC) digitizes the times of signal edges, after which the hit time and charge information is easily calculated by a field programmable gate array (FPGA). *All* digitized hits are then recorded, in $\sim 17\mu$s recording windows. One of the main advantages of this system, compared to the previous implementation is its ability to record every PMT hit up to very high event rates. This means that it does not suffer from any loss of signal up to very high frequencies, which is of great importance for SN burst detection, as well as $\mu \to e$ decay events. In addition, it allows the recording of high-rate, low-energy hits, which would earlier have been discarded, as they are usually associated with dark noise (on the order of a few kHz/PMT). Keeping these hits enables detection of very low energy ($< \sim 5$ MeV) events, notably the 2.2 MeV signal associated with a neutron, crucial, among other things, to DSNB searches. In other words, the dead time of the electronics is now low enough to be negligible for the majority SK events. The timing resolution of the electronics is 0.3 ns for a single photoelectron, well below the 2.8 ns resolution of the ID PMTs. A charge resolution of 2 pC for a single photoelectron (a 5% resolution) is achieved. The new system enables an impressive dynamic range in charge: the QBEE channels can handle signals up to 2500 pC before saturating. This is a five-fold increase over the previous system, improving the energy resolution for events $\gtrsim 1$ GeV to less than 4%. Finally, the maximum throughput of the electronics is of 11.8 MB/s for each QBEE board, equivalent to ~ 80 kHz per channel. This would allow for lossless recording for event rates of up to 130 kHz (theoretically enabling

the lossless detection of up to 6 million SN burst-like events over 10 seconds), an impressive 100-fold improvement over the previous setup. These capabilities allow for the lossless detection of burst SN signals up to 7500 larger than a SN at center of the galaxy. Larger SN signals would suffer from data losses due to the inherent dead time of the electronics. This would only be the case for an especially-close SN burst, for example in the case of the collapse of Betelgeuse, at 0.2 kpc from Earth.

The PMT hits, having been recorded with virtually perfect efficiency, become part of an SK event if they pass the threshold for an event trigger. Several triggers are defined according to the physics needs of the analyses. The standard triggers are defined by a threshold in the number of ID PMT hits inside a 200 ns coincidence window. For example, the High Energy (HE) trigger is activated when a threshold of 52 hits is reached. A value of 200 ns is chosen because it is approximately the maximal time needed by a photon to traverse the ID volume, and therefore represents the maximal time extension of the hits from a single vertex. A single OD trigger is also similarly defined, with a threshold of 19 hits in 200 ns. If an ID trigger is activated close in time to the OD trigger, the hits from the two triggers are combined in a single event. In the absence of a corresponding ID trigger, ID hits are still stored in the event of an OD trigger. Additional types of triggers are defined, for example to look for neutron signals, or for calibration purposes. Once one of the trigger conditions is met, the hits inside the corresponding event time window are stored. A typical SK event from a standard lasts 40 μs, starting 5 μs before the trigger time and ending 35 μs afterwards. Some triggers produce longer-lasting events, typically to enable detection of neutron signals, which can occur several hundreds of microseconds after the trigger time. An overview of the main trigger schemes used at SK is given in Table 2.1.

Trigger	Abbr.	Hits	Event duration (μs)	Notes
Super Low Energy	SLE	35	$-5 \rightarrow 35$	-
Low Energy	LE	49	$-5 \rightarrow 35$	-
High Energy	HE	52	$-5 \rightarrow 35$	-
Super High Energy	SHE	60	$-5 \rightarrow 35$	-
After Trigger	AFT	-	$35 \rightarrow 535$	SK-IV onward
Outer Detector	OD	22		Only available OD trigger

TABLE 2.1: Overview of the main SK trigger schemes, showing the abbreviations used in this text, along with hit thresholds and the duration of the subsequent event. Note that the exact hit thresholds are subject to change over the course of SK-livetime, often within the same data-taking period. The values given here are the ones used in November 2022. The AFT trigger is automatically initiated after an SHE event.

Supernova monitors

Before the collaboration has any chance to analyze the data, online monitoring systems work around the clock to automatically perform (relatively) simpler checks and operations on the incoming data stream. These include not only everyday tasks such as data quality and data transfer checks, but also a real-time SN monitor, designed to independently identify nearby SN burst events as promptly as possible. Our current understanding of SN burst emissions, informed both by Kamiokande's experience with SN1987A and by extensive theoretical predictions, dictate the strategy for the search algorithm. In particular, the algorithm searches for emissions lasting over a total timescale of a few 10s of seconds, with the majority of the events occurring

during the first one or two seconds. Another important feature of an SN burst observation is its spatial distribution. SN burst neutrino events are expected to be very uniformly distributed in the tank volume, allowing for a clear discrimination of clustered backgrounds, such as cosmic-ray-induced radioactive decay. Additionally, with the introduction of Gd in its water, SK can now more reliably detect the neutron signal in an IBD interaction from a SN neutrino, providing an additional handle on such events. When enough such candidates, passing the LE trigger threshold and various noise reduction filters, have been identified, an SK-internal alarm is prompted. If the events withstands further analysis, the alert is made public, interfacing with the aforementioned international SNEWS network for SN alerts. The alert can be triggered by bursts up to ~ 100 kpc away. A preliminary alert can be generated in less than a minute for burst at > 10 kpc. At closer distances, where the high event rate requires longer processing times an answer can be reached in minutes ($\sim$5 minutes for bursts at 3 kpc, for example). The sheer amount of expected neutrino events also allows us to determine the location of the SN in the sky with good precision. While most of the SN events will be IBD interactions, roughly 100-150 forward scattering events $\nu + e^- \rightarrow \nu + e^-$, are expected, yielding an angular uncertainty of $\sim 3^\circ$.

In addition to the burst itself, SK has the capability to detect *pre-Supernova* neutrinos, and it is set up with a real-time monitor for such events [45]. As the name suggests, these are neutrinos emitted before the explosion itself. Pre-SN emissions happen during the final stages of the star's life, as it progresses from burning Hydrogen to increasingly heavier elements, with a significant increase in neutrino emissions from thermal and nuclear processes, dominated by pair annihilation, $e^- + e^+ \rightarrow \nu + \overline{\nu}$. As the Si in the core is burned into Fe—over the span of just a few days—and the final iron core is produced, the neutrino flux becomes large enough to be observed by SK, detecting the $\overline{\nu}_e$ via the IBD interaction. To date, no pre-SN neutrino has been observed. A positive pre-SN detection would not only herald an imminent SN event, but also be a measurement of interest in its own right.

2.3 The SK-Gadolinium upgrade

Starting in 2020, SK began a new chapter in its career, changing for the first time the nature of its detection medium. The water, until then ultra-pure, was, at this time doped with Gadolinium (Gd), greatly enhancing the signature of a neutron capture and paving the way to a new era of SK analyses. The idea for such an upgrade was already envisioned in 2002. The delightfully-named GADZOOKS! (*Gadolinium Antineutrino Detector Zealously Outperforming Old Kamiokande, Super!*) project set about determining the feasibility of the undertaking [46]. The use of Gd for neutron identification was already common for <1 kton detectors, but it was unheard of for a detector at the scale of SK. The main physical motivation for the addition of Gd to the water was, above all, improved SN detection—both burst and DSNB signals—through efficient neutron identification. Other targets included the detection of $\overline{\nu}_e$ produced in the Sun, and the separation of atmospheric neutrinos and antineutrinos (producing preferentially a proton and a neutron, respectively). Two properties of Gd make it particularly attractive. First, a neutron being captured by a Gd atom produces a γ cascade with total energy of $\sim$8 MeV, a dramatically higher than the 2.2 MeV of the single photon emitted during neutron capture on the Hydrogen in pure water. This increase allows much more reliable detection of the neutron signal component of the IBD interaction. Second, the neutron-capture cross-section of Gd is a whopping 49000 barns, much higher than the 0.3 barns for neutron capture on

free protons. This means that a very small amount of Gd—limiting costs and adverse detector effects—is needed to ensure that a majority of neutron captures occur on Gd. Doping the water with 0.1% Gd concentration by mass would ensure 90% of the neutron captures occurring on Gd, though a concentration of 0.01% is already enough to have half of the neutron captures on Gd (see Fig. 2.6). The specific form of Gd chosen for the doping was Gd sulfate, $Gd_2(SO_4)_3$. Gd sulfate is water-soluble (unlike pure Gd), provides good Cherenkov transparency (unlike Gd nitrate, $Gd_2(NO_3)_3$), and is less reactive than other options (such as Gd chloride, $GdCl_3$). In addition, the Gd sulfate is octahydrated for better solubility, $Gd_2(SO_4) \cdot 8H_2O$.

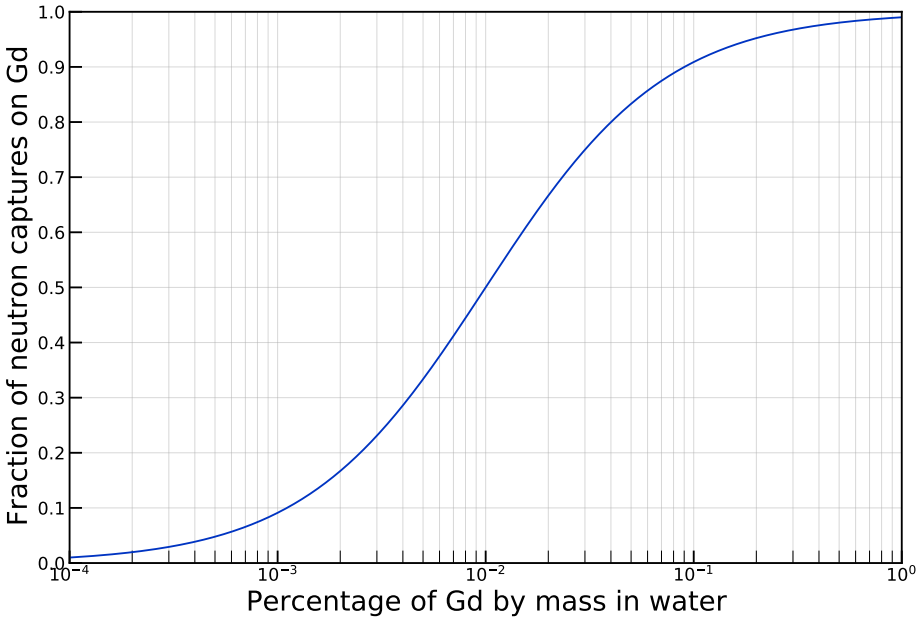

FIGURE 2.6: Fraction of neutron captures occurring on Gd as a function of the Gd concentration by mass in water. The nominal concentration target for SK is 0.1% Gd, corresponding to 90% neutron captures on Gd. Thanks to the non-linear nature of the function, most of the gains can be achieved with proportionally low concentrations: for 0.01% and 0.03% Gd concentration, the fraction of capture on Gd is 50% and 75%, respectively.

Before proceeding with the upgrade, a large-scale R&D program was conducted to test the idea with an SK-like detector system. The project, called EGADS (*Evaluating Gadolinium's Action on Detector System*) [47], was funded in 2009 and involved building a fully-functional 200-ton cylindrical water-Cherenkov detector in the Kamioka mine. The goals were to show that Gd sulfate could be efficiently and economically added to the detector, keeping a constant concentration with no negative impact on the water quality, and the detector components. In addition, it strived to demonstrate that the addition would not negatively affect existing SK analyses, and how to deal with possible new backgrounds introduced by the Gd. A large component of the project was a modification of the water purification system. The system described in Section 2.1, if left unchanged, would happily and efficiently remove the Gd along with any other impurity. Therefore, the process was modified to intelligently pick out the Gd from the purification stream and add it back to the water after it has undergone the main purification steps. The EGADS project proceeded with caution.

Phase	SK-I	SK-II	SK-III	SK-IV
Begin	Apr. 1996	Dec. 2002	July 2006	Sep. 2008
End	June 2001	Nov. 2005	Sep. 2008	June 2018
ID PMTs	11,146	5,182	11,129	11,129
Electronics	ATM	ATM	ATM	QBEE
Trigger	Hardware	Hardware	Hardware	Software
DSNB trigger	SHE	SHE	SHE	SHE+AFT
Water	pure	pure	pure	pure

Phase	SK-V	SK-VI	SK-VII	Total
Begin	Feb. 2019	July 2020	June 2022	Apr. 1996
End	July 2020	June 2022	(running)	(running)
ID PMTs	11,129	11,129	11,129	-
Electronics	QBEE	QBEE	QBEE	-
Trigger	Software	Software	Software	-
DSNB trigger	SHE+AFT	SHE+AFT	SHE+AFT	-
Water	pure	0.01% Gd	0.03% Gd	-

TABLE 2.2: Overview of the SK data-taking phases to date. Data taking phases at SK are signposted by significant changes in detector capabilities. The first of such changes was an unfortunate one, wherein the SK-I data-taking era stopped in 2001 due to an accident destroying over half of the PMTs in the tank. Data-taking resumed in January 2002 with SK-II, albeit with just 47% of the PMT coverage of SK-I. SK-III, starting July 2006, marked the return to a fully-functional detector, reclaiming its original PMT coverage. With SK-IV, starting September 2008, the data-taking and electronics systems were overhauled to the current setup, as described previously. Given are the start and end times, the number of ID PMTs, the type of electronics and trigger system, the trigger scheme used to search for DSNB events, the status of the water, and the total livetime available for SN burst searches.

The detector started running its water filtration system in mid-2011 with pure water. The collaboration then gradually instrumented and monitored the 200-ton tank until, by mid-2015, EGADS was a fully-fledged detector, with 240 PMTs and a water volume doped at 0.1% Gd. After more than two years of stable runtime, EGADS concluded that it was possible to achieve a constant Gd concentration, with no impact on water transparency, and no discernible damage to the components of the detector. Additional backgrounds due to impurities in the Gd were found to be reducible to acceptable amounts with specially-produced highly-radiopure Gd sulfate.

In 2018, the collaboration started refurbishing the SK detector in preparation of the Gd loading. Over two months in the summer of 2020, 13.2 tons of Gd sulfate were dissolved in the tank, reaching a Gd concentration of 0.01% on August 17th [48]. After three months of water recirculation, the water transparency stabilized to pre-loading levels. With this milestone, the experimented finally entered the *SK-Gd* era. Based on the previous success, a further 26 tons of Gd sulfate were added in May 2022, achieving 0.03% Gd concentration, and a 75% fraction of neutron captures on Gd. This brings us to the current state of the detector, at the time of writing, in the SK-VII data-taking phase, started in June 2022. An overview of the previous phases of SK, along with the main changes characterizing each one, is given in Table 2.2.

3
Data selection strategy for DSNB search with SK-IV

The first contribution of this thesis work was to the search of the DSNB using data from the SK-IV period, lasting almost 10 years. As mentioned in the previous chapter, a major improvement of this data-taking period over its predecessors is the improved throughput of its electronics. This improvement enables, among other targets, the storage of high-frequency, low-energy information and therefore allows for the search of weaker signals than would otherwise be possible. In particular, this improvement paved the way for neutron detection (or *neutron tagging*) to be possible in SK for the first time. The neutron, representing half of the detectable DSNB signal in SK, is key to maximizing our chances of detection. The next major improvement in the SK detector, the dissolution of Gd sulfate from SK-VI onwards, is of course the full realization of SK's neutron tagging dreams, by making the neutron not merely *detectable*, but *efficiently so*. Indeed, the neutron tagging efficiencies attainable in SK-IV (with pure water) are modest compared to those possible in the Gd era. Nonetheless, the SK-IV DSNB analysis demonstrates the feasibility of—and lays the foundations for—future searches. In this chapter, the SK-IV DSNB analysis is presented, with particular attention given to this thesis' contributions, and the role these might play in the future of the detector. First, the detection channel and all the relevant backgrounds sources are discussed in detail, along with the modelling strategy employed for each one, in Sections 3.1 and 3.2. Then, event reconstruction and preliminary quality cuts are described in Section 3.3. The second set of cuts is dedicated to the reduction of spallation events. While these cuts are used by the analysis presented here, their development and tuning was not part of this work. They are therefore sketched in Appendix A. For a more detailed description of the spallation backgrounds, their treatment and their reduction in the DSNB analysis, the reader can turn to Refs. [49, 50]. The third set of cuts, aimed at the characterization of the IBD positron, is described in Section 3.4. Finally, section 3.5 presents the analysis' neutron tagging strategy. The statistical treatment of the final data sample will then be given in Chapter 4.

3.1 Detection channel

Since neutrino emission from SN is evenly distributed across neutrino flavors, we could in principle look for several DSNB neutrino species, through a variety of interactions in the detector's water (see Figure 3.1) [51, 52]. In practice, the current sensitivity of SK allows for the search of SN neutrinos through one principal channel: the Inverse Beta Decay (IBD), already mentioned in Section 1.1. The second-highest cross-section—only a few times smaller than the IBD cross-section—comes from

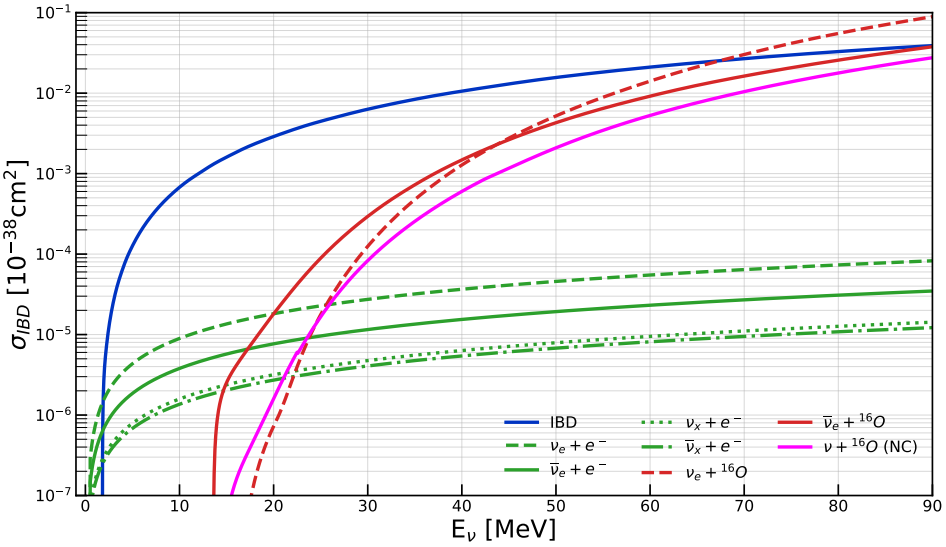

FIGURE 3.1: Summary of the cross-sections for the main processes used to detect SN neutrinos in water Cherenkov detectors. The neutrino can interact with the protons in the Hydrogen atoms via IBD (blue line), with the electrons in the Hydrogen atoms via charged current interactions (green lines), with the Oxygen-16 atoms via charged current interactions (red lines), or with the Oxygen-16 atoms via neutral current interactions (pink line). In this plot, ν_x refers to either ν_μ or ν_τ, and the NC cross-section on oxygen can refer to any ν or $\bar{\nu}$ species. The cross-sections are calculated using Ref. [52].

the coherent neutral-current scattering on protons [53, 54]. This process is however inaccessible at SK, as detection would entail the measurement of minute proton recoils, far below the detector's energy threshold. Within the energy range of the DSNB (below ∼30 MeV), the cross-sections of remaining processes are several orders of magnitude below that of the IBD. Potentially, a high-event-rate galactic SN burst could detect some of the more infrequent processes. For the rare DSNB signal, however, detection of an electron antineutrino via the IBD reaction will be our only way forward:

$$\bar{\nu}_e + p \rightarrow e^+ + n, \tag{3.1}$$

wherein a DSNB $\bar{\nu}_e$ interacts with a free proton in the water.

The advantages of IBD channel are manifold. Already mentioned is its (relatively) large cross-section. Protons are also of course copiously abundant in a water Cherenkov experiment (roughly 7.5×10^{33} protons in SK's 22.5-kton fiducial volume). In addition, its energy threshold E_{thr}, given by

$$E_\nu \geq E_{\text{thr}} = \frac{(m_n + m_e)^2 - m_p^2}{2m_p^2} \approx 1.806 \text{ MeV}, \tag{3.2}$$

is quite low, meaning that it has a relatively small impact on the visible DSNB energy range. Here E_ν is the neutrino energy while m_n, m_e, and m_p are the neutron, positron, and proton masses, respectively. The positron energy E_e is also strongly correlated with E_ν, as most of the energy is transferred to the positron, so that, in principle, the form of the DSNB spectrum is directly accessible. Finally, theoretical descriptions of the process are by now well-established, with accurate computations available down to the energy region of interest (the same cannot be said for neutrino reactions with ^{16}O, for example). Viable low-energy (∼ MeV scale) cross-section calculations were put forward by Beacom and Vogel in 1999 [56], and improved upon by Strumia and

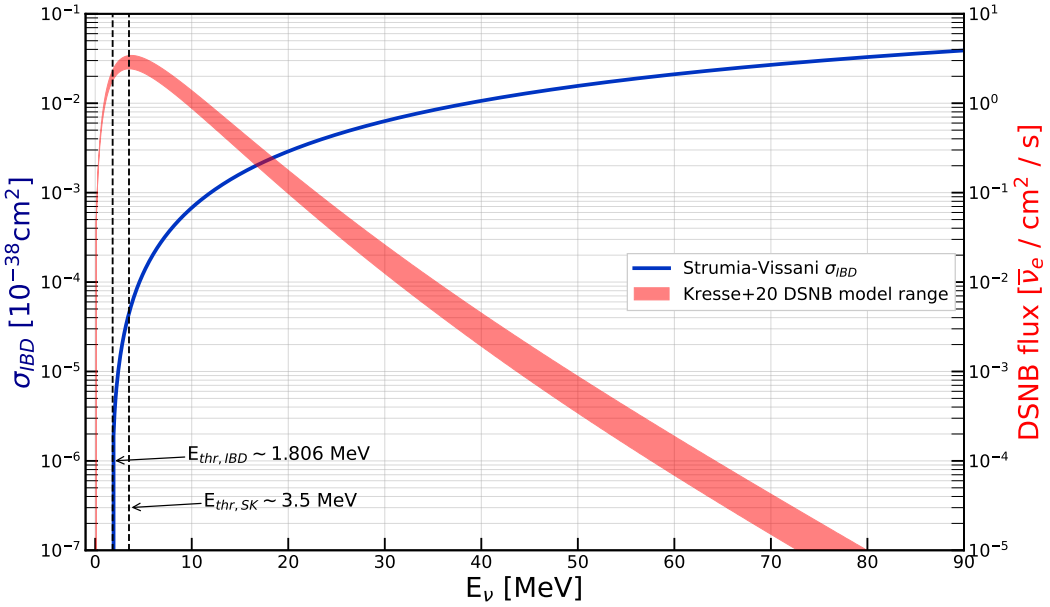

FIGURE 3.2: Strumia-Vissani cross-section [55] for the IBD reaction, along with the DSNB flux prediction range for the model from Kresse et al. [25]. Also shown are the energy thresholds for the IBD interaction, and for detection at SK.

Vissani in 2003 [55]. The latter's calculations are used in this work. The differential IBD cross-section $d\sigma/dE_e$, averaged over initial polarizations, as function of the neutrino and positron energies, is given by

$$\frac{d\sigma}{dE_e}(E_\nu, E_e) = \frac{G_F m_p \cos^2\theta_C}{\pi(s-m_p^2)^2}|\mathcal{M}^2| \qquad \text{if } E_\nu \geq E_{\text{thr}} \tag{3.3}$$

$$\mathcal{M} = \overline{v}_{\nu_e}\gamma^a(1-\gamma_5)v_e \cdot \overline{u}_n\left(f_1\gamma_a + g_1\gamma_a\gamma_5 + if_2\sigma_{ab}\frac{q^b}{2M} + g_2\frac{q_a}{M}\gamma_5\right)u_p. \tag{3.4}$$

Here, θ_C is the Cabibbo angle, σ_{ab} are the Pauli matrices, $M = (m_p + m_n)/2$, $s = (p_\nu + p_e)^2$, $q = p_\nu - p_e = p_n - p_p$, and f_i, g_i are adimensional form factors. Ref. [55] calculates $|\mathcal{M}|^2$ at NNLO in E_ν/m_p, including low-energy radiative corrections and a factor accounting for final-state interactions. The total cross-section, plotted in Fig. 3.2, has an overall 0.4% uncertainty, and is more than precise enough for our purposes.

The signature of an IBD interaction at SK is a characteristic double signal from the positron and neutron. The positron signal is straightforward: Cherenkov radiation is promptly emitted and readily detectable. The positron can be produced at all angles with respect to the neutrino, with a preference to the forward direction. It will carry away most of the neutrino energy, with the average positron energy $\langle E_e\rangle \approx E_\nu - m_n + m_p \approx E_\nu - 1.293$MeV, as a rough approximation. This will sometimes be referred to as the *prompt* signal, and results in a detector response of $\mathcal{O}(100)$ PMT hits. The neutron signature is trickier. No Cherenkov radiation is emitted by neutral particles, so no immediate signal is produced. The low-energy neutron will thermalize and scatter in the water, until it is finally captured by a Hydrogen atom. This thermal neutron capture will produce a Deuterium (^{2}H) atom with emission of a 2.224 MeV photon. The photon itself will typically undergo Compton scattering on an electron, allowing the detection of the Cherenkov radiation produced by this recoil electron. The thermal neutron capture cross-section on Hydrogen atoms is small enough for

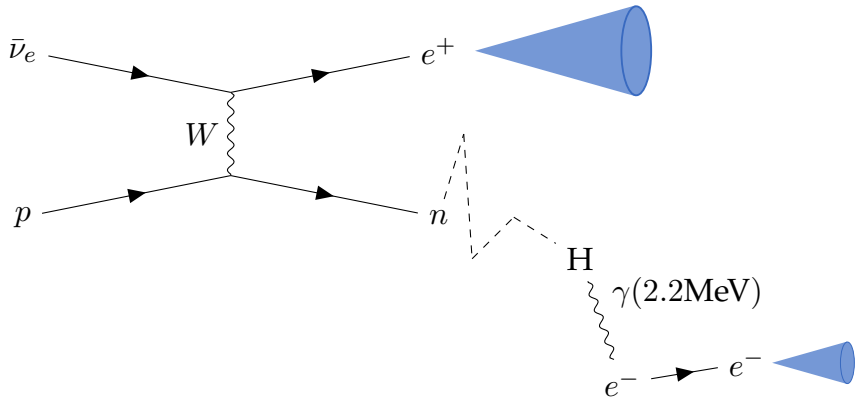

FIGURE 3.3: The detection channel for DSNB events at Super-K. The IBD process generates a positron and a neutron in its final state. The positron's Cherenkov emission (here represented as a blue cone) is directly detected. The neutron will thermalize and eventually be captured by a Hydrogen atom in the water, with emission of a 2.2 MeV photon. The detectable neutron signal will then come from the recoil electron from the Compton scattering of the photon. The resulting Cherenkov signal will be significantly weaker than the positron signal. Due to the neutron capture on H having a large characteristic timescale of $\tau \sim 200\mu$s, the two signals will appear with a significant delay between them.

the process to take longer than a typical SK event. The characteristic time of neutron capture on Hydrogen, $\tau_{\text{ncap},H} \approx 204\mu$s, will be a defining feature of our search. The neutron signal, also referred to as the *delayed* signal, will typically activate only < 10 PMTs. A diagram of the signal chain is given in Fig. 3.3.

In the IBD process the positron is emitted preferentially in the forward direction with respect to the neutrino. However, this angular dependence is not strong enough to allow for strong claims on the event's directionality. The differential cross-section from Ref. [55], as a function of the scattering angle θ between the neutrino and the positron, is given by

$$\frac{d\sigma}{d\cos\theta}(E_\nu, \cos\theta) = \frac{p_e \frac{E_\nu}{m_p}}{1 + \frac{E_\nu}{m_p}(1 - \frac{E_e}{p_e}\cos\theta)} \frac{d\sigma}{dE_e}. \tag{3.5}$$

Above about 2.8 MeV, all directions ($\theta \in [0, \pi]$) are allowed, with only a slight forward tendency. This forward/backward asymmetry rises with neutrino energy, but it is still not useful for single-event information. Only with large samples of supernova neutrinos from a single explosion, as in the case of a nearby galactic burst, could one infer a direction, on a statistical basis. Note that in the case of a burst event, while the IBD events could potentially constrain the source direction, most of the directional information would still come from strongly forward-peaked interactions with Oxygen.

To calculate the DSNB event rate expectations at SK, plotted in Fig. 3.4, the DSNB flux models such as the ones presented in Section 1.2.1 are convolved with the IBD interaction cross-section and the energy resolution of the detector. Due to the variety of flux models, the range of expected number of IBD events roughly spans an order of magnitude, from ∼0.5 to 5 per year for modern predictions.

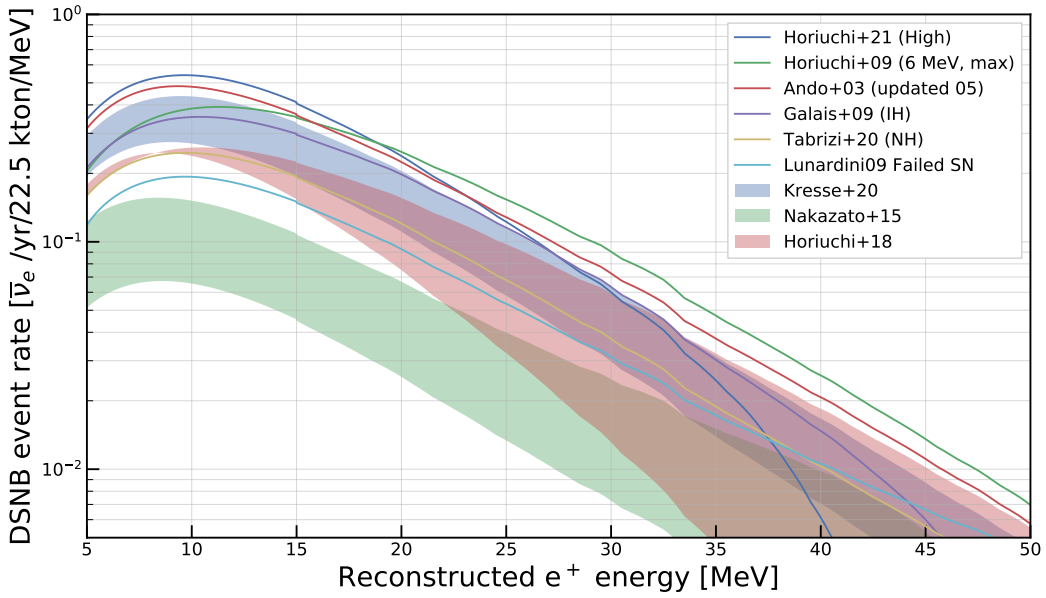

FIGURE 3.4: Expected DSNB rate at in the 22.5-kton fiducial volume of SK, as a function of reconstructed positron energy. Here, each DSNB flux is model is convolved with the Strumia-Vissani cross-section. To model the detector response, a Gaussian smearing corresponding to the energy resolution of SK-IV is applied.

SHE threshold	AFT duration	E_{thr}	livetime
70 hits	350 μs	10 MeV	25.0 days
70 hits	500 μs	10 MeV	869.8 days
58 hits	500 μs	8 MeV	2075.3 days

TABLE 3.1: Variation of the SHE trigger threshold and AFT event duration over time in the SK-IV data sample.

Data sample

To search for the DSNB through the IBD signature, we make use of all the available data from the SK-IV period, lasting nearly a decade, from September 2008 to June 2018. After accounting for detector downtime and calibration work, test runs, and otherwise poor quality data, we are left with roughly 8 years of useful data for physics analyses (2970.1 days, to be exact), including SK data runs in the run number range [61525, 77958]. The data consists of events triggered with the SHE+AFT trigger scheme. The prompt signal triggers an SHE event, comprising data from 5 μs before to 35 μs after the trigger time. In order to detect the delayed event, the longer AFT event is also saved, starting immediately after the end of the SHE event. For the majority of the SK-IV runtime, the PMT hit threshold for an SHE event was 58 hits in 200 ns, and the AFT event duration was 500 μs. These values have however been subject to change over time (see Table 3.1). In particular, the SHE threshold, initially set to 70 hits, was only lowered to 58 hits in early 2011. The hit threshold represents the hard energy threshold for the analysis, which is either 10 MeV or 8 MeV, depending on the time period.

IBD Monte Carlo simulation

Simulating the DSNB signal starts with a model for the kinematics of the IBD interaction. Since the DSNB is an isotropic signal, IBD vertices are randomly generated with uniform distribution in the entire SK volume. In the analysis detailed later, the section of the volume that is less than 2 m away from the detector wall will be ignored to minimize radioactive backgrounds. Nevertheless, it is important to simulate our signal in the entire volume to include events that may be reconstructed inside this *fiducial volume* despite originating outside of it.

Each IBD vertex is then associated with a positron and a neutron. Working in the reference frame of SK, in which the IBD proton is stationary, the four-momenta p_ν, p_e and p_n, of the neutrino, positron and neutron respectively, can be expressed as

$$p_\nu = (\mathbf{p}_\nu, \quad 0, \quad 0, \quad \mathbf{p}_\nu) \tag{3.6}$$

$$p_e = (E_e, \quad \mathbf{p_e} \sin\theta \cos\phi, \quad \mathbf{p_e} \sin\theta \sin\phi, \quad \mathbf{p_e} \cos\theta) \tag{3.7}$$

$$p_n = (E_n, \quad \mathbf{p_n} \sin\theta' \cos\phi, \quad \mathbf{p_n} \sin\theta' \sin\phi, \quad \mathbf{p_n} \cos\theta'). \tag{3.8}$$

Here, we have aligned the z axis in the neutrino direction, and defined the scattering angles θ (between the neutrino and the positron), θ' (between the neutrino and the neutron), and ϕ (the azimuthal rotation around z). The neutrino and positron momenta are related by

$$\mathbf{p}_\nu = \frac{2m_p E_e + (m_n - m_p)^2 + 2m_p(m_n - m_p) - 2m_e^2}{2(m_p - E_e + \mathbf{p}_e \cos\theta)}. \tag{3.9}$$

The IBD interaction itself is not unique to the DSNB signal. As will become clear in the following section, neutrinos from nuclear reactors, a background for our analysis, also interact in this fashion. In addition, neutron-producing β decays from certain spallation isotopes have a similar final state with analogous kinematics. The energy spectrum of the DSNB signal itself also varies according to flux model chosen. For these reasons, a generalized simulation of the IBD vertices can be reused for the aforementioned backgrounds and for any DSNB flux model. To this end, the IBD vertices are generated with uniform E_e, θ, θ', and ϕ distributions, with E_e ranging from 1 to 90 MeV. An appropriate reweighting will then be applied to the simulated events to match the desired spectrum and angular distribution. In the case of the DSNB signal, the positron energy spectrum is obtained by convolving the Strumia-Vissani IBD cross-section with the specific DSNB flux model.

The positron and neutron final states are then used as input to a detector simulation, modeling subsequent interactions in the detector volume, Cherenkov light emission and propagation, and the response of the SK PMTs and electronics. The dedicated SK simulation package is *SKDetSim*. This software package is based on the GEANT-3 framework, and is a tuned by a dedicated group inside the SK collaboration to reproduce a wide range of detector effects observed during calibration data-taking runs. These effects include, among others, variations in gain and dark noise intrinsic to the PMTs, asymmetries in the detector (in particular top-bottom asymmetries), and variations in the water transparency. These features are tuned in a time-dependent fashion, so that each data run is assigned a slightly different tuning from the next. Since the DSNB flux is time-independent, this is where time-dependence is first introduced in the simulation. Each available data run in the SK-IV period is simulated, namely SK data runs in the range [061525, 077958]. In doing so, the observables of the simulated sample will reflect their expectation values over the SK-IV running time.

3.2 Background sources

At low energies, background sources that can mask the DSNB are both varied and abundant. Working near the edge of SK's energy range means facing considerable data selection and reduction challenges. As we will see, highly-performant algorithms will therefore be necessary to tease out our faint signal. Our approach to the modeling of each background depends on the available simulation software and data, as well as our theoretical understanding of the background source in question. Unlike the IBD interaction, which is well-understood, some backgrounds are, at present, particularly difficult to simulate. Our background models will therefore be a patchwork of Monte Carlo simulations and data-driven strategies. In this section, the background sources considered in the analysis, the challenges they pose, and the models used to describe them, are detailed.

3.2.1 Detector backgrounds

Background events due to low-energy effects inherent to the SK detector itself are plentiful. First and foremost is the noise in the PMTs, amounting to about 4kHz per PMT. In addition, impurities in the water are the source of additional backgrounds, in particular through radioactive decay of Rn isotopes. Despite the experiment's excellent water purification system, the remaining impurities still produce visible background activity. Radiation in the PMTs of the detector, in its walls, and even in the rock surrounding it, is also responsible for low-energy backgrounds. Finally, so-called *flasher* events, in which neighboring PMTs are observed to, on occasion, fire simultaneously without a corresponding incoming photon. These kinds of backgrounds can be separated from the prompt event fairly simply, so that they do not pose a serious threat to the identification of the IBD positron. However, they are much more difficult to differentiate from the faint delayed signal. When such backgrounds are misidentified as a neutron signature, we talk of *accidental coincidences*. The performance of neutron tagging will be crucial to their removal.

Accidental coincidence modelling

An accurate account of low-energy backgrounds is needed to develop a robust neutron tagging. At the time of writing, the SK simulation software does not fulfil this requirement. While SKDetSim is capable of simulating the PMT dark noise, it provides a simplified picture of the low-energy backgrounds, missing their subtleties and multifaceted nature. By and large, the noise simulation in SKDetSim was made to contaminate existing detector signatures, and not to be studied as a background *per se*. A neutron tagger however needs to characterize the properties of accidental coincidences in their own right to distinguish them from the neutron signature.

For studying the positron signature, SKDetsim's noise simulation is sufficient, so it is used for contaminating simulated primary events. Rather than using simulation to also contaminate the delayed events, accidental coincidence backgrounds are drawn directly from data. In particular, they are modelled using data samples obtained with the so-called *dummy* trigger. These are 1-ms events that are triggered at random times. Since they are not associated with a primary signal, they are ideal for modelling accidental coincidences, which are totally uncorrelated with the primary event. Dummy trigger data is therefore overlaid onto the simulated IBD signal, in the time window where the neutron search will take place. This allows us to obtain a complete picture of the detector response.

3.2.2 Solar neutrinos

By far the source of the largest neutrino flux on Earth is our own Sun. As a Main Sequence star in the middle of its lifetime, it supports its mass through constant fusion reactions inside its core. According to the standard solar model, the main reaction chain contributing to neutrino production is the *pp chain* responsible for the conversion of Hydrogen into Helium, starting with the fusion of two protons:

$$p + p \rightarrow {}^{2}\mathrm{H} + e^{+} + \nu_e. \qquad (pp) \tag{3.10}$$

This reaction alone represents the source of more than 90% of all solar neutrinos. Nevertheless, these ν_e are far too low in energy (< 1 MeV) to be visible in SK. Detectable neutrinos are produced through further fusion reactions involving the deuteron:

$${}^{2}\mathrm{H} + p \rightarrow {}^{3}\mathrm{He} + \gamma. \tag{3.11}$$

Fusion of He atoms and protons then produces *hep* neutrinos:

$${}^{3}\mathrm{He} + p \rightarrow {}^{4}\mathrm{He} + e^{+} + \nu_e, \qquad (hep) \tag{3.12}$$

this time well above the detectable energy threshold of SK, peaking around 10 MeV. Finally, ^{8}B neutrinos are produced through fusion of Helium atoms into heavier elements, namely Beryllium and Boron. The latter will then produce a ν_e upon beta decay:

$${}^{3}\mathrm{He} + {}^{4}\mathrm{He} \rightarrow {}^{7}\mathrm{Be} + \gamma \tag{3.13}$$

$${}^{7}\mathrm{Be} + p \rightarrow {}^{8}\mathrm{B} + \gamma \tag{3.14}$$

$${}^{8}\mathrm{B} \rightarrow {}^{8}\mathrm{Be}^{*} + e^{+} + \nu_e \qquad ({}^{8}B) \tag{3.15}$$

These neutrinos comprise the bulk of the solar neutrinos, representing a flux roughly three orders of magnitude larger than that of *hep* neutrinos. The pp chain of neutrino-producing reactions are summarized in Fig. 3.5, while the neutrino fluxes from each reaction are sketched in Fig. 3.6.

Solar ν_e events at SK occur primarily via neutral current scattering on electrons. Despite the relatively low cross-section, they are visible due to their enormous flux, around the lower end of the DSNB energy range. Thankfully, solar neutrinos are well-studied in Super-Kamiokande, and fairly straightforward to remove. The scattering reaction is strongly forward, so that the direction of the incoming neutrinos can be inferred on an event-by-event basis and cross-checked against the current location of the Sun in the sky. In addition, no neutrons are produced in the final state, so a good neutron tagging algorithm should also independently take care of them.

3.2.3 Reactor neutrinos

While the Sun is the chief natural source of neutrino flux on Earth, nuclear reactors are the dominant man-made source. In a nuclear reactor, energy is extracted from the fission reactions of heavy atoms (usually Uranium or Plutonium isotopes). When the fissile isotope is induced to undergo fission, the nucleus splits into lighter daughter nuclei. These, in turn, can decay to more stable species via beta decay with emission of an electron antineutrino ($n + e^{-} \rightarrow p + \overline{\nu}_e$). A single fission, along with the decay of the subsequent fissile products, releases about 200 MeV of energy. While most of it

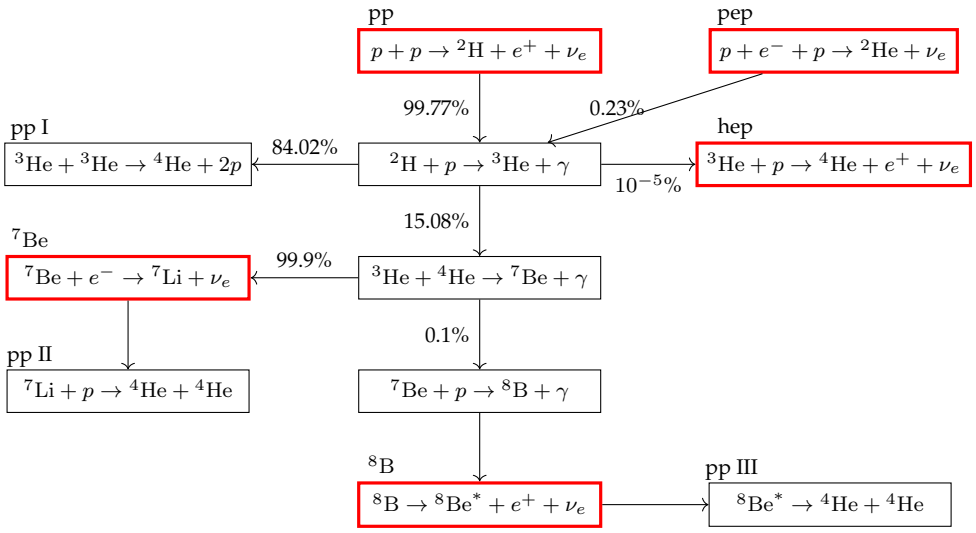

FIGURE 3.5: Schematic of the *pp chain*, the main group of fusion reactions responsible for the generation of Helium atoms from Hydrogen atoms. Highlighted in red are the neutrino-generating processes. The pp chain is composed of four main branches, the *ppI*, *ppII*, *ppIIII*, and *hep* branches. A separate set of Hydrogen fusion reactions exists, namely the CNO cycle, not pictured here. The CNO cycle produces no neutrinos that would be visible by SK.

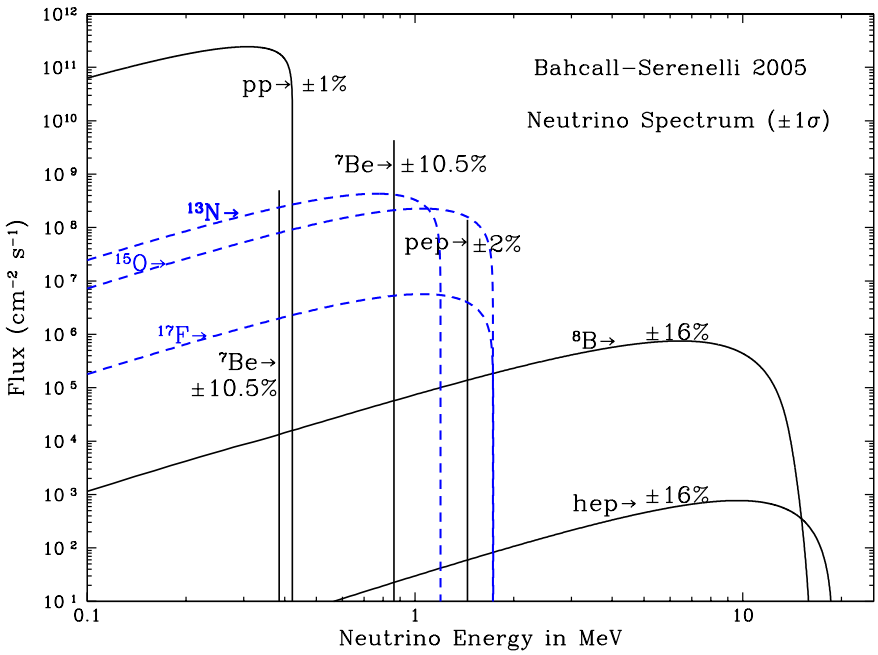

FIGURE 3.6: Solar neutrino fluxes on Earth, separated among the different neutrino production reactions. The dashed blue lines, all well below the SK energy threshold, are from the CNO cycle. The solid vertical lines, also below threshold, correspond to the only allowed energies for the relevant two-body decay in the pp chain. The remaining pp chain fluxes are shown as solid curves. Of these, only ^8B neutrinos and *hep* neutrinos will be visible at SK. Figure from Ref. [57].

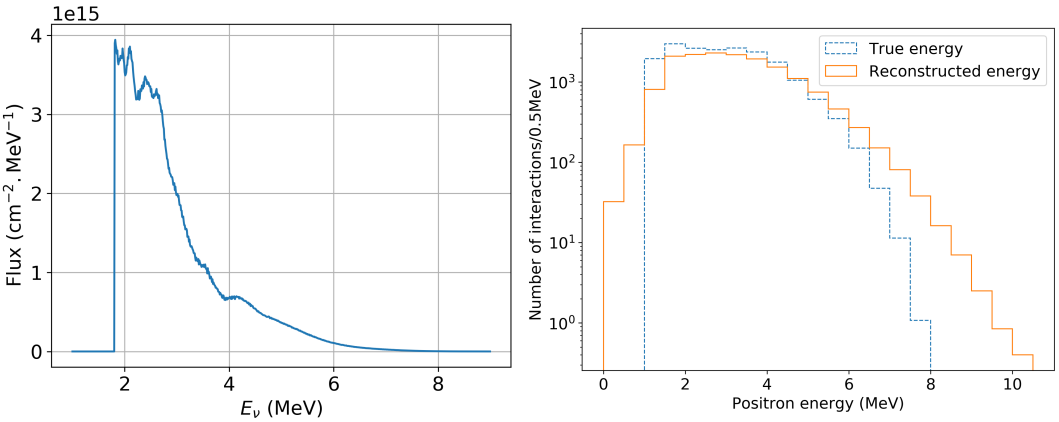

FIGURE 3.7: Left: reactor neutrino flux at Super-Kamiokande over the runtime of SK-IV. The different peaks reflect the variety of reactors in the relative vicinity of the experiment. Right: expected event number of reactor neutrino events over the runtime of SK-IV, as a function of visible (positron) energy. In blue is the true energy spectrum, while in orange is the reconstructed energy spectrum. These events are irreducible, so their presence raises the effective analysis threshold above the nominal SK threshold. While the true reactor spectrum is virtually non-existent above 8 MeV, due to the intrinsic resolution of SK, we expect significant contributions from this background until about 10 MeV.

remains inside the reactor as heat (to be extracted downstream), about 5% escapes, carried away by the antineutrinos.

Reactor neutrinos played a key role in the history of neutrino observation. In fact, they were the first kind of neutrinos to ever be detected, with Cowan-Reines experiment in 1956. Due to their predictable and plentiful flux, they continue to be used for precision measurements in neutrino physics by a long list of experiments. For our search, however, they are unfortunately a nuisance. The exact spectrum of detectable reactor neutrinos depends on the fissile isotopes used, and, of course, on one's proximity to the reactor in question. Nuclear power makes up a considerable portion of Japan's electricity production (albeit this has decreased significantly since the 2011 accident in Fukushima), and several reactors are located within relative proximity of Super-Kamiokande. While the reactors neutrinos spectrum is sharply peaked at very low energies, the tail of the distribution will still impact the analysis through resolution effects, and because of the sheer magnitude of the incoming reactor flux. As a consequence, the lower end of the DSNB spectrum is completely overwhelmed by these events. Since they are $\overline{\nu}_e$, they are irreducible. They will then effectively set the lower energy threshold of our analysis, below which the significance of our signal will be negligible.

Simulating the reactor background is a matter of reweighting the IBD simulation, as detailed above. The reactor flux used is generated by the software tool *SKReact* [58], developed for the detailed study of reactor neutrinos with SK. The reactor model used is from Ref [59], itself based on IAEA records [60]. This tool can calculate the reactor flux from any nuclear reactor in the world, up to a maximum neutrino energy $E_\nu^{max} = 9.2$ MeV, taking into account each detector's time-varying activity, and applies oscillation effects. Constraints from the Daya Bay experiment [61] confirm that any contribution above E_ν^{max} can be neglected. The incoming reactor neutrino flux, and the resulting event rate (both real and reconstructed) are shown in Figure 3.7. Above a positron energy of 6 MeV, most of the contribution comes from misreconstructed events, and the reactor neutrino rate will still exceed the DSNB by roughly two orders of magnitude, setting the most optimistic search threshold at around

8 MeV.

3.2.4 Radioactive spallation backgrounds

The next background source is one of the crucial ones to master for our analysis. While reducible in principle, it will pose significant algorithmic and computational challenges in practice. The background in question is a family of processes referred to as *radioactive spallation*. In general terms, *spallation* denotes a reaction in which debris (or *spall*) is ejected from a target as a result of a collision with a projectile. For our purposes, a spallation reaction starts with a highly-energetic muon (the "projectile") penetrating the experiment's rock cover to reach the detector's water (the "target"). While inside SK's tank, the muon can transfer a significant amount of its energy to destabilize the atoms in the water, forming unstable (radioactive) isotopes. The decay of these unstable species will then cause visible light emission at low energies, easily mistakable for the signature of a positron.

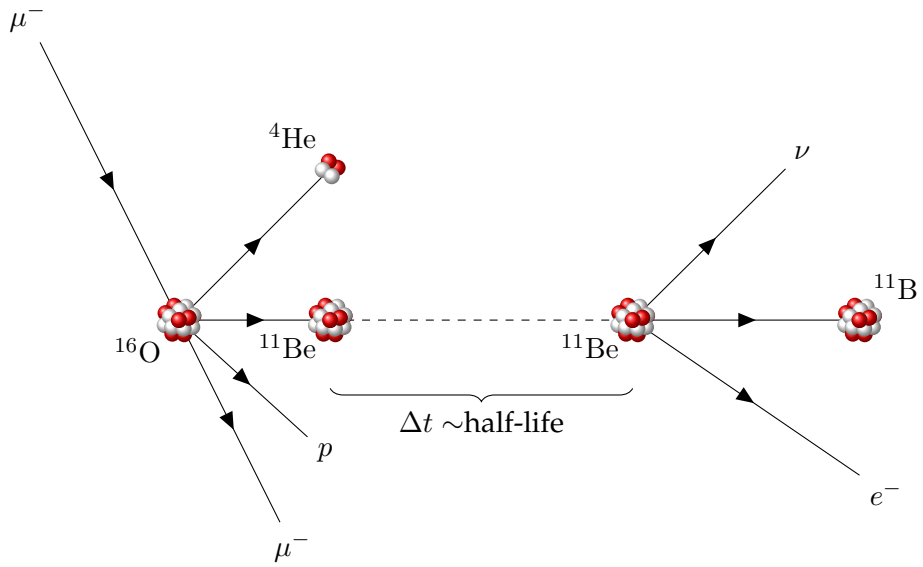

FIGURE 3.8: Simplified picture of a spallation event induced by an incoming cosmic muon, incident on an Oxygen atom in the water. In this example, an energetic muon causes the fission of an Oxygen atom into Helium and an unstable isotope of Beryllium. The latter will undergo β-decay, with a time delay according to the isotope's half-life. The β-decay electron then mimics the prompt signal from the IBD interaction.

The muons themselves are easily recognizable. Especially at higher energies, their high momentum produces hallmark, clear Cherenkov rings that allow for accurate reconstruction of energy and track direction. To check whether a given signal candidate is a result of spallation, then, we have to cross-check the set of recent muon events to determine whether these are compatible with our signal. The first challenge to this approach comes from the fact that many of the isotopes produced during spallation are remarkably long-lived (see Fig. 3.9). The lifetime of ^{16}N, for instance, can be measured in *seconds*—an eternity compared to the μs-scale SK event. Effectively, this means that for each DSNB candidate, we will have to cast a very wide net, looking for potential spallation muons in a large time slice preceding our event. The number of muon-to-DSNB candidate pairs, and the time required to analyze them, can quickly balloon out of control. Particularly long-lived isotope can also potentially escape our detection.

In addition, simulating spallation reactions is delicate, due to the involved nuclear physics, which is associated with significant uncertainties. Simulation of spallation processes is therefore an important research target at SK. While significant headway has been made, at the time of the analysis no complete spallation simulation was available. This impairs our ability to manage and make predictions about this particular source of background. Having an accurate simulation would potentially allow, for example, to model the spallation decay energy spectrum with manageable systematic uncertainties. Since this was not yet possible, the approach had to shift to eliminating the background as much as possible, even at the cost of relatively harsh cuts.

Notably, most spallation-induced decays do not include neutrons in the final state. Neutron tagging will therefore be part of the arsenal at our disposal to tackle radioactive spallation. Due to their abundance, some of these decays will typically still be present even after a performant neutron tagger, so this cannot be our only strategy. Just as important, the $^{9}\mathrm{Li}$ isotope, decaying in one of two modes, can produce a neutron:

$$^{9}\mathrm{Li} \xrightarrow{178.2\ \mathrm{ms}} {}^{9}\mathrm{Be} + e^{-} \tag{3.16}$$

$$^{9}\mathrm{Li} \xrightarrow{178.2\ \mathrm{ms}} {}^{8}\mathrm{Be} + e^{-} + n. \tag{3.17}$$

The $^{9}\mathrm{Be}$ isotope is stable, and $^{8}\mathrm{Be}$ quickly decays into two stable $^{4}\mathrm{He}$ atoms, so neither produces a recognizable signal. While the former of the $^{9}\mathrm{Li}$ decays will have

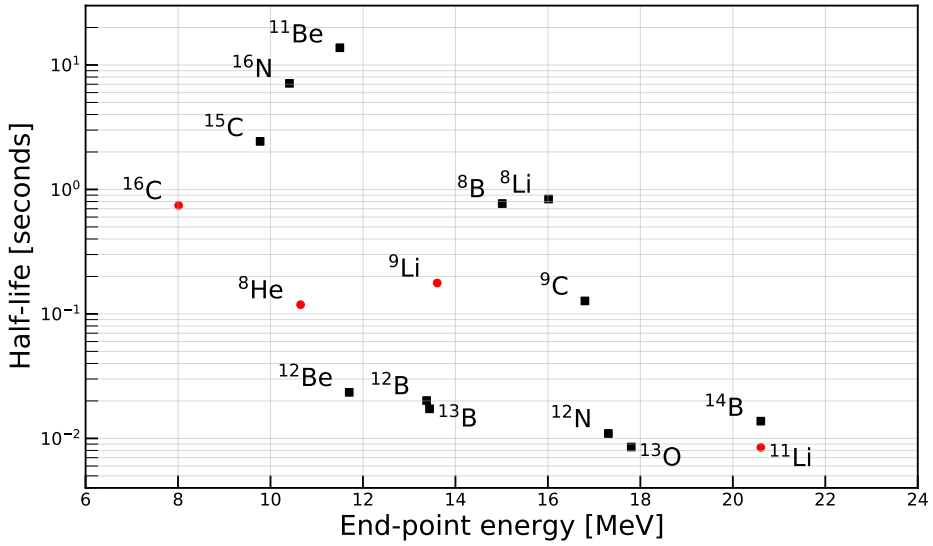

FIGURE 3.9: Unstable isotopes produced during radioactive spallation processes induced by cosmic muons interacting in the SK water. Here, each isotope is plotted according to its half-life and the end-point of the decay product's energy spectrum. At low half-lives the time difference between the offending muon and the DSNB candidate is low, and the correlation clear, meaning these events can be usually eliminated. The end-point energy determines how much of the energy spectrum is contaminated by spallation events. Note that the impact of a given isotope extends beyond its end-point energy due to resolution effects. The isotopes in red are the ones that include a neutron in the final state of the decay, and are therefore irreducible, even with the inclusion of a neutron tagging cut. Of these, $^{9}\mathrm{Li}$ is the only one representing a threat to the analysis, due to its long lifetime and high end-point energy.

a straightforward electron signature, the latter will also include a neutron. Since an electron produces the same Cherenkov signal as a positron, this decay mode will mimic both prongs of the IBD coincidence signal and cannot be removed by neutron tagging alone.

3.2.5 Atmospheric neutrinos

The last background source to discuss is from neutrinos generated in the Earth's atmosphere. Atmospheric neutrinos are produced when a cosmic ray, typically a proton, interacts with the atmosphere, producing hadronic showers and a stream of neutrinos of all flavors. This encompasses a broad range of processes, with neutron energies spanning $\sim$ 6 orders of magnitude, from the MeV scale up to the TeV scale. It is important to take into account high-energy atmospheric neutrinos, as these can leak into our energy range due to the reconstruction of low-energy decay products, secondary particles, and misreconstruction effects. In the following, the main atmospheric contributions contaminating our energy window are discussed.

Charged-current ν_μ events

Representing the main contribution above $\sim$ 15 MeV, ν_μ's and $\overline{\nu}_\mu$'s interact in the water with production of a muon, which then decays weakly, producing an electron:

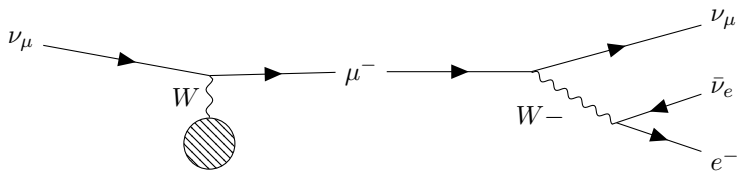

(3.18)

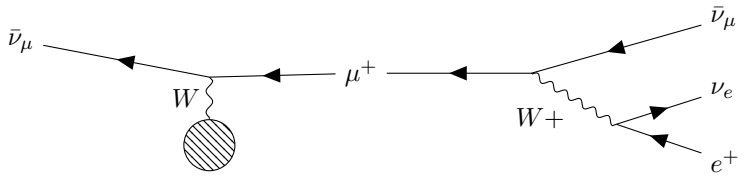

(3.19)

where the shaded circle stands for the range of possible nuclear interactions in the water. If the muon is above the energy threshold for Cherenkov light emission, it can trigger an event and be mistaken for the prompt signal. In this case, an unusually high level of PMT activity is observed, due to the signal from both the muon and electron. If the muon is below threshold, it will be undetectable (hence we talk of *invisible muons*) and the final state will mimic an IBD prompt event. These electrons, sometimes referred to as *Michel electrons* are well-studied, with a well-known energy spectrum, and are mostly found between $\sim$15 and 60 MeV. The nuclear interaction portion of the process is more complex and will not be discussed in detail in this text. One important detail, however, is that these nuclear interactions can produce one or more secondary neutrons. These can be detected in coincidence with the rest of the event, and are therefore a source of irreducible backgrounds.

Charged-current ν_e events

Incoming atmospheric $\bar{\nu}_e$ events, interacting via charged-current, will produce a positron in the final state, mimicking the prompt signal. If the interaction is accompanied by the production of a neutron, both prompt and delayed signals are mimicked.

On an event-by-event basis, these kinds of events are nearly indistinguishable from the DSNB. The same goes for atmospheric charged-current ν_e events, producing an electron in the final state in a similar process. On a statistical basis, the charged-current ν_e energy spectrum is however different from our signal, which increases, within our energy range, roughly linearly as a function of the electron energy.

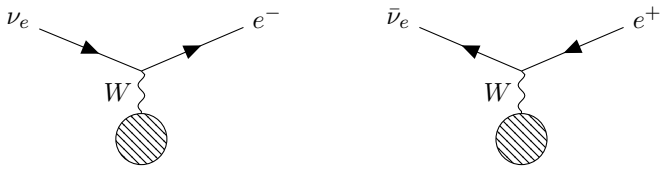

(3.20)

Neutral-current quasi-elastic interactions

Neutral current interactions of atmospheric neutrinos are only detectable due to the nuclear interactions induced by the exchange of the Z boson with the water's nuclei. Most commonly, these are neutral current quasi-elastic (NCQE) interactions with ^{16}O nuclei. The visible signature of this background comes from a primary photon produced in the NCQE interaction itself, and a shower of secondary photons from additional nuclear interactions induced by the production and exchange of nucleons in these reactions. The diagram below gives an example of such an interaction, where the unmarked fermion lines represent nucleons.

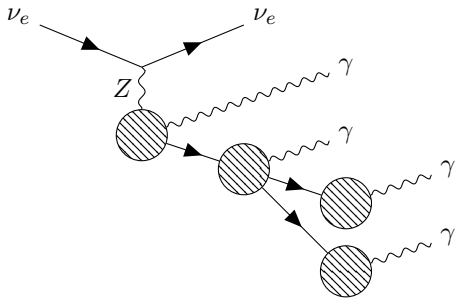

(3.21)

One or more of the photons can be misreconstructed as our prompt positron. The signature of the secondary photons, as well as any real neutron accompanying the nuclear interactions, can also cause a delayed signal. Due to the fact that the NCQE signature is dominated by complex nuclear interactions, it is more difficult to model. In addition, this particular background peaks at lower energies (< 20MeV), making it especially bothersome for our search.

Pion-producing events

Nuclear interactions induced by atmospheric neutrinos in any of the above categories can also produce pions. If the final state particles from the rest of the interaction are below threshold, or otherwise low in energy, the pion itself can trigger a prompt event. These types of backgrounds are fairly unproblematic, as their Cherenkov signature can be distinguished from an electron signature thanks to their higher mass and momentum.

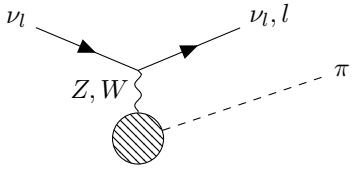

(3.22)

In the case of charged pions, if the pion itself is also below threshold, a prompt event can be triggered by its decay. In particular, the main decay channel of the charged pion,

$$\pi^- \to \mu^- + \overline{\nu}_e \tag{3.23}$$
$$\pi^+ \to \mu^+ + \nu_e, \tag{3.24}$$

produces a muon in its final state. This muon can contribute to the atmospheric backgrounds in the same way as the prompt muons from charged-current ν_μ interactions.

Atmospheric neutrino simulation

Unlike for other background sources, simulation samples for atmospheric neutrinos are readily available in SK, as they are computed by a dedicated group inside the collaboration. The reference atmospheric neutrino flux used in this simulation is the HKKM-2011 flux, calculated by Honda, Kajita, Kasahara, and Midorikawa [62] for a variety of neutrino detector locations. This calculation, constrained by experimental results, accounts for the location of SK and the detector's rock cover. The flux is used as input to the NEUT 5.3.6 [63] event generator, which computes the nuclear interactions in the water and the corresponding final states. These are then processed by SKDetSim, as with the IBD simulation, to compute the detector response. In this simulation, the prompt event is contaminated with simulated noise, as in the IBD simulation. Here, too, no noise is simulated for the delayed event. Instead, we add dummy-trigger noise from data in the neutron search window for a proper account of the accidental coincidences.

3.3 Event reconstruction and preliminary reduction

In this section, the first reduction steps applied to the data and background simulation samples detailed previously are described. Namely, the prompt event is reconstructed, and preliminary reduction cuts are applied to remove some of the more straightforward background events. The reference algorithm for reconstructing low-energy events (< 100 MeV) at SK is BONSAI [64] (short for Branch Optimization Navigating Successive AnnealIng operations). The best-fit primary vertex is found by maximizing the likelihood of the PMT hit times, with the Cherenkov opening angle of the event constraining the best-fit event direction. The software owes its name to its intelligent event vertex search approach. Rather than naively scanning across the entire detector, BONSAI starts with an initial vertex guess and progressively searches the detector by building and traversing a search tree, with new vertex guesses branching off from the previous ones. For optimal performance, the size of the tree is limited by pruning ineffective branches. Once the best-fit vertex and track direction are found, the energy of the event is estimated by calculating the effective number hits N_{eff} in a 50-ns time-of-flight-subtracted time window, defined as

$$N_{\text{eff}} = \sum_{i}^{N_{50}} \left[(X_i + \epsilon_{\text{tail}} - \epsilon_{\text{dark}}) \times \frac{N_{\text{all}}}{N_{\text{normal}}} \times C_i \times e^{(r_i/L_{\text{att}})} \times \frac{1}{QE_i} \right], \tag{3.25}$$

where each PMT hit i out of the N_{50} hits in the time window is assigned a weight based on a variety of parameters. In particular:

- X_i is the occupancy factor. For higher-energy Cherenkov events, some PMTs will detect multiple photoelectrons, in which case neighboring PMTs are likely to also be activated. X_i then is a measure of the probability for the observed number of photoelectrons in the ith PMT, given the activation of the neighboring PMTs, and the total number of PMTs in the event.
- ϵ_{tail} is the correction term accounting for the proportion of particularly-delayed signal hits that are expected to be found outside the 50-ns window. The delay can be caused by a combination of scattering, reflection, as well as absorption and re-emission effects.
- ϵ_{dark} is the correction term for the expected contribution of dark noise hits in the time window, and is measured in data with a dummy-trigger setup.
- $\frac{N_{\text{all}}}{N_{\text{normal}}}$ is the ratio of the number of total PMTs to the number of normally-functioning PMTs, and corrects for the detection inefficiency due to malfunctioning PMTs in the current data run.
- C_i is a factor correcting for the inefficiency due to the effective photocoverage of PMT i, taking into account the incident hit angle.
- $\exp(\frac{r_i}{L_{\text{att}}})$ corrects for the unobservable photons that get absorbed in water, for PMT at distance r_i from the vertex, according to the characteristic attenuation length L_{att}. The attenuation length is measured run-by-run using calibration data.
- $\frac{1}{QE_i}$ corrects for the gain and quantum efficiency of PMT i, and is also measured in calibration.

The value of N_{eff} is therefore an estimate of the true number of Cherenkov photons produced in the event. The energy of the event is reconstructed starting from N_{eff} by applying a polynomial function, whose parameters are tuned during calibration. More information about BONSAI can be found in Ref. [64]. In this analysis, BONSAI 3.3 is used to reconstruct the positron event in the SHE event window. Note that the prompt event energy obtained in this way is the visible energy of the event. Its relationship to the kinetic energy of the particle emitting the Cherenkov radiation depends on the particle in question. Notably, the visible energy from a positron will include the 1.022 MeV annihilation energy, unlike in the case of an electron. For the sake simplicity and consistency, the term *reconstructed energy* will refer to the BONSAI-reconstructed visible energy.

Muon fitting

In parallel to the event reconstruction of the SHE-triggered events, we reconstruct muon events that could be correlated with the prompt event. This step is necessary to identify any muon associated with a spallation-induced prompt event. Due to the potentially long decay timescales of the spallation isotopes, we are forced to look for muons very far in advance of our prompt event. Indeed, the muon search extends to 30 s before each SHE trigger. In addition, uncorrelated muon events are also searched for, by analyzing events within 30 s *after* the SHE trigger (refer to Appendix A for details on the role of these muon samples). For each event in this timeframe, an SLE software trigger is applied across the event to find candidate sub-events, selecting sub-events with more than 500 photoelectrons deposited inside the ID.

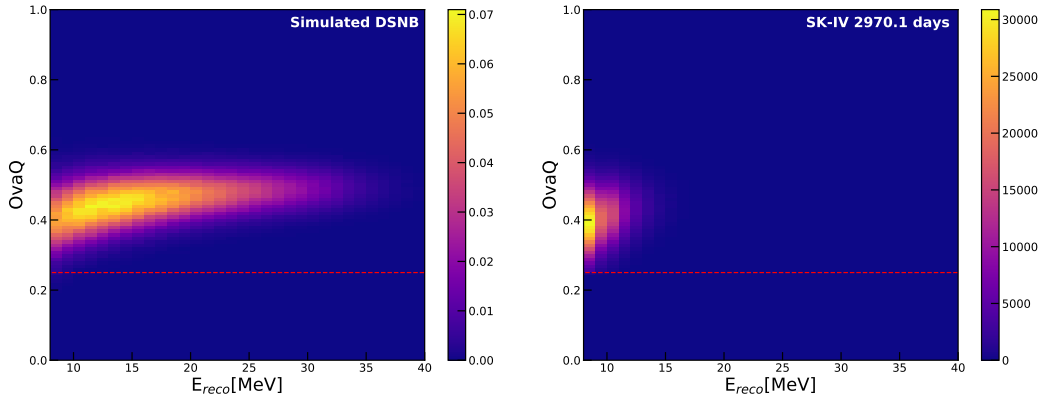

FIGURE 3.10: Distribution of the fit goodness variable $(\text{BONSAI goodness})^2 - (\text{dirKS})^2$ (here labeled as *ovaQ*) as a function of reconstructed energy. On the left, the expected number of events from a simulated DSNB signal (here, the Horiuchi [20] model), for the SK-IV running time. On the right, the events in our data sample. The red dashed line indicates the cut point. The vast majority of the data occurs at low energies, primarily from spallation backgrounds, as well as contribution from detector backgrounds and reactor neutrinos.

The dedicated muon fitter Muboy [65], developed by the SK collaboration, is used to reconstruct the tracks of incoming cosmic muons. In addition to identifying muon events, the fitter separates the muons into topological categories. The most common muon events are single through-going muons, traversing the tank from one end to the other, comprising 83.8% of the total in our dataset. Other types of muon events include multiple-track events (7.2%), muons stopping inside the tank (4.7%), corner-clipping muons (4.7%), and misfit muons for which no category could be determined (0.3%). Corner-clipping muons are simply through-going muons with a short track near the detector corner, which have a limited potential for energy deposition, and cause virtually no spallation. These categories can inform our spallation reduction strategy down the line. More details about the Muboy algorithm can be found in Ref [65]. Poorly-fitted muons, corresponding to a Muboy goodness < 0.4, are also fitted with a more resource-intensive brute-force fitter. Due to the large timescales considered, the muon-fitting process is a particularly time-consuming step in the analysis, taking up to a month of processing time to go through the SK-IV dataset.

Preliminary cuts

After applying reconstruction software to the prompt events and the associated muon events, we can proceed with the preliminary reduction step. First, a series of cuts is applied to the data sample to remove backgrounds only found in data. Events not suited for physics analysis are removed, including calibration runs, test runs, pedestal events, incomplete events, and events with data quality problems. Any event accompanied by an OD trigger, indicating a cosmic-muon-induced event, is removed, as well as events with a total charge deposition greater than 1000 photoelectrons. Events triggered by the electrons from the decay of visible muons are removed by requiring that the trigger time be > 50 μs after an observed muon event. Sometimes, events can be triggered by random noise fluctuations. To remove these events from the data, we require that the number of PMT hits with charge deposition < 0.5 photoelectrons amount to less than 55% of the total: $N_{Q<0.5}/N_{tot} < 0.55$. These cuts are found to have negligible impact on the simulated samples.

Cut	Value	Apply to simulation?
Physics event selection	/	
OD trigger	veto	
Total photoelectrons	< 1000	
Time after muon	> 50 μs	
$N_{Q<0.5}/N_{tot}$	< 0.55	
d_{wall}	> 2 m	✓
$(\text{goodness})^2 - (\text{dirKS})^2$	> 0.25	✓
E_{rec}	$[10(8), 100]$ MeV	✓

TABLE 3.2: First set of cuts applied to the data and simulated samples. Some cuts only address backgrounds that are not present in simulation, and have negligible impact on the signal, so they are only applied to the data sample.

The rest of the cuts at this stage are applied to both data and simulated samples. Only events inside SK's fiducial volume are kept, meaning that the distance from the detector wall, d_{wall}, of the reconstructed prompt vertex has to be greater than 2 m. The quality of the BONSAI fit is also evaluated. The BONSAI fit goodness of the vertex $\vec{x}$ is defined by a weighted average of a Gaussian function of the time-of-flight-subtracted time of each PMT hit:

$$\text{BONSAI goodness}(\vec{x}) = \frac{\sum_i w_i e^{-\frac{1}{2}\frac{(t_i - \text{TOF}_i)}{3\text{ ns}}}}{\sum_i w_i} \qquad w_i = e^{-\frac{1}{2}\frac{(t_i - \text{TOF}_i)}{60\text{ ns}}}, \tag{3.26}$$

where PMT hit i is associated with time t_i from the trigger time, time of flight TOF_i with respect of vertex $\vec{x}$, and the weight w_i. In other words, a well-fitted vertex should have hits with time residuals within ~ 3 ns (the expected time resolution for Cherenkov events), not taking into account hits that are more than ~ 60 ns and are unlikely to be part of the real vertex. Finally, the uniformity of the hits along the azimuthal angle ϕ around the fitted event direction is used to determine fit quality, since a Cherenkov event produces light uniformly with respect to ϕ. We compute the Kolmogorov-Smirnov statistic d_{KS}, measuring the maximum distance between the cumulative ϕ distribution of the event and of a uniform ϕ distribution. The *dirKS* variable is then defined as d_{KS}/N_{hit}. This variable is used in conjunction with the goodness, requiring $(\text{BONSAI goodness})^2 - (\text{dirKS})^2 > 0.25$. Figure 3.10 plots *dirKS* for both the simulated signal sample and the data sample. At this point we also eliminate events with reconstructed energy $E_{\text{rec}} > 100$ MeV, and with $E_{\text{rec}} < 10$ MeV (for events with an SHE threshold of 70 hits), or $E_{\text{rec}} < 8$ MeV (for events with an SHE threshold of 58 hits). The cuts until this point are summarized in Table 3.2. After this preliminary reduction, we still have roughly 1.4 M events in the data sample.

3.4 Positron identification

The third set of reduction cuts is aimed at the characterization of the prompt positron. The cuts in this section target the removal of atmospheric and solar neutrino backgrounds, as well as any remaining events caused by radioactivity, and only consider the prompt event.

Travel distance

Despite selection of the fiducial volume in section 3.3, some radioactive decays in the walls and in the rock surrounding the detector still remain. Further reduction of the fiducial volume would significantly eat away at our signal efficiency. Rather, we define the *effective* distance from the wall d_{eff}, calculated as the distance from reconstructed vertex to the wall *along the reconstructed track direction*, as sketched in Figure 3.11. This observable represents the distance that a particle produced in a radioactive decay would have to travel to induce an event at the observed vertex.

As discussed in Section 3.2.1, no simulation is available for the low-energy radioactivity inherent to the detector setup. To estimate the prevalence of such backgrounds, we use the data sample and compare its d_{eff} distribution to that in the IBD simulation, as shown in Figure 3.12. The latter represents a set of homogenous and isotropic events, and can therefore be used as the reference d_{eff} distribution. To compare the data and simulation distributions, we normalize them to have the same area above 10 m, where no radioactivity-induced event is thought to survive. Radioactivity can then be observed as an excess at low d_{eff} values. Here, the cut is energy-dependent, as this background is much more prevalent at lower energies, defining the threshold as follows:

$$d_{\text{eff}} > \begin{cases} 500 - 50(E_{\text{rec}} - 16\text{MeV}) \text{ cm} & E_{\text{rec}} < 20\text{MeV} \\ 300 \text{ cm} & E_{\text{rec}} \geq 20\text{MeV}. \end{cases} \tag{3.27}$$

Pre- and post-activity cuts

These cuts are aimed at the removal of multiple Cherenkov cones produced in the tank by a single atmospheric neutrino. As described in Section 3.2.5, atmospheric neutrinos can produce a delayed event through the electron from a muon or pion decay. These events will display a clear double-peak structure in their charge deposition, unlike the IBD positron, which is only accompanied by a much fainter neutron peak. Either of the double peaks can trigger an event, so we can look for such backgrounds by analyzing the PMT activity both before and after the trigger time.

To look for activity before the trigger time, we analyze the available 5 μs that are saved preceding each event in SK. We calculate the time-of-flight required for each hit to correspond to the event vertex, and search for PMT hits clustered in TOF-subtracted time. We compute the maximum number of hits in a 15-ns window,

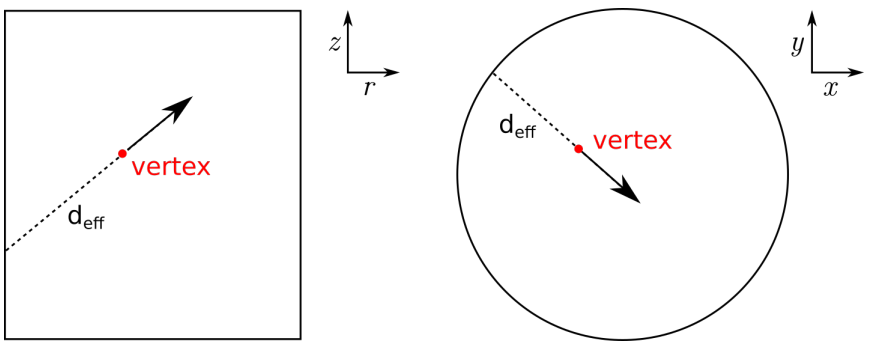

FIGURE 3.11: Schematic of the d_{eff} observable computation. The red dot and solid arrow represent the reconstructed vertex and track direction, respectively. Left: side view of the detector. Right: top-down view of the detector.

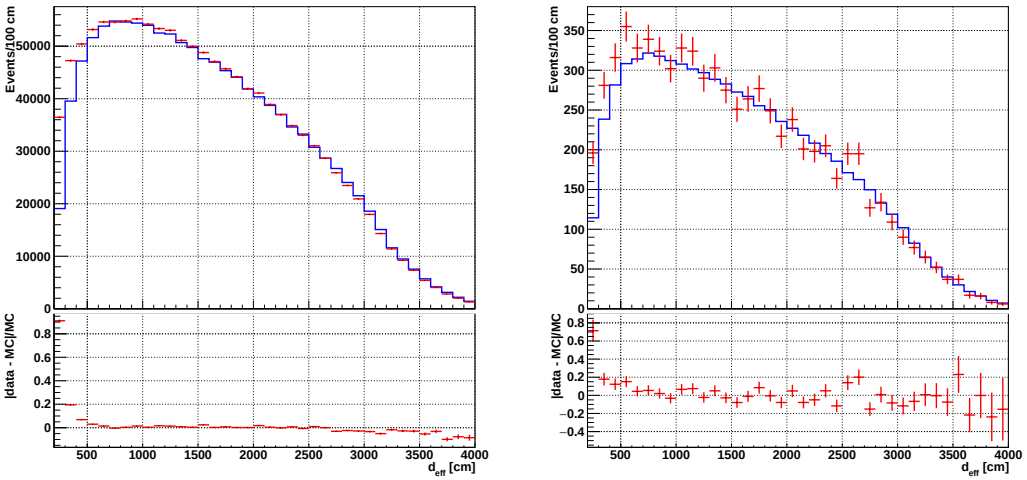

FIGURE 3.12: Plot of the d_{eff} distribution for data (red) and DSNB simulation (blue), along with the relative difference between the two in the lower panel. Left: data with $E_{\text{rec}} \in [8, 30]$ MeV. Right: data with $E_{\text{rec}} \in [16, 80]$ MeV.

	Low sideband	Signal region	High sideband
θ_{C}	[20°, 38°]	[38°, 50°]	[78°, 90°]
Dominated by	μ/π	e^+/e^-	NC events

TABLE 3.3: Signal and sideband Cherenkov angle regions.

$N_{\text{pre}}^{\text{max}}$, found at least 12 ns before the trigger time, and require $N_{\text{pre}}^{\text{max}} < 12$. A similar strategy is employed to search for post-event activity, using an algorithm developed for previous SK analyses that reconstructs the electrons from muon or pion decay [66]. The number of reconstructed electrons, $N_{\text{decay-e}}$ is required to be 0.

Cherenkov angle

A cornerstone observable of the analysis is the opening angle θ_{C} of the Cherenkov cone produced by the prompt event. For highly-relativistic particles such as the IBD positron, θ_{C} is around $42°$ in SK's pure water. Heavier particles, in particular muons and pions, that are observed near their Cherenkov-threshold energy, will instead be associated with smaller opening angles. Neutral-current interaction events from atmospheric neutrinos, on the other hand, tend to be associated with signals from multiple photons. These will result in multiple Cherenkov cones, and can be misreconstructed as a single Cherenkov cone with an especially large opening angle. The θ_{C} angle is therefore a very powerful discriminator.

To estimate the opening angle from the PMT hits in the prompt event, we first calculate the opening angle of the cones defined by all possible three-hit combinations. For each event, we scan the distribution of these Cherenkov angles with a $2.7°$ window, and define θ_C as the center of the window maximizing the number of hits. In addition to a signal region, defined between $38°$ and $50°$, we define two sidebands as in Table 3.3, which will be used in the final statistical analysis.

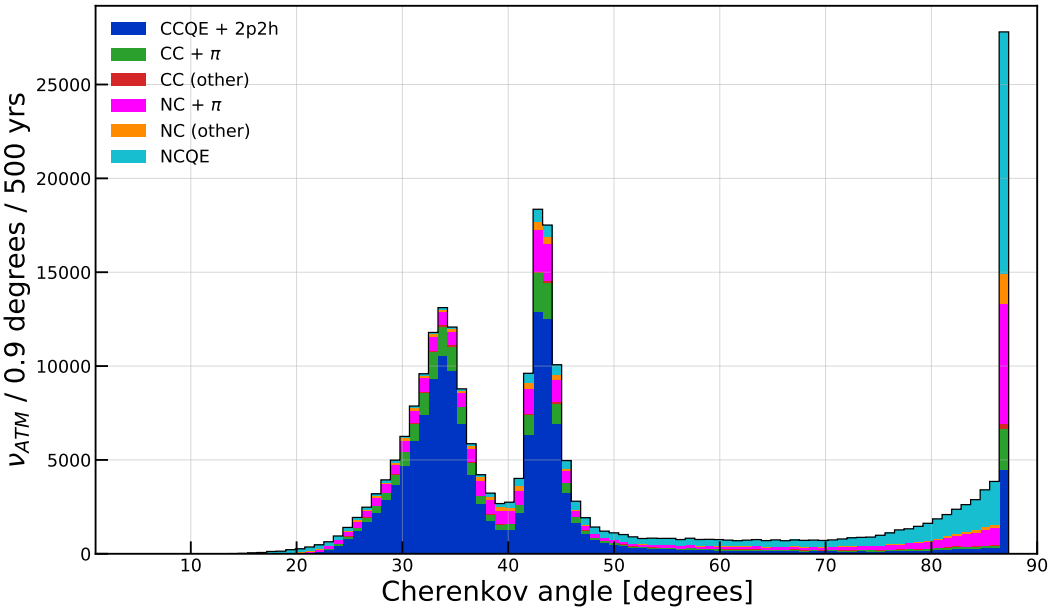

FIGURE 3.13: Cherenkov opening angle distribution for the atmospheric neutrino simulation sample. Note that the cutoff of the distribution at the high end is an artifact of our computation method. The last bin corresponds to the center of the last angle window sampled.

Ring clearness

In addition to estimating the Cherenkov angle itself, we compute a measure of the ring's clearness, i.e. how closely the Cherenkov ring resembles a perfect circle. This is motivated by the fact that due to their relatively low momentum, electrons tend to scatter in the detector's water as they emit their Cherenkov radiation, resulting in a misshapen ring. In contrast, higher-momentum particles will hardly deviate from their path and produce a clearer Cherenkov ring. This principle is especially useful to discriminate visible pions, and is referred in this text as the pion likeness R_{PION}. This observable is calculated starting from the same distribution of three-hit opening angles used to compute θ_C, by comparing the number of triplets in two different angular windows around θ_C:

$$R_{\text{PION}} = \frac{N_{\text{triplets}}(\theta_C \pm 3°)}{N_{\text{triplets}}(\theta_C \pm 10°)} \tag{3.28}$$

The distributions of R_{PION} for atmospheric backgrounds and the DSNB signal are shown in Figure 3.14. A cut of $R_{\text{PION}} < 0.36$ was chosen.

Average charge deposition

Similarly, since higher-momentum particles suffer fewer deviations in their path, they also tend to deposit more charge in the same individual PMTs. In comparison, electrons and positrons at our energy range will cause hit patterns that are more dispersed, with lower charge deposition per PMT, on average. Since we are interested in the average charge deposition caused by the Cherenkov radiation of the prompt event, using all hits in the event in the calculation would dilute the observable by including e.g. hits caused by dark noise. Therefore, we take the 50-ns time-TOF window maximizing the number of PMT hits, and calculate the average charge

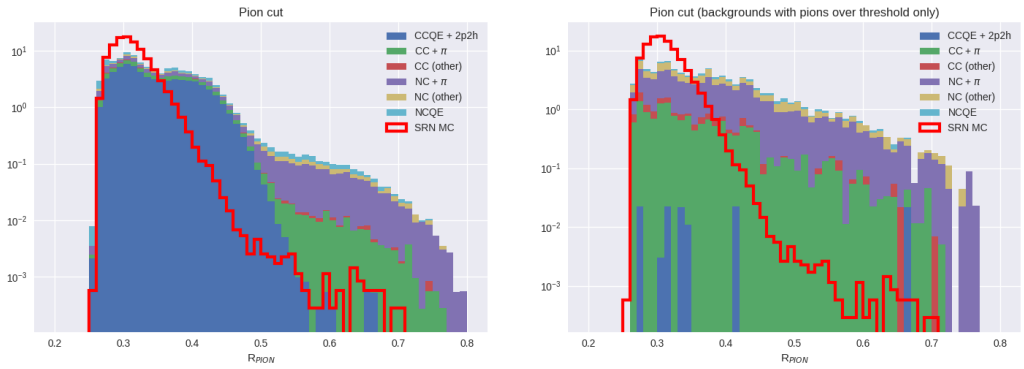

FIGURE 3.14: Distribution of R_{PION} for atmospheric neutrino (filled histogram) and DSNB simulation (empty histogram). Left: distribution for the entire atmospheric neutrino simulation. Right: distribution for atmospheric neutrino event containing visible pions.

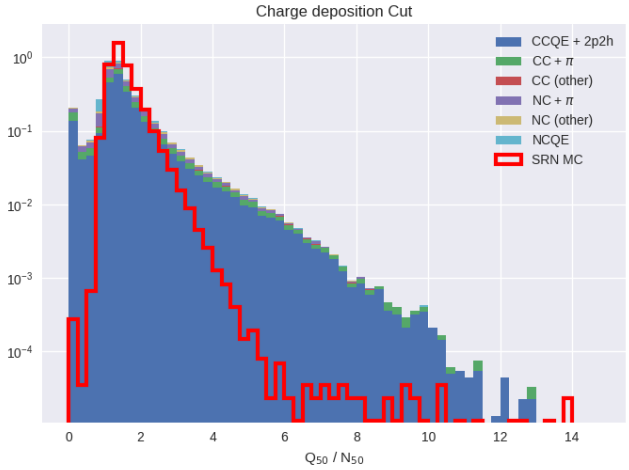

FIGURE 3.15: Distribution of q_{50}/n_{50} for atmospheric neutrino (filled histogram) and DSNB simulation (empty histogram).

deposition in this window, q_{50}/n_{50}. The distribution of the observable is plotted in Figure 3.15. For this analysis, a cut $q_{50}/n_{50} < 2$ is chosen.

Solar cut

Solar neutrinos, as mentioned in 3.2.2, are mainly detectable at SK via NC interaction of the ν_e produced in the ^{8}B production branch. Since no neutron is produced in these interactions, the neutron tagging algorithm (detailed in the following section) easily removes any solar neutrino background. However, in this analysis we also want to take into account events that fail the neutron tagging cut, in order to maximize our effective exposure. Therefore, a solar neutrino cut will be applied to those events for which no neutron is found.

This analysis uses the same cuts used by the SK-IV solar neutrino analysis [67]. In particular, the angle θ_{sun}, between the reconstructed event track direction and the current direction of the Sun, is the key observable here. Solar neutrino events will be clustered around $\cos\theta_{\mathrm{sun}} = 0$, with kinematic and reconstruction effects broadening the peak of the distribution. Since well-reconstructed events cluster more tightly around $\cos\theta_{\mathrm{sun}} = 0$, different cuts are chosen based on the reconstruction's

Cut	Value	Signal eff.	Background eff.
Preliminary reduction	Various	$> 99(1)\%$	—
θ_C	$\in [38°, 50°]$	$93.9\ (0.7)\%$	35.0%
R_{pion}	< 0.36	$97.3\ (0.2)\%$	62.6%
q_{50}/n_{50}	< 2	$98.9\ (1.2)\%$	75.5%
$N_{\text{pre}}^{\text{max}}$	< 12	$> 99\%$	63.6%
$N_{\text{decay-e}}$	< 1	$> 99\%$	42.5%
d_{eff}	Energy-dependent	$98.4(0.1)\%$	93.3%
θ_{sun}	MSG-dependent	$92(1)\%$	$< 0.1\%$

TABLE 3.4: Preliminary reduction and positron characterization cuts, along with signal and background efficiencies. The background efficiency of the solar neutrino cut refers to solar neutrino backgrounds only.

goodness-of-fit. The latter is parametrized by the Multiple Scattering Goodness (MSG), an observable specifically developed for the solar neutrino analysis. For more information about this cut, refer to Refs. [49, 67].

Cut optimization

Where simulation of the background is available, we use the simulation samples to maximize the $S/\sqrt{S+B}$ figure of merit. As mentioned, due to the absence of simulated low-energy radioactive background, the d_{eff} distribution is taken from data, and is compared to a homogenous and isotropic signal. The estimation of the systematic uncertainty associated with these cuts—not part of this work—is detailed in Reference [49]. The reduction cuts discussed until this point, along with their efficiencies and systematic uncertainties, are summarized in Table 3.4.

3.5 Neutron tagging

The identification of the IBD neutron marks an important milestone for Super-Kamiokande. With the SK-IV data-taking period, the detector achieved the capability to record the neutron captures on Hydrogen for the first time. As will become clear, neutron tagging is a delicate process. In this section, this final set of cuts, teasing out the coincidence neutron signal, is detailed.

3.5.1 Candidate preselection

The standard reconstruction strategies used at SK, even for low-energy events, are not adequate to detect the very faint signal produced by the 2.2 MeV neutron capture photon. A typical neutron capture on Hydrogen will produce < 10 PMT hits in total, which could easily be lost among the sea of low-energy detector backgrounds, representing a whopping $\mathcal{O}(10^5)$ hits in each event. The neutron capture itself occurs at a timescale of $\tau \sim 200\mu$s after the prompt event, and during this time the neutron will be able to travel a considerable distance (~ 50 cm). After capture, the photon emission and the visible Compton scattering occur in rapid succession—instantaneously, as far as our time resolution is concerned. This means that we can in principle resolve two interaction vertices: the delayed vertex associated with the visible neutron signal—practically the same as the vertex associated with the neutron capture—and the prompt positron emission vertex. Attempting to directly fit the

neutron capture vertex, using the BONSAI low-energy fitter will be associated with large uncertainties and potentially fail to fit the event altogether, due to the scarce number of PMT hits available.

Before searching for the neutron itself we apply two cuts to obtain a cleaner event. First, a cut on the maximum number of hits in a 200-ns window $N_{200}^{\max}$ is placed at $N_{200}^{\max} < 50$. This cut removes any PMT activity that is uncharacteristically high for a neutron capture. This cut has a virtually perfect signal efficiency, while removing residual events that contain visible muons in the neutron search window. Secondly, we remove PMT hits that are consistent with *continuous dark noise*. This is a kind of dark noise likely associated with scintillation in the glass cover of a PMT, wherein multiple noise hits are observed to originate within the same PMT on timescale of about 10 μs. Therefore, pairs of hits on the same PMT within 6 μs of each other are removed.

The first key that will allow for detection of the neutron is the fact that the typical spatial separation between the prompt and delayed vertices is actually close to the spatial resolution of the low-energy vertex fit that we obtained for the prompt vertex. This means that, as a first approximation, we can set the locations of two vertices equal to each other. This first approximation can then be directly used to eliminate a very large fraction of the low-energy accidental background that would be largely incompatible with this vertex. More specifically, the vertex estimate can be used to calculate the corresponding time of flight (*TOF*) to each PMT hit in the event. The distribution of time-TOF for the PMT hits associated with the true neutron capture will then display a peak at the time of neutron capture. Further refinement of the neutron vertex is then possible by testing a limited subset of the detector volume around the prompt vertex, attempting to maximize the time-TOF peak. The procedure is sketched in Fig. 3.16.

To look for this timing peak, we define a window of 10 ns that we slide across the time-TOF of the event hits and count the number of PMT hits it contains, N_{10}, plotted in Fig. 3.17. A neutron capture candidate is then defined as a cluster of hits associated with a local N_{10} maximum. A *preselection* cut,

$$N_{10} > 5 \tag{3.29}$$

is applied to the candidates. Without such a cut, the amount of accidental coincidences would simply overwhelm the neutron signal and significantly slow down processing

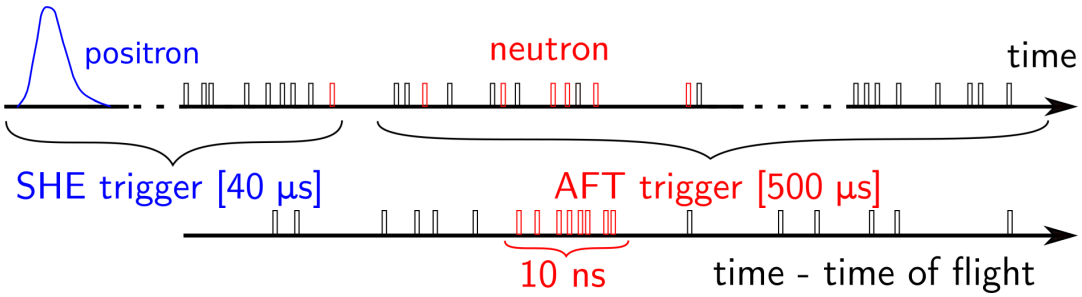

FIGURE 3.16: Sketch of the selection of neutron signal candidates following a prompt signal from a possible positron event. In blue, the prompt signal hits as a function of time. In red, the true delayed signal hits. In black, background hits, potentially causing an accidental coincidence with the prompt signal. On top, the hits are ordered with respect to the true PMT hit time. On the bottom, they are ordered with the time-TOF, approximating the delayed vertex as the prompt vertex.

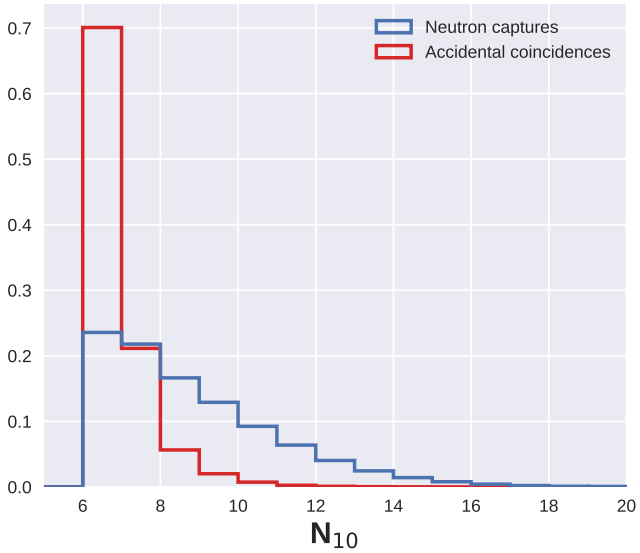

FIGURE 3.17: Probability density function of N_{10}, for accidental coincidences (in red) and true neutrons captures (in blue). The amount of background contamination rises sharply with decreasing N_{10}.

time. The impact of this cut is significant, removing a hefty fraction of the accidental coincidences, and leaving us with an average of 11 accidental coincidences per prompt signal. On the flip side, it also eats away at more than half of the delayed signal, leaving us with just 45% of the true neutrons after preselection. In principle, the looser preselection cut, the more signal we can hope to extract down the line. In practice, neutron signatures with low values of N_{10} represent extremely faint signals that are hardly distinguishable from noise, at least with the method described in this text. These low-hit signals are therefore—as far as our cut procedure is concerned—essentially lost from the start, and the associated inefficiency can be attributed to the intrinsic weakness of the neutron signal in pure, undoped water. This hypothesis was tested by varying the preselection cut and applying the following neutron tagging cut (described below). The neutron tagging procedure achieved maximum efficiency for $N_{10} > 5$, with no further efficiency gains (or further background reduction) at lower N_{10} values, for which much longer computing time is required due to the exponentially greater number of candidates for each reduction of the N_{10} threshold.

3.5.2 Observable calculation

This first reconstruction step is hardly enough for a complete algorithm, as even after this drastic reduction in candidates, accidental coincidences outnumber the true signals about 24-to-1, for the average IBD event. Further characterization of the event is therefore needed. Despite the low number of PMT hits, a surprising amount of information can be inferred that can help us sift out the noise from our neutron signals. In total, 22 discriminating observables are defined, attempting to draw out information about the candidate hits' spatial topology, their time distribution, and charge deposition.

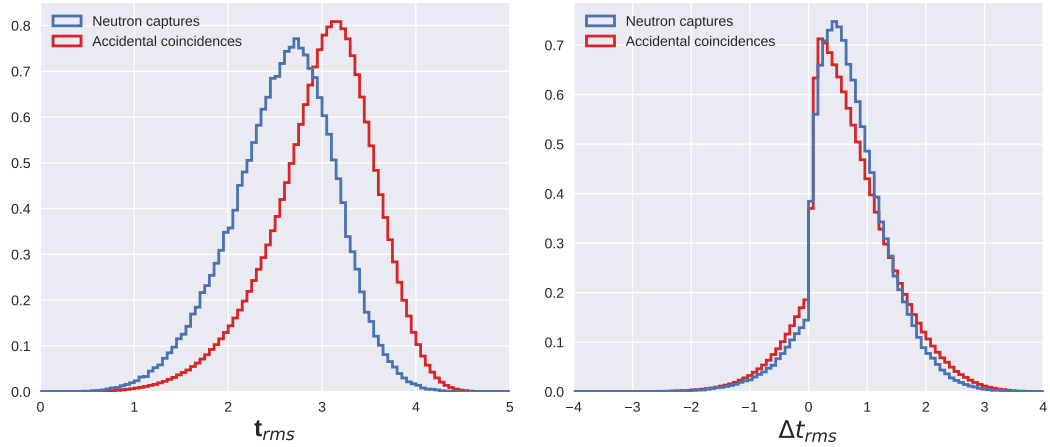

FIGURE 3.18: Probability density function of t_{RMS} (left) and Δt_{RMS} (right), in nanoseconds, for accidental coincidences (red) and true neutrons captures (blue). PMT hits from true neutron captures display lower time dispersion, and see this dispersion reduced to a greater extent when a better vertex is found. Note that, due to its definition, Δt_{RMS}, corresponds to the *decrease* in t_{RMS} after the vertex improvement.

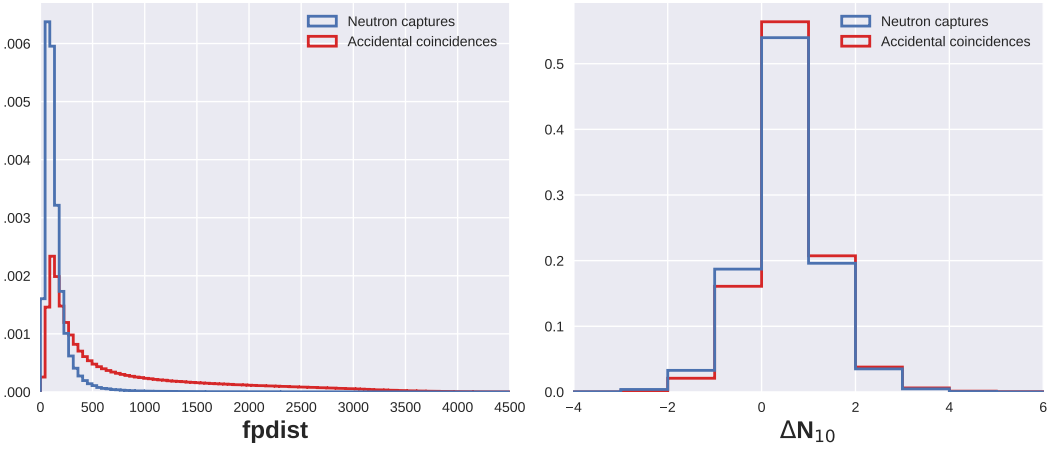

FIGURE 3.19: Probability density function of $\mathrm{fp}_{\mathrm{dist}}$, in cm (left) and ΔN_{10} (right), for accidental coincidences (red) and true neutrons captures (blue). As expected, true neutron candidates are reconstructed close to the prompt vertex, while accidental coincidences can be reconstructed at larger distances. Here, the neutron vertex reconstruction is extended to the entire detector to take advantage of the long tail of the background distribution.

Neutron capture vertex determination

As mentioned previously, the neutron captured vertex is considered the same as the prompt vertex as a first guess. To obtain a more accurate estimate, we first calculate the root mean squared deviation of the neutron candidate hit cluster:

$$t_{\text{RMS}} = \sqrt{\frac{\sum_{i=1}^{N_{10}} (\bar{t} - t_i)^2}{N_{10}}}, \tag{3.30}$$

with t_i being the t-TOF of the ith PMT hit, and $\bar{t}$ the average t-TOF for this neutron candidate. This is then a measure of the time dispersion of the hits, and a proxy for the quality of the neutron vertex estimation. A perfect vertex reconstruction would yield a t_{RMS} near 0, so we look for a new vertex around our initial guess, $\mathbf{x}_{\text{prompt}}$, with the goal of minimizing t_{RMS}. We find a new best vertex position, $\mathbf{x}_{\text{fit}}$, by scanning over a 2-meter-radius sphere around the initial guess, and find the new set of hits maximizing N_{10} for this vertex. We calculate the new time dispersion t'_{RMS}, and the new number of hits, N'_{10}, and define the difference from their original values as observables:

$$\text{fp}_{\text{dist}} = |\mathbf{x}_{\text{fit}} - \mathbf{x}_{\text{prompt}}| \tag{3.31}$$

$$\Delta t_{\text{RMS}} = t_{\text{RMS}} - t'_{\text{RMS}} \tag{3.32}$$

$$\Delta N_{10} = N'_{10} - N_{10}. \tag{3.33}$$

In addition, we can use the results from fitting the delayed vertex with BONSAI, the standard low-energy fitter—with the caveat that they incur large uncertainties. The fit is performed taking into account PMT hits within 1.3 μs around the prompt vertex. In addition to the new vertex $\mathbf{x}_{\text{BONSAI}}$, the algorithm provides an estimate of the energy of fitted event. Rather than work directly with this energy estimate, which is harder to interpret at low energies, we use the *effective* number of hits N_{eff}, an estimate of the number of photons produced at the vertex, as detailed in Section 3.3. Since this variable is correlated with the BonSai-reconstructed Energy, we also refer to it as `bse`.

The discrepancy between the BONSAI-fitted neutron vertex and the previous fit defines a new observable:

$$\text{bp}_{\text{dist}} = |\mathbf{x}_{\text{BONSAI}} - \mathbf{x}_{\text{fit}}|. \tag{3.34}$$

This value can be interpreted as a measure of our confidence with the vertex fit, and larger values of bp_{dist} indicate a higher likelihood that the neutron candidate is an accidental coincidence.

Some reconstructed neutron vertex positions are inherently more likely to originate from background due to their proximity to the detector walls. While radioactivity from the materials in and around SK is largely removed by the fiducial volume cut, some will remain, potentially causing an accidental coincidence near the wall. For this reason, a relevant observable is the distance of the reconstructed neutron vertex from the wall,

$$\text{f}_{\text{wall}} = |\mathbf{x}_{\text{fit}} - \mathbf{x}_{\text{wall}}| \tag{3.35}$$

$$\text{b}_{\text{wall}} = |\mathbf{x}_{\text{BONSAI}} - \mathbf{x}_{\text{wall}}|. \tag{3.36}$$

For the calculation of the observables that follow, the new vertex $\mathbf{x}_{\text{fit}}$ will be used, along with the new set of PMT hits.

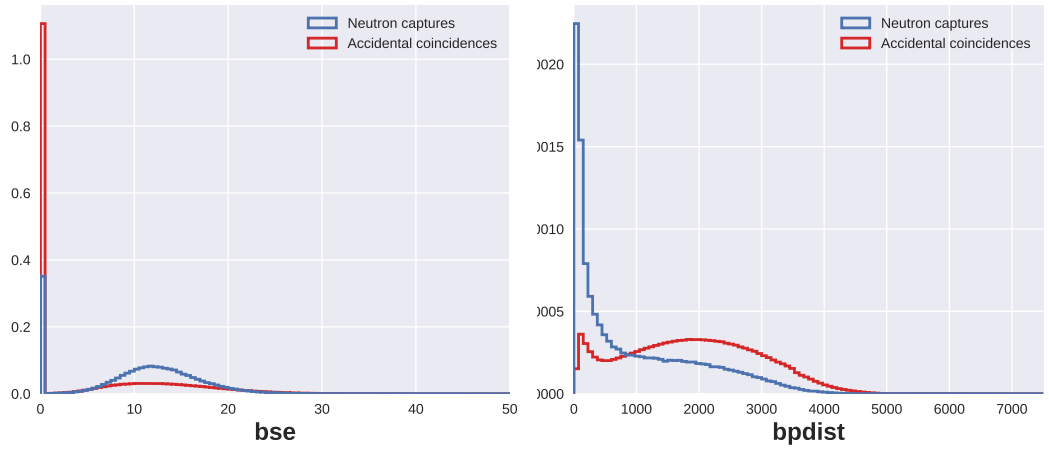

FIGURE 3.20: Probability density function of `bse` (left) and bp_{dist}, in cm (right), for accidental coincidences (red) and true neutrons captures (blue). Values of 0 for `bse` indicate that the fit could not converge and are dominated by background. For vertex fits that converged successfully, background candidates display a flatter `bse` distribution. Here, bp_{dist} is shown for candidates with non-zero `bse`. As expected, true neutrons display much greater consistency between the two vertex fits. In contrast, for accidental coincidences the two fits often differ by several tens of meters.

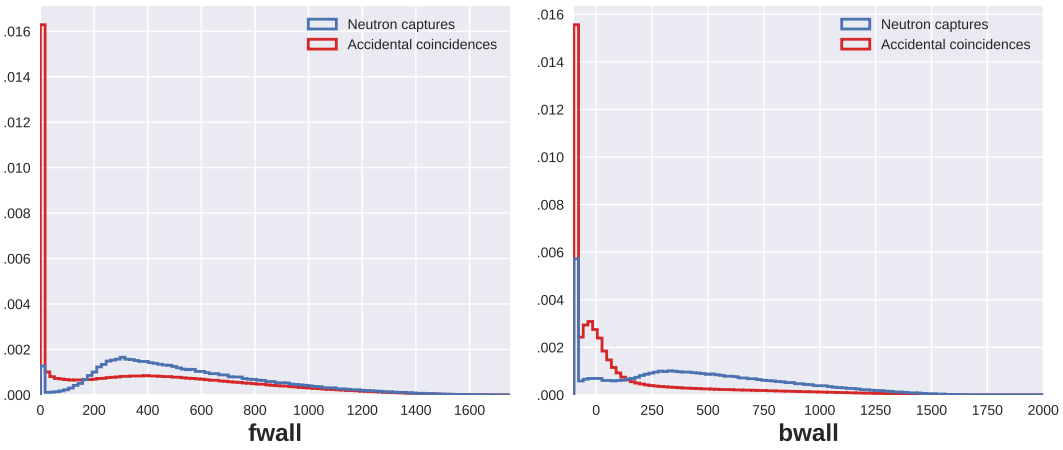

FIGURE 3.21: Probability density function of f_{wall} (left) and b_{wall} (right), both in cm, for accidental coincidences (red) and true neutrons captures (blue). A large number of background candidates are reconstructed with their vertex near the detector walls. The BONSAI fitter can reconstruct the vertex outside the detector itself, unlike the t_{RMS}-based fit, leading to different behavior at the lower end of the distribution.

Cherenkov event topology

After improving the fit on the neutron capture vertex, we further characterize the neutron candidate based on our expectations of its geometry. Our meager portion of PMT hits will of course not allow us to appreciate a Cherenkov cone *per se*. Nevertheless, we can look for the general features of the Cherenkov pattern that we expect to see, on average, in the neutron signal.

First, we can directly attempt to measure the opening angle of the cone. Since the Cherenkov light is emitted by a Compton electron, the opening angle is in principle the same as the one observed for the positron, i.e. about $42°$. To estimate the opening angle, we first define the average hit direction, $\mathbf{v}_{mean}$, and define the opening angle θ_{mean} as the mean angular difference of the PMT hits from $\mathbf{v}_{mean}$:

$$\mathbf{v}_{mean} = \frac{\sum_{i=1}^{N_{10}} (\mathbf{v}_i - \mathbf{x}_{\text{fit}})}{|\sum_{i=1}^{N_{10}} (\mathbf{v}_i - \mathbf{x}_{\text{fit}})|} \tag{3.37}$$

$$\theta_{mean} = \sum_{i=1}^{N_{10}} \arccos(\mathbf{v}_i \cdot \mathbf{v}_{mean}), \tag{3.38}$$

where $\mathbf{v}_i$ is the position vector of the ith PMT hit. While we will not be able to resolve the angle to the same extent as the prompt event, a large subset of the accidental coincidences will display very different reconstructed opening angles. To further characterize θ, we calculate the root mean squared deviation of the opening angle:

$$\theta_{\text{RMS}} = \sqrt{\frac{\sum_{i=1}^{N_{10}} (\theta_{mean} - \theta_i)^2}{N_{10}}}. \tag{3.39}$$

This can be thought of as measure of the *fuzziness* of the Cherenkov ring. For a true neutron signal is determined by the inherent ring deformations due to the possible scattering of the Compton electron, as well as the resolution of the neutron capture vertex. The angular characterization of the candidate is completed by the calculation of its azimuthal deviation, ϕ_{RMS}. The azimuthal angle ϕ_i of PMT hit i is defined as taking the mean direction $\mathbf{v}_{mean}$ as the axis of rotation. We then calculate how uniformly the PMT hits are distributed around this axis:

$$\Delta\phi_{mean} = 2\pi / N_{10} \tag{3.40}$$

$$\phi_{\text{RMS}} = \sqrt{\frac{\sum_{i=2}^{N_{10}} (\Delta\phi_{mean} - (\phi_i - \phi_{i-1}))^2}{N_{10}}}. \tag{3.41}$$

We also keep track of how many PMT hits are found in unlikely positions, given our estimation of the neutron vertex. With $N_{\text{low}\theta}$ we count the number of hits within a $20°$ opening angle around the mean PMT direction. These PMT hits are unusually forward compared to the expected Cherenkov opening angle. Similarly, N_{back} counts the numbers of hits in the backward direction, with respect to the mean hit direction. Finally, we assign a weight to each PMT hit, based on the probability of the PMT's position PMT in relation to the vertex:

$$\text{wgt}_i = f(\theta_i^{\text{inc}}) \frac{e^{-r_i / L_{\text{att}}}}{(r/\text{cm})_i^2}. \tag{3.42}$$

Here, PMTs that are further away from the vertex are penalized, accounting for the

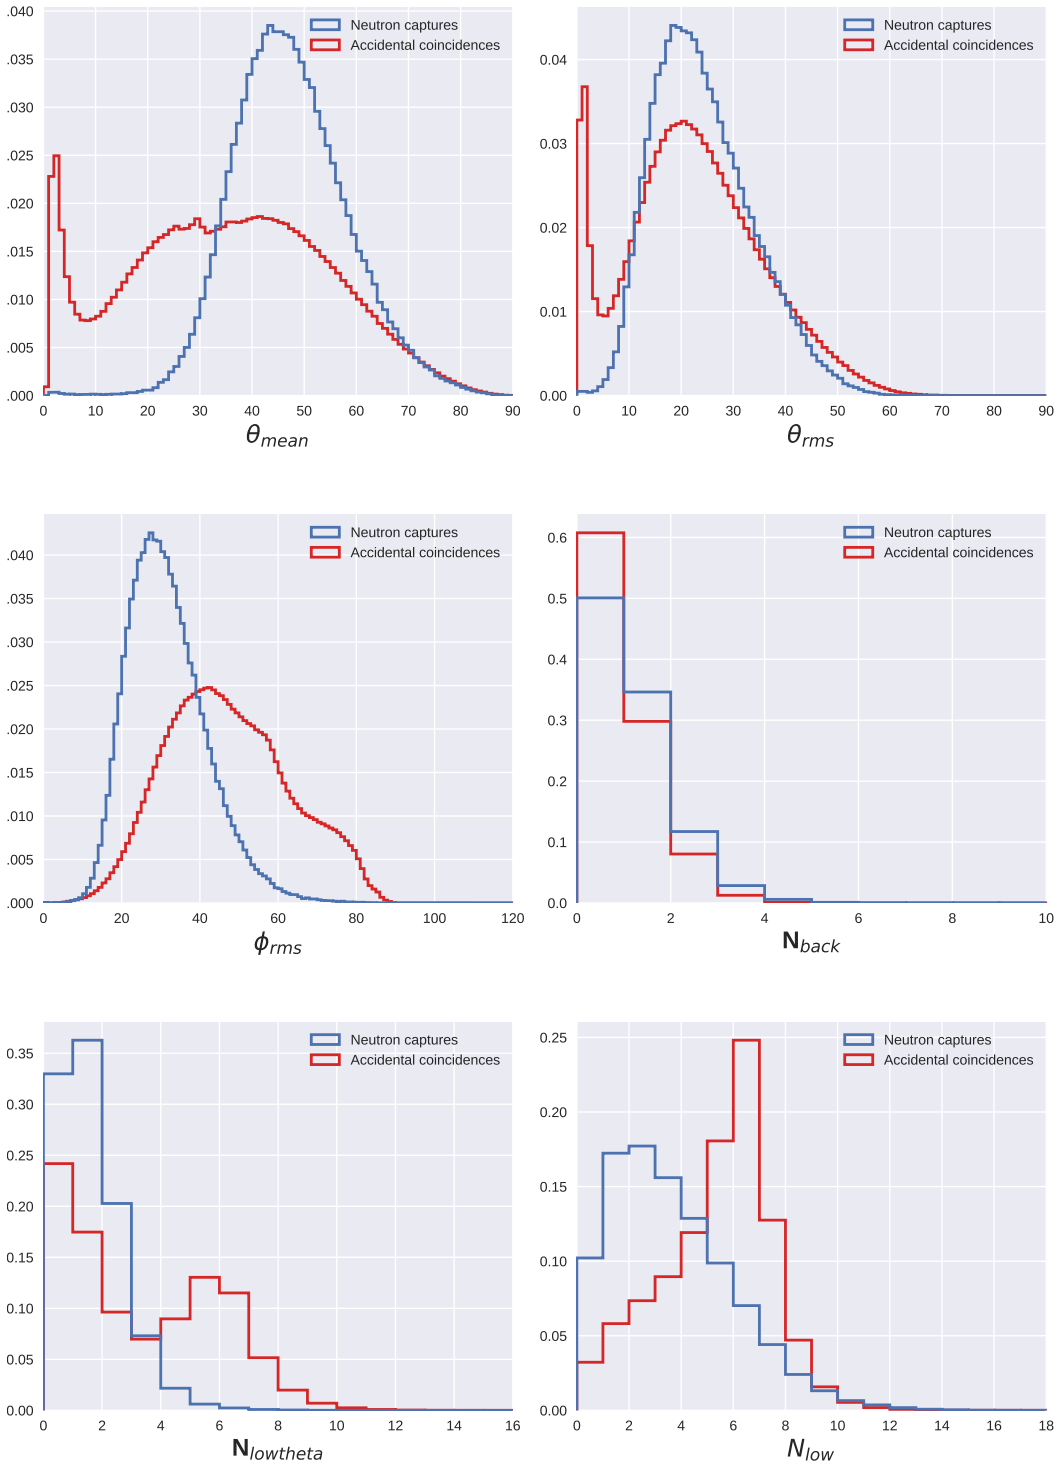

FIGURE 3.22: Probability density function of θ_{mean} (top left), θ_{RMS} (top right), and ϕ_{RMS} (middle left), in degrees, as well as N_{back} (middle right), $N_{\text{low}\theta}$ (bottom left), and N_{low} (bottom right), for accidental coincidences (red) and true neutrons captures (blue). While our resolution does not allow us to precisely determine the Cherenkov angle of the neutron candidate, we can still clearly see the θ_{mean} peak at about $42°$. Highly-localized backgrounds will be reconstructed with θ_{mean} and θ_{RMS} near 0. As expected, accidental coincidences contain more hits that in the backwards direction, at low reconstructed opening angles, or in otherwise low-probability positions.

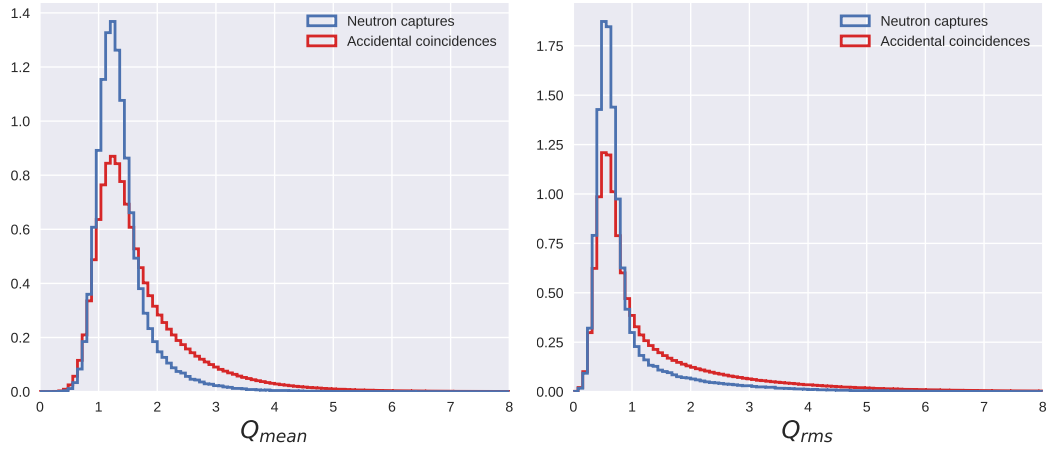

FIGURE 3.23: Probability density function of Q_{mean} (left) and Q_{RMS} (right), both in photoelectrons/PMT, for accidental coincidences (red) and true neutrons captures (blue). Since we are affected by low-energy backgrounds, most of the accidental coincidences will display a low charge deposition. However, unusually-high deposition—visible in the tail of the Q_{mean} distribution—are more likely for background hits.

inverse-square decay of the signal and the attenuation in the water. Further, the angle of incidence of the PMT hit with respect to the wall, θ^{inc}, enters the probability calculation through the scaling function f. In particular, photons striking the PMT perpendicularly ($\theta_{\text{inc}} = 0$) are detected with greater efficiency that photons at high θ_{inc}, so PMT hits at high incidence angles are further penalized. N_{low} is then defined as the number of PMT hits below a chosen weight threshold.

Noise characterization

The last subset of observables is defined to take advantage of specific aspects of the low-energy backgrounds that cause accidental coincidences. First, PMT hits that record a high charge deposition are a hint that the candidate is comes from background. True neutron captures, due to their low energy will, by and large, cause low PMT activation ($\sim$1 photoelectron/PMT). Therefore, the mean charge deposition per PMT Q_{mean}, and its RMS deviation Q_{RMS}, are used as observables. We also define N_{highQ}, as the number of PMT hits associated with three or more photoelectrons.

Certain categories of backgrounds can also be tightly clustered, activating multiple nearby PMTs. These are typically caused by radioactivity near the walls, or flasher events inherent to the PMTs themselves. To look for such backgrounds, we define two hits as belonging to the same hit *cluster* if the two hit vectors (calculated from the neutron vertex) are within $14.1°$ of each other. Due to the geometry of the detector, this is the maximum angle between two hit vectors associated with two neighboring PMT hits, assuming the vectors originate inside the fiducial volume. We also only consider clusters containing three or more hits. We can then define N_{clus}, the total number of PMT hits that belong to a cluster.

It is also possible that the background causing the accidental coincidence causes prolonged PMT activity, exceeding the $\sim$10-ns time extent of the neutron signal. In this case we can check the number of hits in a larger time window around our candidate. Namely, the observable N_{300} counts the number of hits in 300-ns window centered around the original 10-ns widow.

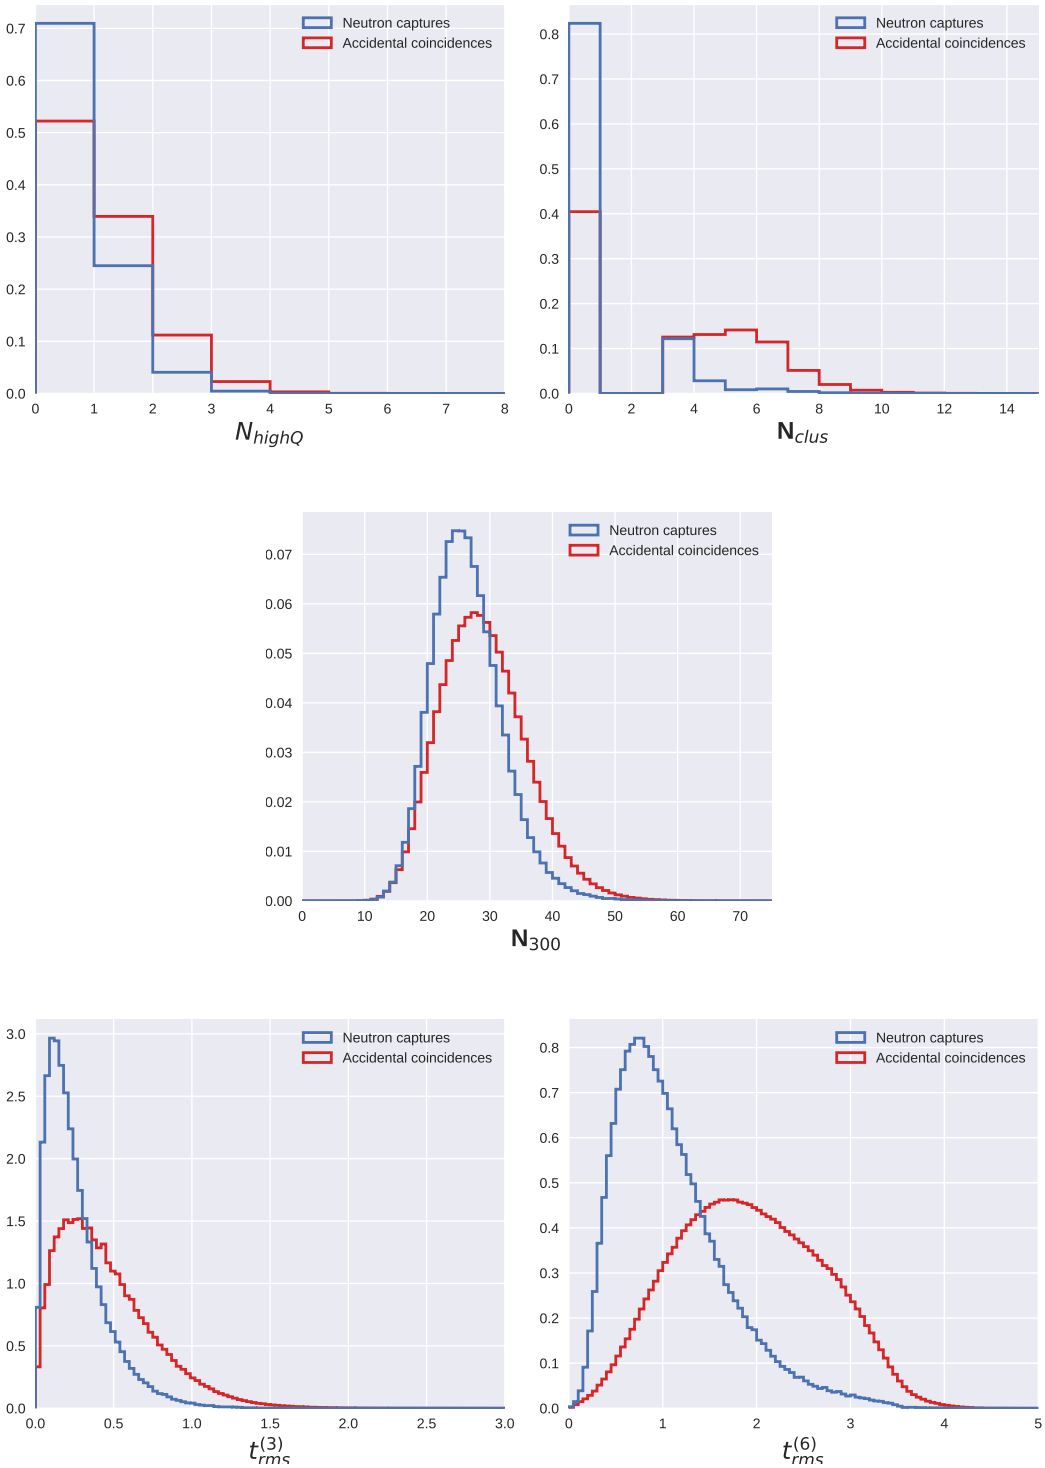

FIGURE 3.24: Probability density function of N_{highQ} (top left), N_{clus} (top right), N_{300} (middle), as well as $t^{(3)}_{\text{RMS}}$ (bottom left) and $t^{(6)}_{\text{RMS}}$ (bottom right), in nanoseconds, for accidental coincidences (red) and true neutrons captures (blue). The discontinuity in the N_{clus} distribution is due to the definition of a cluster as containing a minimum of three hits. A sizeable portion of the background is found at $N_{\text{clus}} > 4$ with nearly no signal, making it a particularly strong discriminator. Restricting the number of hits results sharpens the time dispersion peak of the signal to a greater extent than the background.

Finally, a subtlety when characterizing the background is that PMT hits from background sources can still be found within a signal. In other words, the signature of a true neutron capture can still be contaminated with one or more background hits. This effect can be observed by minimizing time dispersion for a subset of PMT hits within our candidate. For a true neutron, taking a subset of hits will allow us to exclude noise-like hits and further sharpen the t_{RMS} peak. Two observables are computed: $t^{(3)}_{\text{RMS}}$, and $t^{(6)}_{\text{RMS}}$, respectively the minimum t_{RMS} value found for a subset of 3 and 6 PMT hits. The neutron capture observables are summarized in Table B.1.

3.5.3 Boosted Decision Tree selection

We have now defined a range of discriminating observables characterizing various aspects of the neutron signal. The most straightforward way to exploit these variables would be to come up with appropriate cut thresholds for each one. One problem with this approach is the relatively large number of variables, many of which display significant correlation and interdependence between one another. Realistically, multidimensional cut criteria would need to be chosen to properly account for these relationships, rapidly increasing the complexity of the problem. This setup a classic use case for the application of multivariate machine learning techniques. In particular, we use a Boosted Decision Tree (BDT) algorithm to determine whether the neutron candidate is signal- or background-like.

BDTs

Using a machine learning algorithm allows, in principle, to exploit complex, multidimensional relationships that would be more difficult to capture with traditional methods alone. In general, a machine learning model $F(x;\lambda)$ with parameters λ is used to approximate a function $f(x)$ by leveraging sample data with known x_i and $y_i = f(x_i)$. This dataset is used to tune the parameters λ of the model F during *training*. The training itself is guided by the goal of minimizing a loss function $L(y, F(x;\lambda))$, which estimates the quality of the model, and whose form is chosen in accordance to the task. In our case, the object of the approximation will be the function that classifies a neutron candidate starting from the above observables. Several types of algorithms can be applied to a *classification* task like this one, each with their own strengths and weaknesses. The main alternative to the BDT in this case is a neural network (NN). In theory, NNs offer greater complexity, at the cost of longer, more resource-intensive, and more difficult training. A BDT, on the other hand, operates at a simpler level, but is easier and faster to train. This is not a trivial advantage, as modifications in the analysis are common, so using an algorithm that is quickly adaptable is desirable. Due to its (relative) simplicity, a BDT is particularly apt when working with variables that don't need extensive transformations applied to them in order to be useful, and therefore relies on smartly-chosen observables. Under these conditions, a BDT can match or even outperform a NN. An additional advantage to a BDT is its *interpretability*. Machine learning models are notoriously difficult to understand, often being described as *black boxes*, whose output is used without any knowledge of the algorithm's inner workings. In this respect, the BDT's simplicity is once again an advantage. Compared to other algorithms like a NN, we can more easily check the impact of each variable on the overall algorithm.

The building block of a BDT classifier is the *binary decision tree*. A binary decision tree is a recursive structure built by successively partitioning the initial data set using pre-calculated discriminating observables. Each partition uses one of the observables

to divide the dataset into two disjoint sets, each of which can be further subdivided in a similar fashion. The goal is to partition the feature domain into subdomains that can each be labeled with a single score. A cut can then be chosen on this score to define the input as signal or background. This process is not unlike a series of interdependent cuts, and is not particularly powerful on its own, even when fine-tuned during training. A BDT improves on this idea with *ensemble learning*, by leveraging the output of multiple decision trees at once. First, a constant function F_0 is chosen as an initial guess. Then, the model is improved iteratively, where at step n we add a decision tree t_n:

$$F_0(x) = c_0 \tag{3.43}$$

$$F_n(x) = F_{n-1}(x) + t_n(x) \tag{3.44}$$

$$F_m(x) = c_0 + \sum_{n=0}^{m} t_n(x), \tag{3.45}$$

where we stop the iteration at step m, obtaining the final BDT model F_m. This iterative model improvement is referred to as *boosting*. At each step, a tree is added such that the ensemble minimizes the loss function in the training dataset $\{x_i, y_i\}$:

$$t_n(x) = t(x; \lambda_n) \tag{3.46}$$

$$\lambda_n = \arg\min_{\lambda} \sum_i L(y_i, F_{n-1}(x_i) + t(x_i; \lambda)). \tag{3.47}$$

Here we have characterized the tree t_n as being parametrized by the generic tree function t with parameters λ_n. In each step n of the training, the new tree t_n of the model will attempt to complement the previous trees, progressively making up for the shortcomings of its predecessors. Since the λ parameters are usually numerous, their optimization is usually not trivial, as it involves finding the global minimum of the highly-dimensional $L(\lambda)$ function.

The neutron tagging BDT is implemented in Python with the `xgboost` software package, which allows us to use highly-optimized functions to develop our model. Two features set `xgboost` apart from similar libraries. First, great importance is placed on *regularizing* the model. This means that model simplicity is preferred: highly-complex models that reflect the training data well run the risk of capturing irrelevant characteristics of the dataset, rather generalizing the target function $f(x)$ (a pitfall known as *overtraining*). To this end a regularization term ω, penalizing model complexity, is added to the loss function in equation 3.47:

$$L(y, F_n(x)) \to O(y, F_n(x), t_n) = L(y, F_n(x)) + \omega(t_n), \tag{3.48}$$

where the aim is now to minimize the objective function O. A smart handling of the regularization is paramount to a successful training. The second key feature of `xgboost` is its optimization strategy. As its name (short for Extreme Gradient Boosting) might suggest, the minimum of O is looked for by estimating the function's gradient and descending this gradient in iterative steps. The algorithm used for this descent, leveraging Newton's method, is significantly more performant than traditional gradient descent strategies.

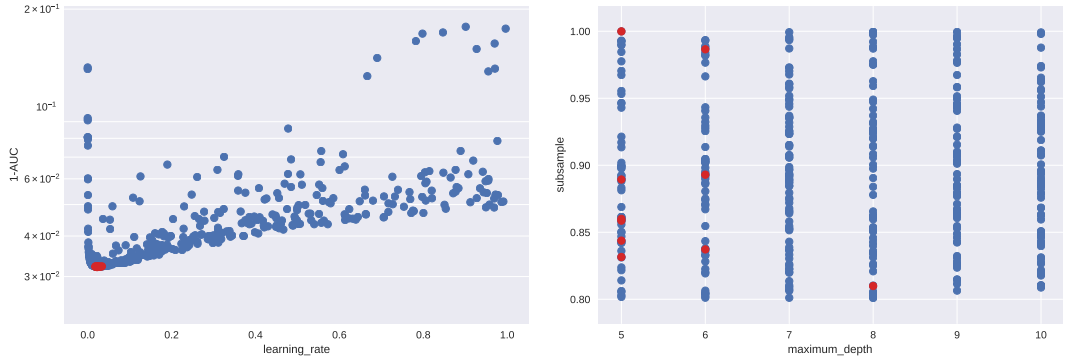

FIGURE 3.25: Example of a hyperparameter search to select the optimal BDT model, where each blue dot represents a single model, trained with a randomly-sampled combination of hyperparameters. Left: $1-$AUC metric as function of the learning rate. Lower values of $1-$AUC indicate better performance. Right: plot of the maximum depth and subsampling fraction tested. In both plots, the 10 best-performing models in this search are highlighted in red.

Model training and selection

Turning our focus back to our problem, the above strategy is applied to the dataset of neutron candidates from the IBD simulation. In total, 25 M simulated IBD events are used for the development of the model. From this sample, we extract 280 M accidental coincidences and 2 M neutron signals. Note that, proportionally, we take a much higher number of accidentals from the dataset. This is due to the inherent asymmetry of the problem: while we can hope to achieve a signal efficiency of $> 50\%$ (excluding the preselection efficiency), the background needs to be reduced by several orders of magnitude. Therefore, a larger sample is needed to estimate the model's performance with acceptable statistical uncertainty. The dataset is separated into three statistically independent, disjoint subsets: a training dataset, a validation dataset, and a test dataset, comprising respectively 50%, 25%, and 25% of the total. The training dataset is used to calculate the loss function and its gradients during training, as in the procedure above. The test dataset is used after training, to evaluate the performance of the model, by applying it to never-before-seen neutron candidates. The validation set is used in during model selection, to select the best model among several that were trained with different *hyperparameters*. Hyperparameters are those parameters that, instead of being tuned during training via gradient descent, are decided on by the user. While the neutron tagging is a binary classification problem, where we try to estimate a function $f(x) \in \{0, 1\}$, we train the BDT to output a score $F(x) \in [0, 1]$. This score can be thought of as capturing the neutron signal likelihood into a single variable. Down the line, this will allow us to cut on the BDT output at our own discretion, based on our signal efficiency and purity requirements.

A range of values was tested for several BDT hyperparameters: the learning rate (determining the size of each learning iteration step), the maximum number of trees, the maximum tree depth (determining the maximum amount of split points a given data point will undergo in each tree), and the event weighting the event weighting (assigning different importance to background and signal during training). In addition, the fraction of trees used during the gradient calculation is also included as a hyperparameter. Using less than 100% of the trees, chosen randomly from the total, is an additional tool to prevent overtraining. At each step of the training, the

Hyperparameter	Tested range	Chosen value
Learning rate	$[0.001, 1.0]$	0.025219
Max tree depth	$[3, 10]$	5
Max number of trees	$[800, 3000]$	1500
Early stopping threshold	/	50
Subsampling fraction	$[0.8, 1.0]$	0.97
Signal weight	$[0.1, 10]$	1.0
Training method	/	`auto`

TABLE 3.5: Hyperparameters chosen for the neutron tagging BDT model. *Early stopping threshold* is the number of iterations after which the training is stopped if no improvement in performance is detected, and was not optimized. The *Subsampling fraction* is the fraction of trees used for the gradient calculation. The *signal weight* determines the relative importance of signal candidates during training (so a weight of 2.0 means that the signal is given twice as much importance as the background). An equal weighting of signal and background was found to be optimal. The training method is a parameter internal to `xgboost` and was left in its default setting.

BDT cut	BDT efficiency	BDT bg acceptance	Overall ntag efficiency	Accidental coincidences/event
0.1	82%	1%	37%	0.11
0.5	67%	0.13%	30%	0.014
0.9	49%	0.014%	22%	0.0015
0.99	28.4%	0.00086%	12.8%	0.000095

TABLE 3.6: Example performance for four BDT cuts. The BDT efficiency and background acceptance are respectively the percentage of signal and background remaining after the cut, calculated on the test dataset. The overall efficiency takes into account the preselection, multiplying the preselection efficiency with the BDT efficiency. The number of accidental coincidences/event multiplies the BDT background acceptance with the accidental coincidences/event after preselection, and thus represents the expected number of accidental coincidences for each prompt event after the BDT cut. The statistical uncertainty on these values is negligible.

validation set was used to track progress. To define the model's performance and track its changes during training, the *ROC curve* is used. A ROC curve is a function defining the relationship between signal efficiency and purity, each a function of the chosen BDT output cut. The area under the ROC curve (AUC) is chosen as a metric of training progress. If the AUC did not improve for more than 50 iterations, the training was interrupted, and the iteration with the best performance was chosen as the final model, computing additional performance metrics with the validation set. Due to the relatively high number of hyperparameters tested, testing every possible hyperparameter combination for our desired range and granularity would be computationally prohibitive. For this reason, a random sampling among possible hyperparameters is used. This method has been shown to have similar efficacy as a thorough hyperparameter search but is much less computationally demanding. The final set of chosen hyperparameters, representing the model with the best AUC calculated on the validation set, is summarized in Table 3.5. The evolution of the AUC metric during training is shown in Fig. B.1, along with the ranking of each observable according to the relative importance in the algorithm.

After selecting the model, the final performance is calculated with the test dataset.

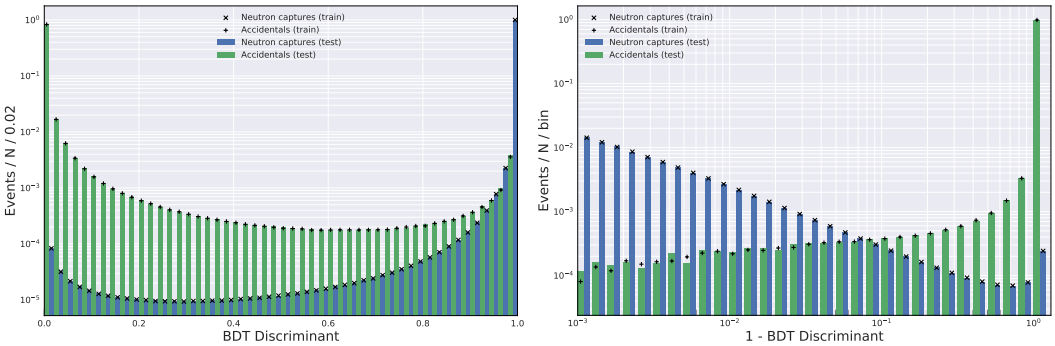

FIGURE 3.26: Left: BDT output value calculated on the test dataset, for signal (blue histogram) and background (green histogram). The BDT output is also plotted for the training dataset, for signal ($\times$) and background ($+$). Each distribution is normalized to an area of 1. Notice that the plot is in log scale, so the signal and background distributions are very strongly peaked at 1 and 0, respectively. The close correspondence between the training and test distributions confirms that the model did not suffer significant overtraining. Right: distribution of (1- BDT output), using a log scale for the x-axis. This is a different visualization of the same information, but better captures the behavior of our model around the high end of the BDT output distribution. Since we are interested in a high-purity sample, this plot is useful when discussing particularly tight cuts.

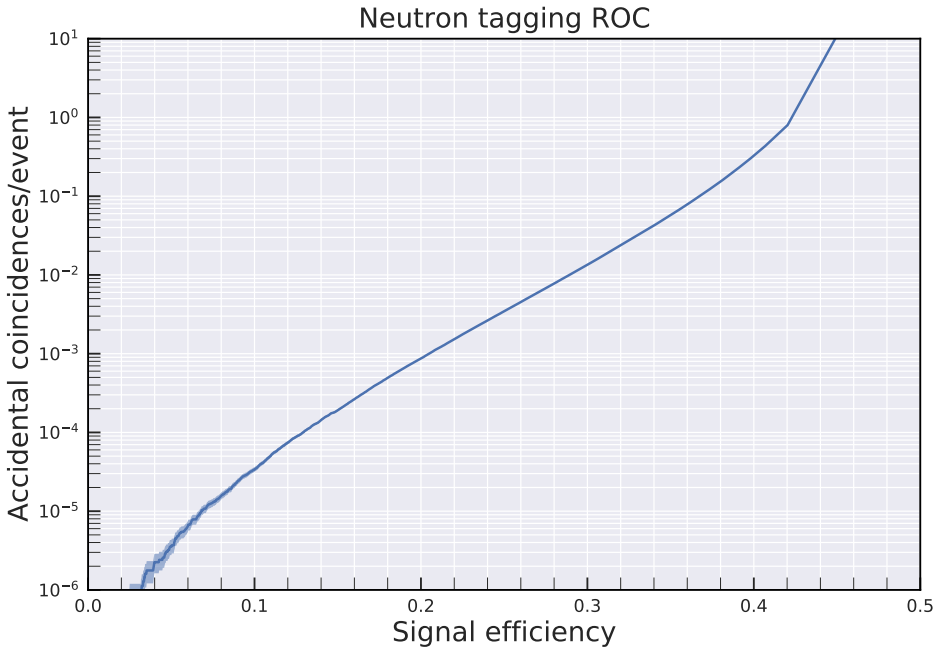

FIGURE 3.27: ROC curve of the neutron tagging BDT model, computing the accidental coincidence background rate as a function of signal efficiency for a range of BDT output cuts. The accidental coincidence background rate is the expected number of accidental coincidences per prompt event and can thus exceed 1. The signal efficiency is normalized by the preselection efficiency. The shaded area represents the statistical uncertainty on the background rate.

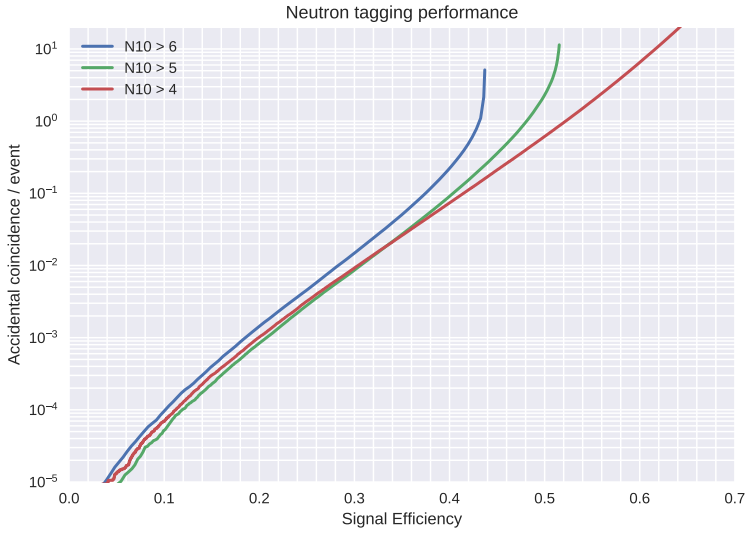

FIGURE 3.28: Comparison between the performance of neutron tagging algorithms for different preselection cuts. Each preselection cut is associated with a different ROC curve. Decreasing the N_{10} cut down to $N_{10} > 5$ brings measurable improvements. Further lowering of this threshold does not result in additional benefits, while making the training of the algorithm more difficult.

The distribution of the BDT output, computed for signal and background visualizes the discrimination power of the BDT. Comparing the same distribution with calculated on the training dataset gives us an impression of the amount of overtraining in our model. These distributions are plotted in Fig. 3.26. Overall, little overtraining was found to have occurred, implying that the model was able to successfully generalize the output function to make predictions about unseen neutron candidates. The ROC curve is shown in Figure 3.27. To give an impression of the performance achieved, the efficiency and background rates for a few example BDT cuts are listed in Table 3.6. Thanks to the large size of the test dataset, the statistical uncertainty on the BDT efficiencies are negligible. Overall, the neutron tagging cut will be the most inefficient one in the analysis. This is not a value judgement of the algorithm itself; rather, it mostly is a result of the inherent limitations of the detector. Much of the inefficiency lies in the preselection: candidates failing this cut contain very few PMT hits and are unlikely to bring significant improvement in performance.

The neutron tagging strategy above was repeated for several values of the preselection cut, namely the N_{10} cut. After training a neutron tagging BDT for each cut value, the relative ROC curves can be plotted together as in Figure 3.28, allowing us to compare the impact of the preselection cut on the final algorithm's performance. It was found that lowering the threshold below $N_{10} > 5$ brought no additional performance gains and made the convergence of the BDT model longer and more difficult.

3.5.4 Systematic uncertainties

The neutron tagging cut is subject to statistical uncertainty due to the sample size of the test dataset, on which signal and background efficiencies are calculated. In the case of the signal efficiency, this uncertainty is negligible. The statistical uncertainty on the background rate (visible in Figure 3.27), on the other hand, only becomes

FIGURE 3.29: Top: background acceptance of the BDT with respect to the BDT cut, computed for different SK-IV periods. Each color corresponds to a different time period, centered at the run number indicated. Bottom: estimated systematic uncertainty in the background rate. For relevant BDT cuts, the uncertainty is below 10%.

relevant at very tight cuts, which are not used in this analysis. The uncertainties associated with this cut are dominated by systematic uncertainties. Since the IBD simulation models neutron captures and accidental coincidences in fundamentally different ways, the strategy for the uncertainty estimation will also be very different between signal and background efficiencies.

Accidental coincidence background rate

Since accidental coincidences are taken directly from data, there is no systematic uncertainty associated with the mismodeling of this background. However, within the data, a significant variation in the dark noise rate was observed over the 10-year running time of SK-IV. The overall accidental rate calculate above therefore averages this variance out. Further, the time period covered by the background data, sampled from dummy trigger runs, does not match the time period covered by the SHE+AFT data in one-to-one fashion, due to differences in the availability of the data runs. For this reason, the variance in the background rate due to time-dependence of the dark noise is taken as a systematic uncertainty. To estimate this variance, the background data is divided into 10 time bins, separating the SK-IV period into subsets of roughly eight months each. For each BDT cut, the uncertainty is computed as the maximal difference in the background rate between two runs periods, at the given cut point.

Signal efficiency

Systematic uncertainty on neutron tagging efficiency stems from mismodeling of the 2.2 MeV neutron capture signal in the simulation samples used for the BDT training. To estimate the magnitude of the uncertainty, the efficiency achieved on the simulated samples is compared with efficiency obtained from real neutron captures in data. To do so, we use data samples containing real neutron captures from Am/Be

Source point	x-position (cm)	y-position (cm)	z-position (cm)
Center	35.3	−70.7	0.0
Y−12	35.3	−1201.9	0.0
Z15	35.3	−70.7	1500.0

TABLE 3.7: Tank positions of the Am/Be radioactive source during calibration. The origin corresponds to the geometrical center of the detector volume.

calibration runs. During Am/Be calibration, a source containing Americium-241 and Beryllium-9, embedded in a Bismuth-Germanium oxide (BGO) scintillator is placed in the detector. The source produces pairs of prompt and delayed neutron captures signals, mimicking the $\overline{\nu}_e$ IBD signal. $^{241}\mathrm{Am}$ decays by emission of an α-particle, which interacts with $^{9}\mathrm{Be}$ with primarily two possible neutron production channels:

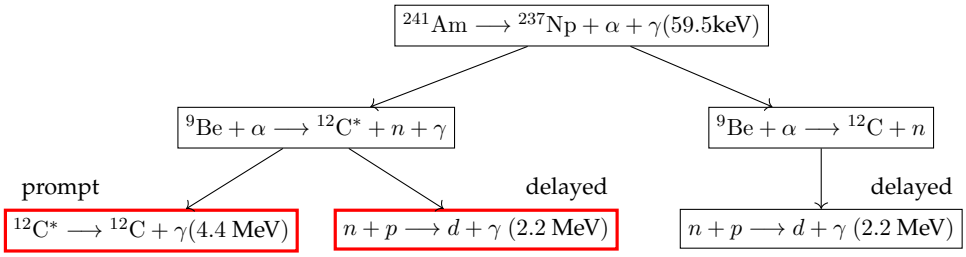

Only the first of the above channels produces a prompt event, wherein the 4.4 MeV photon is used as the primary event SHE trigger, and after which a 2.2 MeV delayed neutron capture signal is produced. This signal is identical to an IBD neutron, and is therefore suitable for testing the tagging algorithm. To search for the coincidence event, the same AFT+SHE trigger setup is used as for the DSNB search, except that the AFT trigger used for these calibration runs saves data in a larger window, lasting 800 μs after the 35-μs SHE event, instead of the 500 μs for regular data runs. The $^{12}\mathrm{C}$ decay, serving as the primary event, occurs inside the BGO scintillator, which is necessary to amplify its 4.4 MeV photon signal, in order to use the SHE trigger. The channel in which $^{9}\mathrm{Be}$ immediately decays to the $^{12}\mathrm{C}$ ground state also produces a delayed neutron, but is not associated with a prompt event trigger. Therefore, the associated neutron will only be detected in coincidence with an uncorrelated prompt event, and thus constitutes a background for our purposes.

The source is placed at three different locations inside the tank: near the center of the detector, near the barrel wall, and near the top wall (Table 3.7). Data is taken from two calibration runs during SK-IV, in 2009 and 2016. In addition, background data runs were taken in both years using a random trigger, with the source in the detector at the center position.

The prompt event itself, producing a scintillation signal, does not resemble the IBD positron, and therefore requires its own selection criteria. The scintillation from the prompt event is associated with a characteristic light deposition, and can be simply selected for using the number of photoelectrons (p.e.) detected inside the

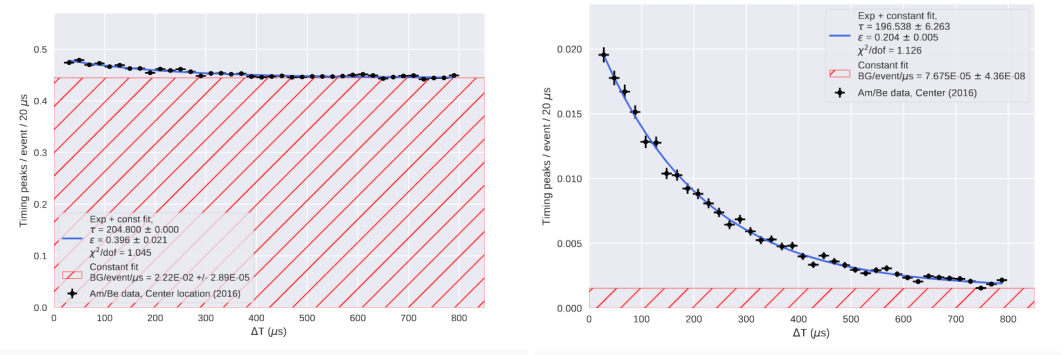

FIGURE 3.30: Time after trigger time (ΔT) of the neutron candidates in the Am/Be calibration data sample taken in 2016 in the center position, fitted with an exponential function plus a constant term. Left: before application of the BDT cut. Right: after application of a chosen BDT cut.

inner detector, depending on the source position:

$$\begin{aligned} \text{Center}: &\quad 750 < \text{p.e.} < 1050 \\ \text{Y} - 12: &\quad 850 < \text{p.e.} < 1150 \\ \text{Z15}: &\quad 900 < \text{p.e.} < 1150. \end{aligned}$$

We further require that each scintillation event be separated by at least 1.5 ms from the previous trigger to avoid contamination from scintillations that are not well-separated in time. Finally, within each event, the maximum of the PMT hits in a sliding 200 ns window after the prompt trigger is measured. If this maximum is greater than 50, the event is rejected, in order to remove events in which a second scintillation occurs shortly after the first, within the same trigger time window. This set of cuts defines our prompt event, for which the neutron candidates are selected using the usual strategy.

To compare the efficiency on the Am/Be calibration data with the efficiency on MC samples, we use simulation samples that are produced to match each calibration data run. At the time when the study was conducted, no available simulation was available that could directly simulate the Am/Be and scintillator setup—specifically, no modelling of the BGO scintillator was available—so a simulation of IBD events, tailored to the calibration, runs was developed. In particular, the simulation samples used for this study are produced following the same procedure as those described in Section 3.1, constraining the prompt vertex to correspond to the Am/Be source location, tuning the simulation to match to the calibration run in question, and injecting background taken from the Am/Be random-trigger background runs. The fact that the simulation does not fully model the Am/Be results in a significant discrepancy between the simulated prompt event and the Am/Be prompt event. It is therefore an important limitation of this study, and needs to be treated delicately.

Since in the Am/Be data we don't know which neutron candidate corresponds to a true neutron capture, and which to an accidental coincidence, we can't evaluate the performance of the algorithm on an event-by-event basis. Note that this is also why these data samples cannot be used in lieu of a simulation to train the BDT. However, given the copious amount of neutrons produced by the Am/Be source, we can extract the efficiency of the algorithm on a statistical basis. To do so, we exploit the time-dependence of the neutron capture, whose probability decays exponentially

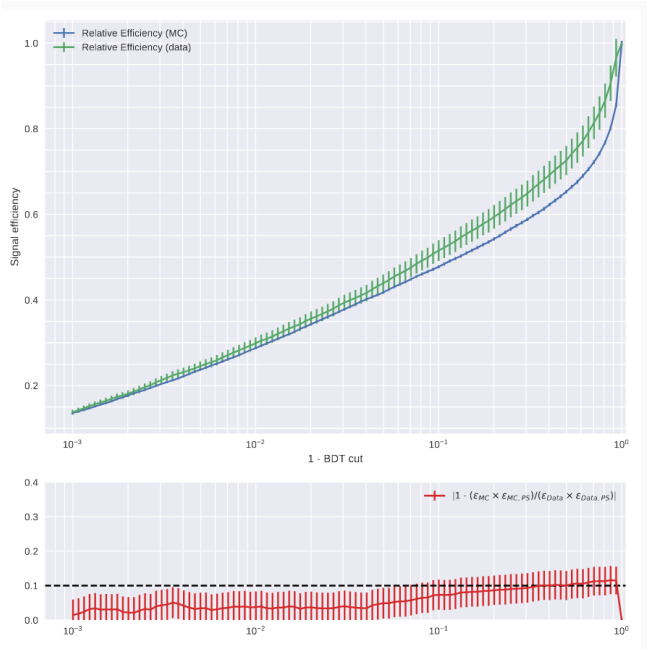

FIGURE 3.31: Relative BDT efficiency on the 2016 Am/Be data samples (green) and on the corresponding simulation (blue), as well as the estimated uncertainty in red, calculated as the relative difference between the two.

according to the capture time constant, $\tau_{cap} \approx 200\ \mu$s. First, each neutron candidate is tagged according to a chosen BDT cut. Then, the distribution of the time difference δT, between the neutron candidates and the prompt event, for those candidates that pass the neutron tagging selection, are fit with a decaying exponential plus a constant term B. While true neutron capture event counts decay exponentially in time from the time of positron emission, uncorrelated backgrounds are constant in time. The efficiency ϵ_{ntag} can then be estimated by taking the ratio of the number of candidates contributing to the exponential component of the fit, N_S, to the total prompt Am/Be events N_{prompt}:

$$\frac{dN_n}{d\Delta T}(\Delta T) = N_0 e^{(\Delta T/\tau_{cap})} + B \tag{3.49}$$

$$\epsilon_{\text{ntag}} = \frac{N_S}{N_{\text{prompt}}} = \frac{\int_{\Delta T} \left(\frac{dN_n}{d\Delta T} - B\right) d\Delta T}{N_{\text{prompt}}}, \tag{3.50}$$

where $\frac{dN_n}{d\Delta T}$ is the number of neutron candidates per ΔT bin and N_0 is a normalization constant. During the fit, N_0, τ_{cap} and B are free parameters. An example of such a fit is given in Figure 3.30. This approach however does not account for the intrinsic difference in the prompt event between data and simulation described above. Specifically, the prompt events selected in the Am/Be data is not a pure sample, including scintillation-induced backgrounds that are not well-understood. This effect increases the value of N_{prompt} and artificially lowers the estimated neutron tagging efficiency.

To tackle this issue, we separate the treatment of the uncertainty on the efficiency into the uncertainty associated with the efficiency ϵ_{PS} of the N_{10} preselection cut, and that associated with the *relative* BDT cut efficiency ϵ_{BDT}:

$$\epsilon_{\text{ntag}} = \epsilon_{\text{PS}} \times \epsilon_{\text{BDT}}. \tag{3.51}$$

BDT cut	Position	Year	Data	MC	MC/Data
0.40	Center	2009	69.4	76.2	1.098
		2016	80.0	72.3	−1.095
	Y−12	2009	74.6	77.4	1.037
		2016	83.8	77.6	−1.072
	Z15	2009	87.6	77.1	−1.120
		2016	80.2	74.1	1.077
0.90	Center	2009	50.0	55.3	1.106
		2016	56.1	52.8	−1.058
	Y−12	2009	56.4	59.6	1.057
		2016	60.8	59.2	−1.026
	Z15	2009	65.8	58.3	−1.114
		2016	57.0	55.4	−1.028
0.99	Center	2009	33.6	35.5	1.058
		2016	34.9	33.2	−1.048
	Y−12	2009	36.9	37.8	1.024
		2016	38.2	37.1	−1.028
	Z15	2009	38.8	34.5	−1.105
		2016	31.9	32.2	1.001

TABLE 3.8: Observed (Data) and predicted (MC) relative efficiencies (in %) for different BDT cuts and for the Am/Be calibration runs considered in this study. Data was taken in 2009 and 2016, at three different locations in the tank: in the center of the detector, near the barrel wall (Y−12) and the top wall (Z15).

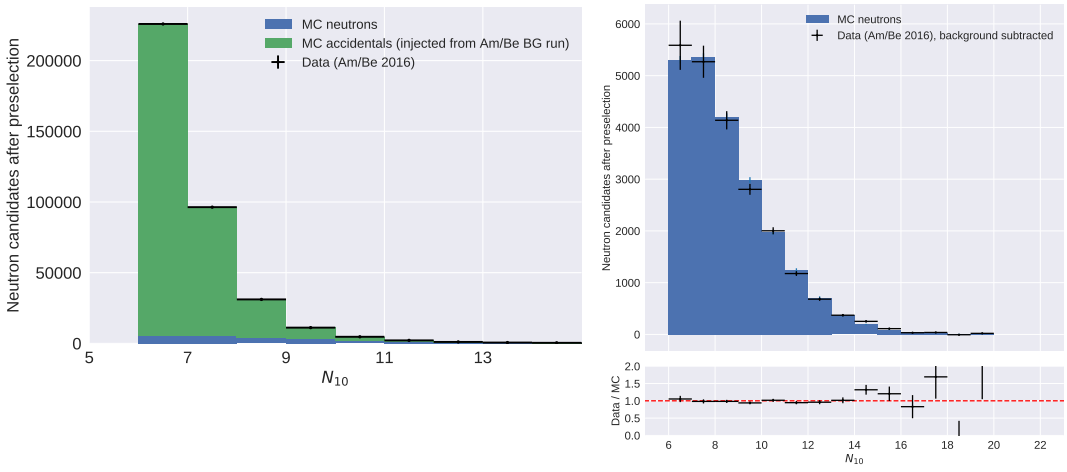

FIGURE 3.32: Left: N_{10} distribution of Am/Be data and simulation samples, for the 2016 run in the center location. The simulated sample is divided into true neutron captures (blue) and accidental coincidences (green). Right: N_{10} distribution for true neutron captures from the simulation with the best-fit QE value, and from background-subtracted data.

Since N_{prompt} is the denominator to both ϵ_{ntag} and ϵ_{PS}, it cancels out in the computation of the BDT efficiency, allowing for direct comparison of simulation and data efficiencies:

$$\epsilon_{\text{BDT}} = \frac{\epsilon_{\text{ntag}}}{\epsilon_{\text{PS}}} = \frac{N_S}{N_S^{\text{PS}}}, \tag{3.52}$$

where N_S^{PS} is the number of neutron captures extracted as in 3.50, but with only the preselection cut applied. The relative BDT efficiency for the 2016 samples in the center position, are compared for data and simulation in Fig. 3.31, for a range of BDT cuts. Table 3.8 gives the same comparison for all the Am/Be runs and source locations, at different BDT cut points. Over all data runs and the considered BDT cut points, the maximum data-simulation discrepancy found of 12%, in relative terms.

The statistical uncertainty on the efficiency ϵ_{PS} of the N_{10} preselection cut, on the other hand, will depend on N_{prompt} and is affected by the distortion effects caused by the scintillation. The scintillation-induced prompt events will be associated with uncorrelated backgrounds, and no correlated neutrons. Therefore, the N_{10} distribution in data will be biased towards a background-like distribution. To correct for this effect, we compute N_{10} for the background run taken with the Am/Be source, and subtract this from the regular Am/Be data sample to obtain a pure-signal N_{10} distribution. Since the number of prompt events is not comparable between data and MC, the *absolute scale* of this signal distribution cannot be used directly for comparison. Therefore, we use the *shape* of the N_{10} distribution as a handle on the data-simulation discrepancy. The N_{10} distribution is based on a PMT hit count, and is therefore highly sensitive to the quantum efficiency (QE) in the simulation. We therefore use the QE parameter in the simulation to parametrize the distortion of the N_{10} distribution. Several QE values are tested around the QE of the regular simulation samples, in order to find the QE value that best fits the N_{10} distribution in data. The best-fit QE is found to be 2.1% higher than the one used in our simulation. The N_{10} distribution associated with this value results in a preselection efficiency that is 3.0% larger than in the standard simulation. The 1σ uncertainty on the best-fit QE, found using a χ^2 fit, is also associated with an additional 2.1% uncertainty on the preselection efficiency. The two uncertainties are added in quadrature to yield an overall preselection efficiency uncertainty of 3.7%. Combining the uncertainties from preselection and BDT, we obtain a maximum overall uncertainty on the efficiency of the neutron tagging algorithm of 12.5%.

3.5.5 Cut optimization

By definition, the neutron tagging cut described above allows for the removal of any background that is not associated with a neutron. In the energy range of interest, solar neutrinos are easily wiped out. The remaining background contributions are then from reactor neutrinos, spallation and atmospheric backgrounds. The main target of the cut will be the removal of spallation events, which—other than the ^{9}Li isotope—produce no neutron in the corresponding β-decay. In addition, by requiring exactly one neutron to be tagged, we can also target backgrounds associated with the production of multiple secondary neutrons. This is relevant for several types of atmospheric backgrounds, which can be characterized by high secondary neutron multiplicities. Since the dedicated cuts described in Appendix A and Ref [49] also target spallation, the neutron tagging and spallation cuts are optimized concurrently.

For the purpose of the statistical analysis in the following chapter, very low levels of spallation are required. As will be discussed, the analysis that follows relies on assumptions about the functional form of the remaining backgrounds after all cuts are applied. Our ability to model these backgrounds is limited by the lack of a simulated spallation sample, giving rise to large uncertainties in the spallation spectrum. Therefore, the spallation and neutron tagging cuts are optimized to bring the spallation event rate number to $O(1)$ level in the data sample. After this prerequisite is met, the optimal neutron tagging cut is defined as the cut maximizing the overall significance, calculated accounting for all sources of remaining backgrounds.

4
Spectral analysis of SK-IV data

The goal of this chapter will be to derive upper limits on the DSNB flux. After having defined, in the previous sections, the background reduction strategies in this DSNB search, we move on to the statistical treatment of the final data sample surviving the cuts. As part of the SK-IV DSNB search using these cuts, two different analyses were carried out. The first was a binned, model-independent analysis that focused on the lower end of the spectrum, in the $8-30$ MeV reconstructed energy range. More details can be found in Refs. [49]. In this text, the second analysis is discussed, namely an unbinned, model-dependent analysis that aims to directly fit the DSNB spectrum to the energy distribution of the data, together with the expected contributions from background spectra. One of the advantages of this method is that each background contribution is treated as a nuisance parameter to the signal. This analysis focuses on a higher energy range than the model-independent search, namely on the $16-80$ MeV reconstructed energy range. The choice of energy threshold is motivated in large part by the large uncertainty associated with the spallation backgrounds at lower energies. At these low energies, spallation is expected to dominate even with tight neutron tagging cuts. As will become clear, our treatment of the data will rely on our understanding of the spectral form of the backgrounds. In the $16-80$ MeV region, the behavior of the final spectrum is expected to be dominated by the well-known energy spectrum of the Michel electrons. This analysis is inspired by similar treatments used in DSNB searches with SK-I, SK-II, and SK-III data [68, 36]. This method will also allow us to combine our results with these analyses, reaching an overall exposure of 22.5×5823 kton·day, and arrive at model-dependent upper limits for the DSNB.

In Section 3.4, we discussed the Cherenkov angle cut, and how, in addition to defining a signal region, it is useful for defining background-dominated sidebands. This strategy was already used in previous analyses of SK data [68, 36]. As part of the SK-IV analysis, this method was adapted for the introduction of neutron tagging. Neutron tagging allows for higher signal purity, but in pure water—before the addition of Gd—this comes at the expense of significant efficiency loss. This loss in efficiency is a disadvantage of the model-independent analysis in Reference [49], as it only takes into consideration events passing the neutron tagging cut. In order to obviate this loss, we take also take into account events that fail the neutron tagging cut. We then extend the three-region paradigm by further dividing the parameter space into an *IBD-like* region, in which a single neutron was found in coincidence with the prompt signal, and a *non-IBD-like* region, in which either no neutron, or multiple neutrons were found in coincidence. These regions are defined for each Cherenkov angle region, obtaining six regions in total, as summarized in Table 4.1. The true signal region—the IBD-like region in the central Cherenkov angle region—will be the most sensitive to the DSNB signal. However, the non-IBD-like region will still hold a large fraction of the signal.

		Cherenkov angle		
		$20-38^\circ$	$38-50^\circ$	$78-90^\circ$
N_{ntag}	0 or >1	μ/π	e^+/e^-	NC
	1	$\mu/\pi + 1n$	$e^+/e^- + 1n$	NC+$1n$

TABLE 4.1: Overview of the regions used in this analysis. We subdivide the parameter space according to the Cherenkov angle and the number of tagged neutrons N_{ntag}. Regions with small and large Cherenkov angles are dominated by events with visible muons and pions, and by NC atmospheric events respectively.

4.1 PDF modelling

After preliminary reduction cuts, spallation cuts, positron identification cuts, and neutron tagging, the backgrounds above 16 MeV are dominated by atmospheric neutrinos. In order to perform spectral fits we define, in addition to the DSNB spectrum, several complementary background categories and derive associated probability distribution functions (PDFs). The backgrounds are separated according to their spectral shapes, such that the overall background contamination can be parametrized as a linear combination of each background PDF. Specifically, we define four atmospheric neutrino background categories and a spallation background category.

1. Decay electrons

The first category comprises electrons and positrons from the decay of invisible muons and pions generated by CC interactions of atmospheric neutrinos. The spectral shape of this background, following the Michel spectrum, can be estimated directly from data, using a sample of cosmic muon events from SK. By using the reconstructed energy distribution of the associated decay electrons from data, this spectrum has negligible systematic uncertainty associated with its shape. As an electron background, it will dominate in the signal Cherenkov angle region, but will leak into the sidebands due to resolution effects. The proportion of background leaking into the sidebands, as well as the contribution to the two neutron-tagging regions, is estimated from the atmospheric simulation.

2. Electrons from CC ν_e interactions

In the second category we find electrons and positrons produced in CC interactions of electron neutrinos and antineutrinos, with no visible muons and pions in the final state. As with the previous category, this electron background is mainly found in the signal angle region. Its shape and contribution to each region is estimated from the atmospheric simulation, using the truth-level information to select the appropriate events.

3. Muon and pions from CC interactions

The third category includes visible muons and pions, produced in a variety of CC processes. As mentioned in Section 3.4, due to their larger momenta muons and pions display a relatively low Cherenkov opening angles in the low-Cherenkov-angle sideband. Their impact will be negligible in other regions. Here, too, we use the

atmospheric simulation to estimate the spectral shape and contribution to the regions. In particular, simulated events interacting via CC, with the production of muons or pions, but no visible electrons, are selected.

4. NC interactions

The last atmospheric background category includes all events from NC interactions of atmospheric neutrinos that do not include visible electrons in the final state. The associated spectrum peaks at low energies, and is therefore one of the most problematic. While the Cherenkov angle cut places most of this background in the high-Cherenkov-angle sidebands—where it vastly dominates—a significant portion leaks into the signal region. Once again, the spectral shape and region contribution is estimated from simulation.

5. Remaining spallation events

Despite the aggressive spallation cuts used in this analysis, as described in Ref. [49], some spallation events could still remain. Modelling the spallation background was not part of this work, but is described here for completeness. In absence of a full spallation simulation at the time of the analysis, a parametric description was developed. The spallation spectrum will be strongly peaked at low energies, and will depend on its isotope composition. In our energy range, three spallation isotopes could contribute to this background, namely ^{8}B, ^{8}Li, and ^{9}C. Each isotope i can contribute a β-decay spectrum S_i, yielding a PDF of the form

$$f_{\text{spallation}}(E) = n_0 S_{^{9}\text{C}}(E) + (1-n_0) n_1 S_{^{8}\text{Li}}(E) + (1-n_1) S_{^{8}\text{B}}(E), \tag{4.1}$$

where n_0 and n_1 describe the relative contribution of each spectrum. Since this background is composed of electrons, with no secondary neutrons, it is found predominantly in the Cherenkov angle signal region, and is not found in the IBD-like region. The treatment of this background can be found in more detail in Ref. [49].

6. DSNB signal

For a given DSNB flux model, the IBD simulation from 3.1 is reweighted according to the model, after applying preliminary reduction and positron identification cuts. Since this is a model-dependent approach, each DSNB model considered will entail a different fit.

From distributions to PDFs

While the DSNB and spallation spectra are already continuous functions and can be readily interpreted as PDFs, the four atmospheric background categories are binned energy distributions, either from simulation or data. Therefore, a piecewise polynomial function is fitted to these energy distributions to obtain the atmospheric PDFs after positron identification cuts. Each PDF $f(E;\theta_{\text{C}}, N_{\text{ntag}})$ is defined to sum up to 1 across the six parameter space regions:

$$\sum_{\substack{\theta_{\text{C}} \\ \text{regions}}} \sum_{\substack{N_{\text{ntag}} \\ \text{regions}}} \int f(E;\theta, N_{\text{ntag}}) dE = 1. \tag{4.2}$$

The atmospheric neutrino PDFs are plotted in Figures 4.1 and 4.2. In the final spectral fit, the relative normalization of the DSNB and of each background source is varied to find the best-fit event numbers, so there is no need to correct for efficiencies that only affect the overall scale of these distributions. Energy-correlated cuts will however distort their shape and need to be taken into account. The atmospheric background distributions are obtained from the simulated sample, after application of preliminary cuts and positron identification cuts, but without the application of the spallation and solar cuts. Since these cuts are energy-dependent, the starting atmospheric PDFs are then convolved with the energy-dependent spallation cut and solar cut efficiencies. A similar procedure is applied to the DSNB PDF.

4.2 Spectral shape fit

Armed with our signal and background PDFs, whose shapes reflect our choice of reduction cuts, we can now proceed to fit the data samples after cuts. For each PDF f_j, where j denotes either the DSNB signal or one of the background categories, we fit the number of observed events N_j. Let's denote the observed reconstructed energies in the data sample as $\vec{E} = \left\{E^1, ..., E^{N_{\text{data}}}\right\}$, for N_{data} events in data. We then use an extended maximum likelihood method to find the best-fit value of the $\vec{N} = \{N_j\}$ vector. In particular, we perform a simultaneous fit on all the parameter space regions, aiming to find the N_j vector that maximizes the following likelihood:

$$\mathcal{L}(\vec{N}; \vec{E}) = e^{-\sum_j N_j} \prod_{i=1}^{N_{\text{data}}} \sum_j N_j \, f_j(E^i; \theta^i_{\text{C}}, N^i_{\text{ntag}}). \tag{4.3}$$

The best-fit number of true DSNB interactions via the IBD channel can then be simply obtained from the signal component N_S^{best} of the best-fit number vector $\vec{N}^{\text{best}}$:

$$R^{\text{best}}_{\text{DSNB}} = \frac{N_S^{\text{best}}}{\varepsilon_S}, \tag{4.4}$$

where ε_S is the overall signal efficiency of the analysis. The ultimate goal of our exercise, the extraction of the DSNB flux can be finally obtained as

$$\Phi^{\text{best}}_{\text{DSNB}} = \frac{R^{\text{best}}_{\text{DSNB}}}{FV_{\text{SK}} \times n_p \times \sigma_{\text{IBD}} \times T_{\text{SK-IV}}}, \tag{4.5}$$

using the IBD cross-section σ_{IBD}, the relevant SK interaction volume (i.e. the fiducial volume) FV_{SK}, the proton number density of water n_p, and the total livetime $T_{\text{SK-IV}}$ of the SK-IV data considered. To obtain the uncertainty of the fit, we can maximize the likelihood over the number of observed background events $\vec{N}_B$, such that we obtain the marginal likelihood $\mathcal{L}_S$, depending only on the number of signal events N_S:

$$\mathcal{L}_S(N_S) = \mathcal{L}(N_S, \vec{N}_B^{\text{best}}(N_S)). \tag{4.6}$$

The width of this marginal likelihood can then be interpreted as a measure of the fit uncertainty.

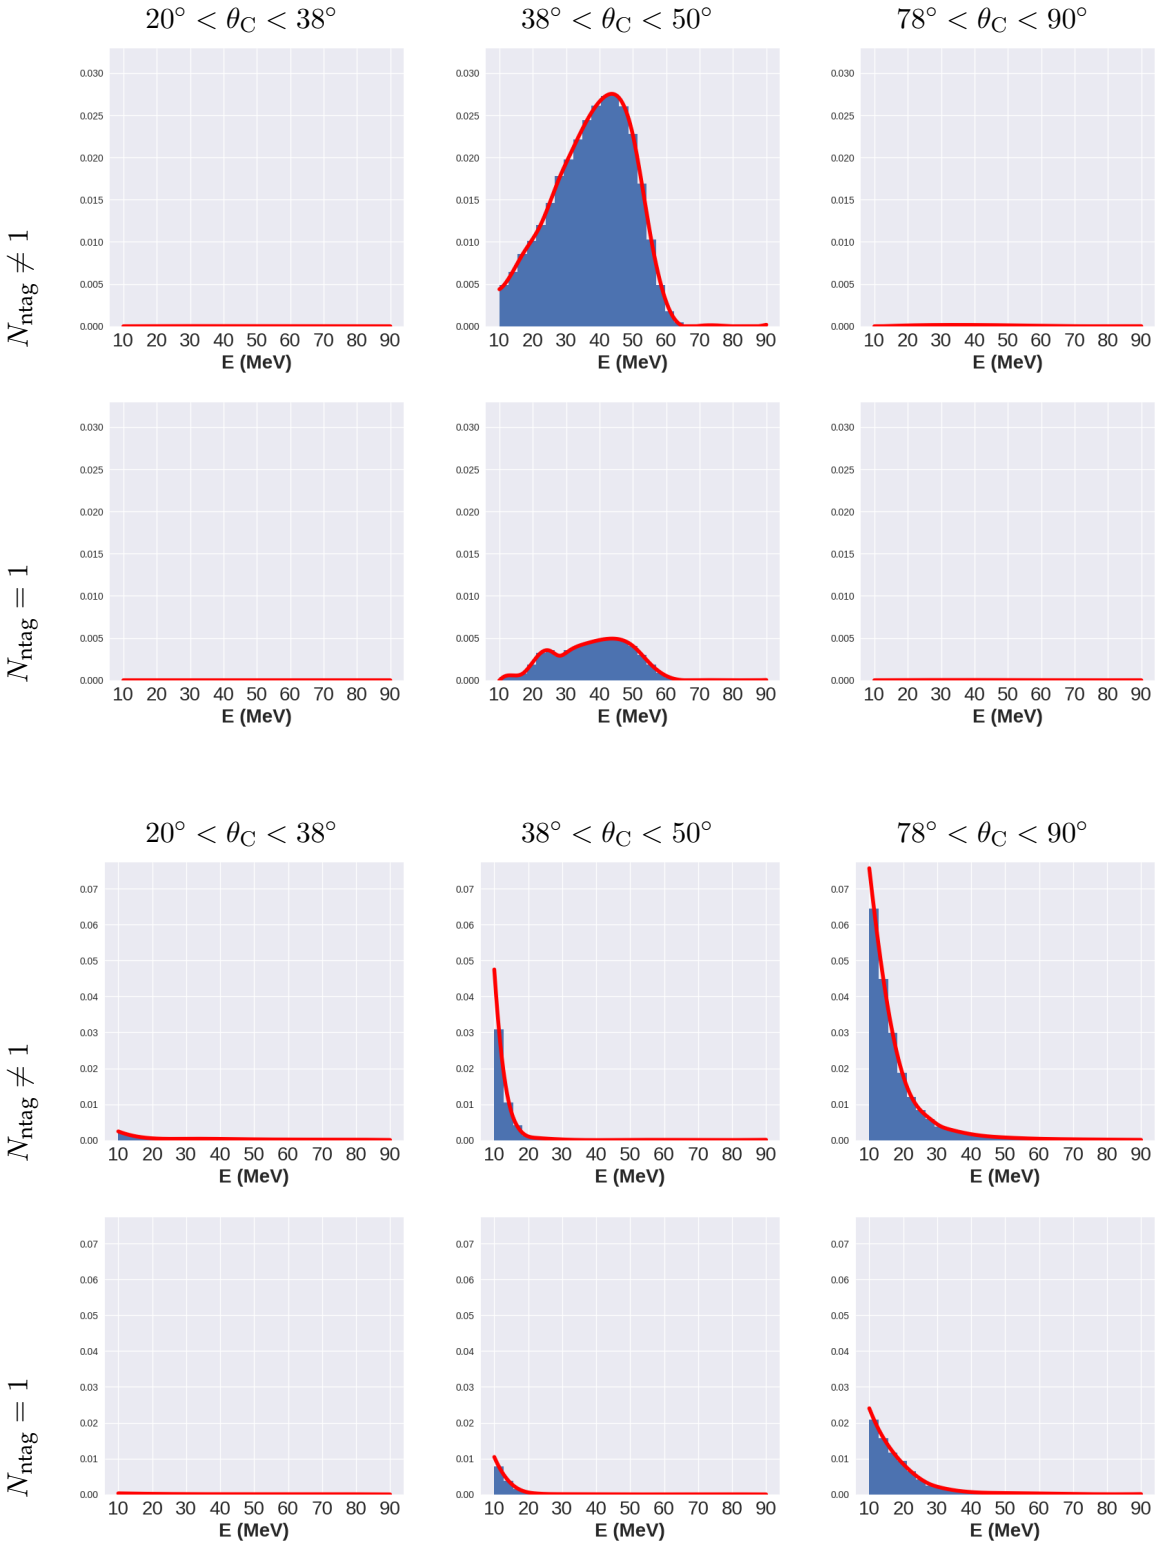

FIGURE 4.1: Decay electron (top) and Neutral Current interaction (bottom) background PDFs in the six fitting regions. The first, second and third columns correspond to the low-Cherenkov angle sideband, signal region, and high-Cherenkov angle sideband, respectively. The first and second rows correspond to the non-IBD-like (0 or > 1 tagged neutrons) and IBD-like (1 tagged neutron) regions, respectively. Each PDF is normalized to 1 over all regions.

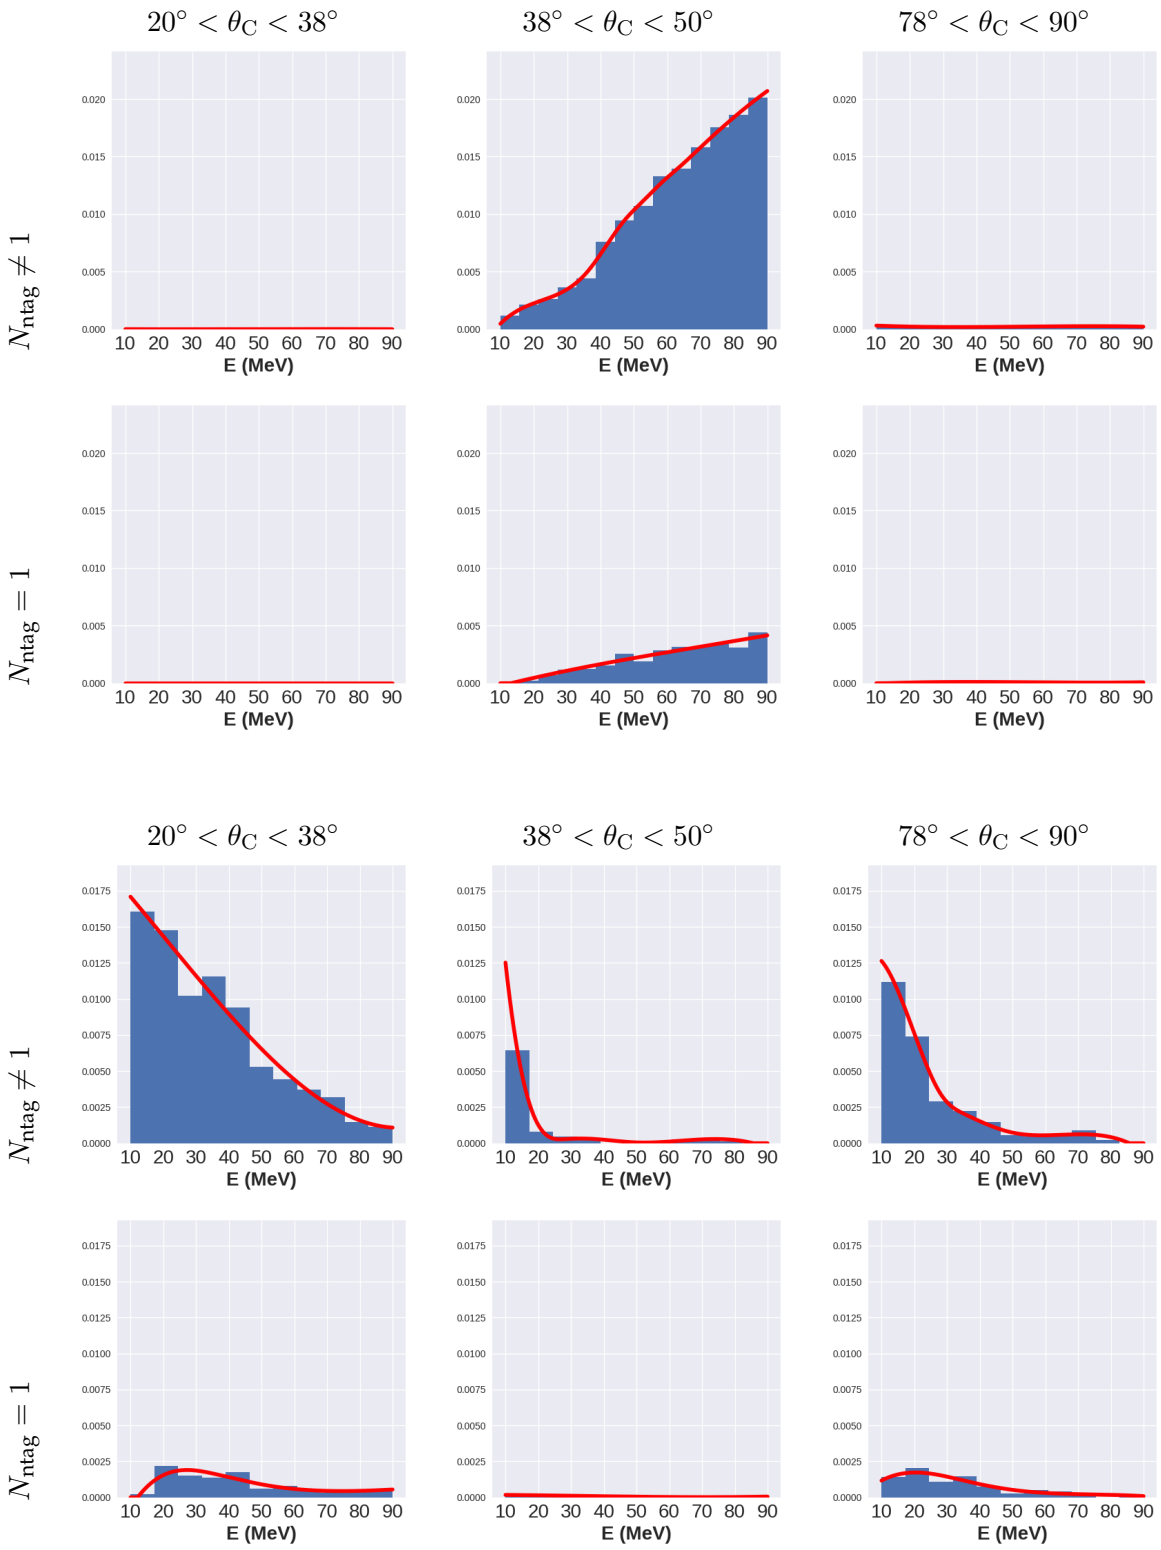

FIGURE 4.2: CC ν_e (top) and μ/π (bottom) background PDFs in the six fitting regions. The first, second and third columns correspond to the low-Cherenkov angle sideband, signal region, and high-Cherenkov angle sideband, respectively. The first and second rows correspond to the non-IBD-like (0 or > 1 tagged neutrons) and IBD-like (1 tagged neutron) regions, respectively. Each PDF is normalized to 1 over all regions.

4.3 Systematic uncertainties

The systematic uncertainties in the spectral fit are implemented as modification to the likelihood function, so that they will be reflected in the uncertainty of the final fit. Each systematic uncertainty is parametrized by a single parameter, and assigned a prior probability density function. The nuisance parameters are then also fit alongside with $\vec{N}$, and the likelihood function in Equation 4.3 is convolved with the prior probabilities of these parameters. Significant uncertainty is associated with the atmospheric simulation spectra due possible mismodeling [36]. Nevertheless, uncertainty in the decay-electron spectrum, taken from data, can be neglected. In addition, the uncertainty in the μ/π spectrum is also neglected, as the spectrum was found to have little impact in the fit.

Energy scale and energy resolution uncertainties

The energy scale and energy resolution uncertainties associated with the simulation are obtained from the SK-IV solar analysis [67], and adapted to account for the higher energy range of this analysis. These uncertainties amount to 2% and 3%, respectively. We model their impact by distorting both the signal and background PDFs, except for the data-derived decay-electron PDF. We parametrize the energy scale uncertainty as an overall shift in the PDFs, while for the energy resolution we apply the distortion functions computed in Ref. [69]. We parametrize the size of the distortion applied to the PDFs with the parameter ϵ, such that $\epsilon = 1$ represents a 1-sigma deviation in the energy scale and resolution. The new likelihood can then be expressed as:

$$\mathcal{L}' = \int_{-\infty}^{+\infty} e^{-\epsilon^2/2} \mathcal{L}(\epsilon)\, \mathrm{d}\epsilon \tag{4.7}$$

where $\mathcal{L}(\epsilon)$ represents the initial likelihood with the PDFs distorted by an amount ϵ. These uncertainties were found to have a negligible impact on the final result.

Atmospheric CC ν_e spectral shape

Below 100 MeV, the spectrum of ν_e CC processes has not been studied in depth to date. To model the systematic uncertainty associated with the spectral shape of the atmospheric CC ν_e backgrounds, we follow the same procedure described in [36]. Since the atmospheric neutrino simulation predicts a spectrum that increases linearly with energy, we model the uncertainty as modifying the slope of the spectrum. Once again, we define the ϵ parameter to parametrize the size of the distortion in units of standard deviations. The distorted PDF is then:

$$f'_{\mathrm{CC}\nu_e}(E;\epsilon) = N\left[1 + \frac{0.5\epsilon(E - 16\ \mathrm{MeV})}{74\ \mathrm{MeV}}\right] f_{\mathrm{CC}\nu_e}(E), \tag{4.8}$$

where N is a normalization constant. The prior for ϵ in this case is asymmetrical, since, as discussed in [36], a -2σ effect would be unphysical. Therefore, we use an asymmetric Gaussian function $G(\epsilon)$, and consider deviations between -1σ to $+3\sigma$. Using a Gaussian prior, the new likelihood is then:

$$\mathcal{L}' = \int_{-1}^{3} \mathcal{L}(\epsilon) G(\epsilon)\, \mathrm{d}\epsilon. \tag{4.9}$$

We apply this distortion to the spectrum in the Cherenkov signal region, for both neutron tagging regions. The distortions for these two regions are assumed to be fully correlated.

Relative normalization for NC backgrounds

The Cherenkov angle cut is associated with large uncertainties in the background efficiency of the NC backgrounds. Uncertainties in the nuclear interaction models translate to uncertainties in the Cherenkov pattern produced by such events. Therefore, we are left with a significant systematic uncertainty on the proportion of NC events leaking from the high-Cherenkov-angle sideband into the Cherenkov angle signal region. Modelling of this uncertainty also follows the method introduced by [36]. We consider a 1σ effect to be a 100% change in the number of NC events in the signal region. We model the impact of this change on the other regions by transferring the change in the number of events from the signal region to the sidebands, and reweight the associated PDFs accordingly:

$$f'^{\text{mid}}_{\text{NC}}(E;\epsilon) = (1+\epsilon) \times f^{\text{mid}}_{\text{NC}}(E) \tag{4.10}$$

$$f'^{\text{high}}_{\text{NC}}(E;\epsilon) = (1 - N_{\text{high}}\epsilon) \times f^{\text{high}}_{\text{NC}}(E) \tag{4.11}$$

$$f'^{\text{low}}_{\text{NC}}(E;\epsilon) = (1 - N_{\text{low}}\epsilon) \times f^{\text{low}}_{\text{NC}}(E). \tag{4.12}$$

Here, N_{high} and N_{low} are normalization constants, obtained from the integrals of the PDF regions, $A^{\text{region}}_{\text{j}}$:

$$A^{\text{region}}_{\text{j}} = \int_E f^{\text{region}}_{\text{j}}(E)dE \tag{4.13}$$

$$N_{\text{high}} = \frac{A^{\text{mid}}_{\text{NC}} A^{\text{high}}_{\text{NC}}}{(A^{\text{high}}_{\text{NC}})^2 + (A^{\text{low}}_{\text{NC}})^2} \tag{4.14}$$

$$N_{\text{low}} = \frac{A^{\text{mid}}_{\text{NC}} A^{\text{low}}_{\text{NC}}}{(A^{\text{high}}_{\text{NC}})^2 + (A^{\text{low}}_{\text{NC}})^2}. \tag{4.15}$$

We finally convolve the likelihood over a Gaussian prior, as in equation 4.9. We apply the same spectral distortions to each neutron tagging region, assuming that the distortions in these regions are fully correlated.

Neutron multiplicity

Unlike the analysis in Ref. [36], our PDFs are sensitive to the number of neutrons in the signal and background. This number, also known as the *neutron multiplicity*, is subject to uncertainties which need to be accounted for in a neutron tagging analysis. In particular, the atmospheric neutrino background simulation is associated with large uncertainties related to the number of neutrons produced in the nuclear processes that accompany the primary event. Calculation of this uncertainty is treated in more detail in Ref. [49]. As described there, we set the systematic uncertainty on the neutron multiplicity to a conservative value of 40%. This particular uncertainty comes from a T2K measurement detailed in Ref. [70], and only depends on the momentum transfer q^2. As such, it is fully correlated between atmospheric neutrino background categories. Therefore, the same value is used across these categories. For the DSNB signal, while there is no uncertainty on the number of neutrons, the neutron tagging

cut contributes an absolute uncertainty of 6%. In our PDFs, this uncertainty relates to the relative proportion of background that is categorized as IBD-like or non-IBD-like. We therefore parametrize the distortion of the PDFs as a change in normalization between neutron tagging regions:

$$f'^{\text{IBD}}_j(E;\epsilon) = (1+\alpha_j\epsilon) \times f^{\text{IBD}}_j(E) \tag{4.16}$$

$$f'^{\text{non-IBD}}_j(E;\epsilon) = (1-\alpha_j N\epsilon) \times f^{\text{non-IBD}}_j(E) \tag{4.17}$$

$$N = \frac{A^{\text{IBD}}_j}{A^{\text{non-IBD}}_j} \tag{4.18}$$

where α_j is the relative 1σ systematic uncertainty for PDF j, and N is a normalization constant. This time, we can take the prior probability of ϵ to follow a standard Gaussian, so the new likelihood will be calculated as in 4.7.

Spallation spectral shape

As mentioned above, since the treatment of the spallation backgrounds was not the object of this work, the modelling of uncertainty associated with spallation spectrum—which is used in the final spectral fit—is sketched here for completeness. The main uncertainty related to the spallation spectrum in Equation 4.1 is in the relative contributions of each isotope to the overall spectrum. Since the ^{9}C isotope has a shorter half-life than the other ^{8}B and ^{8}Li, and is produced closer to the muon track, it is easier to remove, and will be subdominant. Therefore, the range of possible values for the scaling n_0 is considered to be $[0, 0.5]$, while it n_1 is taken to take any value in $[0, 1]$. Modifying the isotope composition of the spallation background has the effect of changing the steepness of the overall spectrum. The central value for the spectral PDF is taken as the average between the steepest and least-steep spectrum, each of which is taken as the 1σ deviation, with the deviation following a Gaussian prior distribution. More details about this uncertainty, are given to Ref. [49].

Combined background shape uncertainty

Combining the distortions applied to the spectra as described above, the shapes of PDFs are allowed to vary according to five parameters: ϵ_E, $\epsilon_{\text{CC}\nu_e}$, ϵ_{NC}, ϵ_{n}, and ϵ_{spall}. The modification to the likelihood function $\mathcal{L}$ from Equation 4.3 can then be summarized as

$$\begin{aligned}\mathcal{L}'(\vec{N}) = \int & \frac{e^{-\epsilon_E^2/2}}{\sqrt{\pi}} G(\epsilon_{\text{CC}\nu_e}) G(\epsilon_{\text{NC}}) \frac{e^{-\epsilon_{\text{n}}^2/2}}{\sqrt{\pi}} \frac{e^{-\epsilon_{\text{spall}}^2/2}}{\sqrt{\pi}} \\ & \times \mathcal{L}(\vec{N}; \epsilon_E, \epsilon_{\text{CC}\nu_e}, \epsilon_{\text{NC}}, \epsilon_{\text{n}}, \epsilon_{\text{spall}}) \\ & \times d\epsilon_E \, d\epsilon_{\text{CC}\nu_e} \, d\epsilon_{\text{NC}} \, d\epsilon_{\text{n}} \, d\epsilon_{\text{spall}},\end{aligned} \tag{4.19}$$

where the integral sign implies the multidimensional integration over the distortion parameters, and G is the weighted Gaussian from above.

Signal efficiency

The systematic uncertainty on the signal efficiency enters into the computation of the DSNB interaction rate from the observed number of signal events in Equation 4.4. Therefore, the systematics for all the relevant reduction efficiencies described in the previous sections are accounted for here. However, note that the Cherenkov

SK phase	I	II	III	IV
Fiducial volume	1.3%	1.1%	1.0%	1.5%
Cross-section	1.0%	1.0%	1.0%	1.0%
Livetime	0.1%	0.1%	0.1%	0.1%

TABLE 4.2: Uncertainties associated with the fiducial volume cut, the Strumia-Vissani IBD cross-section calculation, and the live time calculation for SK-I to IV. The uncertainties for SK-I,II,III have been taken from the SK solar analyses [71, 69, 72]. For SK-IV the uncertainties have been obtained from the antineutrino Monte-Carlo simulation used for the DSNB analysis.

angle cut and the neutron tagging cut only affect the distribution of the signal across the six fitting regions, and thus need not be considered. Therefore, the overall signal efficiency ε_S is the product of the efficiencies from the preliminary cuts, spallation cuts, and positron identification cuts (except the Cherenkov angle cut). The associated uncertainties are added in quadrature. In addition, the final flux calculation in Equation 4.5 relies on the value of σ_{IBD}, on the fiducial volume cuts, and on the SK-IV livetime, so the associated uncertainties must also be considered. All of these systematic uncertainties are added in quadrature, to yield the overall signal efficiency systematic uncertainty σ_S:

$$\sigma_S^2 = \sigma_{\text{cuts}}^2 + \sigma_{\text{x-sec}}^2 + \sigma_{\text{FV}}^2 + \sigma_{\text{livetime}}^2. \tag{4.20}$$

These uncertainties are summarized in Table 4.2. To model the impact of this systematic uncertainty, we use a similar approach to the spectral shape uncertainties, by distorting the likelihood function according to a prior probability. First, the modified likelihood $\mathcal{L}'(\vec{N})$ from Equation 4.19 is marginalized, as in Equation 4.6 to obtain the signal number likelihood $\mathcal{L}'_S(N_S)$. Rewriting in terms of the total number of IBD interactions $R = N_S/\varepsilon_S$ a Gaussian prior centered at ε_S with standard deviation σ_S is applied to the marginal likelihood:

$$\mathcal{L}''_S(R) = \int_0^1 \mathcal{L}'_S(\varepsilon_S R)\, e^{-(\varepsilon-\varepsilon_S)^2/(2\sigma_S^2)}\, d\varepsilon, \tag{4.21}$$

where the normalization, such that the integral between 0 and 1 of the prior is equal to 1, is implied.

4.4 Combination with SK-I,II,III

The fitting method presented here differs from the one used in Ref. [36] in a few ways, the most salient of which is the inclusion and treatment of the neutron tagging cut. Nevertheless, the statistical approach is similar, and, crucially, both analyses are dominated by statistical uncertainties. The results from the two analyses can therefore be incorporated into a combined observation—for the entire SK-I-II-III-IV period—with a higher statistical power. To use data from the previous SK phases, we use the IBD simulation produced in Ref. [36], applying the same cut criteria, and use the same functions to model the atmospheric neutrino backgrounds. We then reanalyze the SK-I, II, and III, with a few changes to the original analysis:

- The *remaining spallation* background, not originally included, was introduced, along with the associated systematic uncertainty.

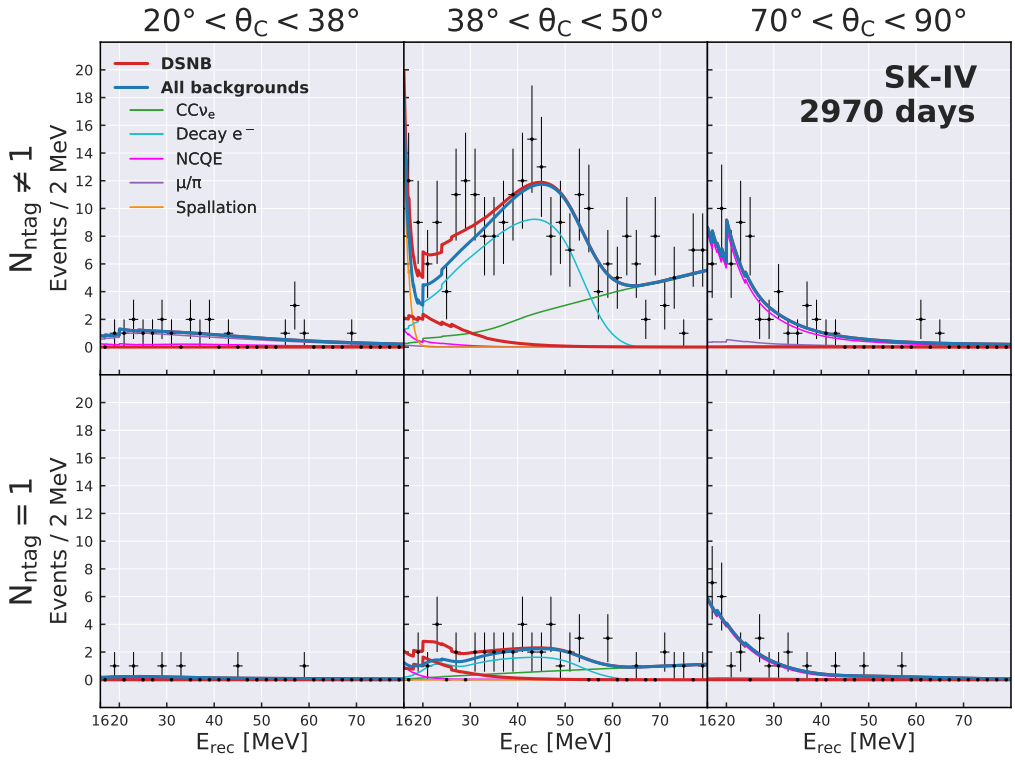

FIGURE 4.3: Best-fit signal and background spectra (solid lines) for the SK-IV period, assuming the DSNB flux predicted by the Horiuchi+09 model [20], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

- A larger variety of DSNB flux models is studied, notably including several predictions that were only formulated over the past few years and are more nuanced in their modelling.
- The calculation of the IBD positron energy was updated. While the original analysis applied a 1.3 MeV shift to the neutrino energy, this analysis uses a more accurate function based on the computation in Ref. [55].

Through this repeated analysis, we obtained three additional marginal likelihoods (as in Eq. 4.21), one per SK period, that we can sum together with the SK-IV likelihood to yield a combined likelihood:

$$\mathcal{L}''^{\text{combined}}_{S} = \mathcal{L}''^{\text{SKI}}_{S} + \mathcal{L}''^{\text{SKII}}_{S} + \mathcal{L}''^{\text{SKIII}}_{S} + \mathcal{L}''^{\text{SKVI}}_{S}. \tag{4.22}$$

4.5 Results

The analysis method above was applied to the full SK-IV period dataset. The combination with the previous eras brings the total exposure to 22.5×5823 kton·day, using a wide range of DSNB predictions to fit the signal. An example of such a fit is

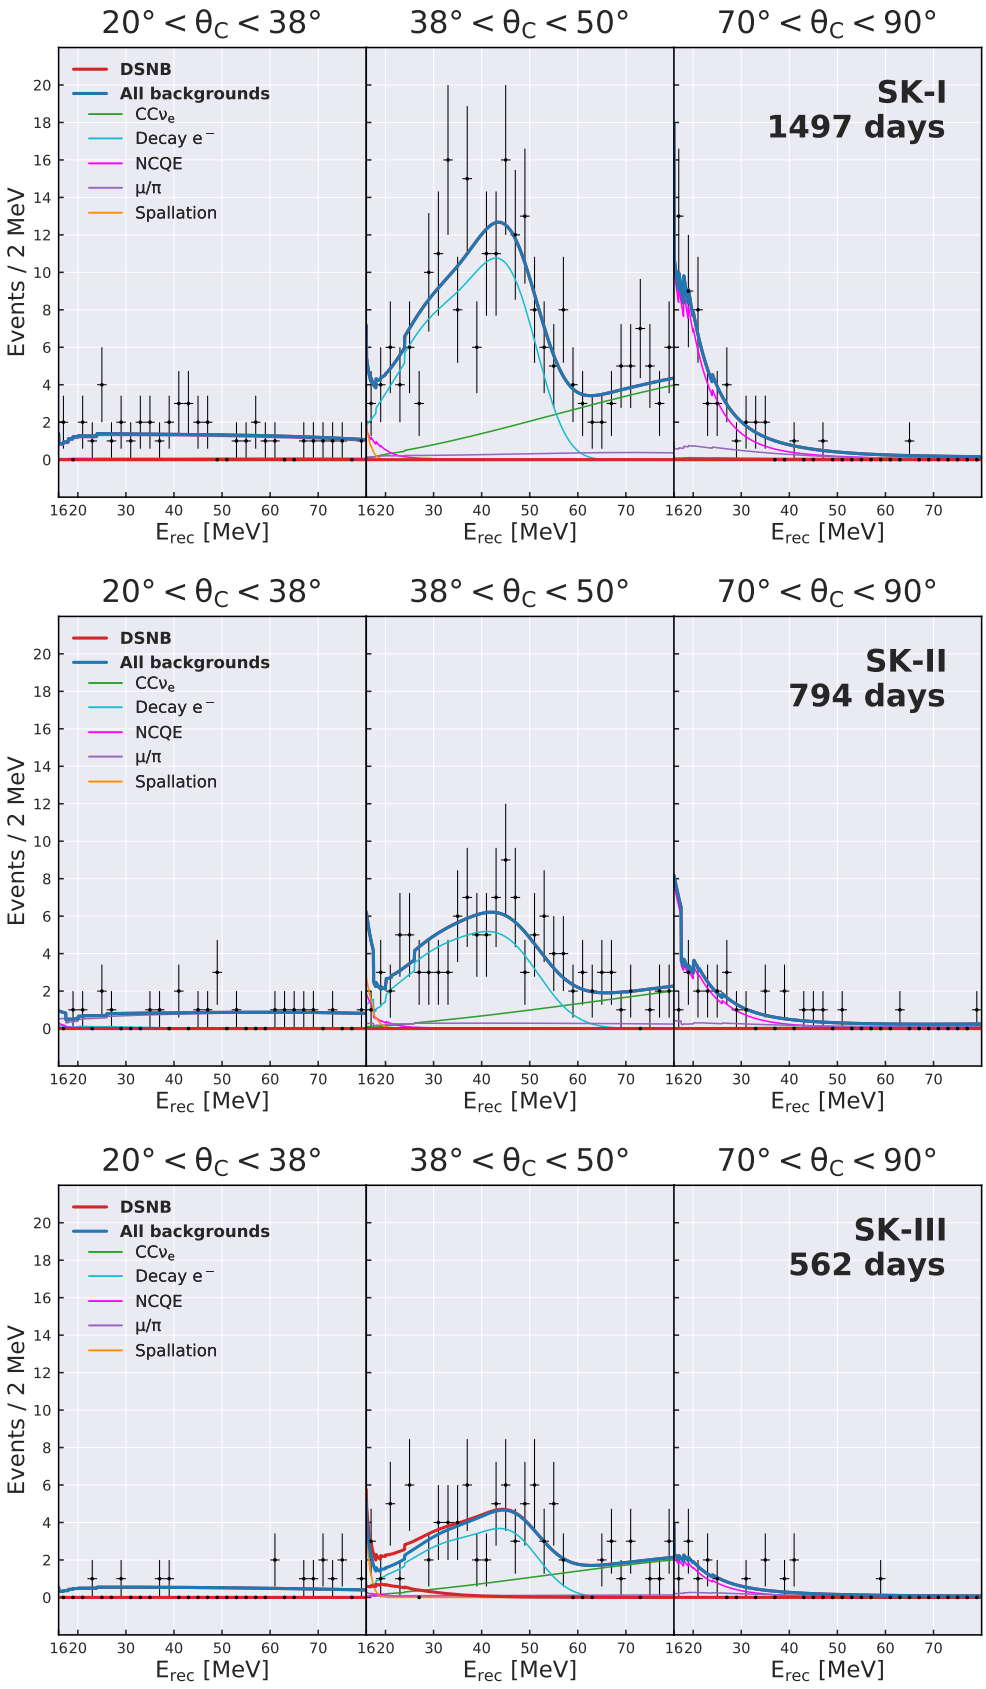

FIGURE 4.4: Best-fit signal and background spectra (solid lines) for the SK-I (top), SK-II (middle), and SK-III (bottom) periods, assuming the DSNB flux predicted by the Horiuchi+09 model [20], along with the energy spectrum of the remaining events in data after reduction. The three parameters space regions are the Cherenkov angle signal region and the two sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

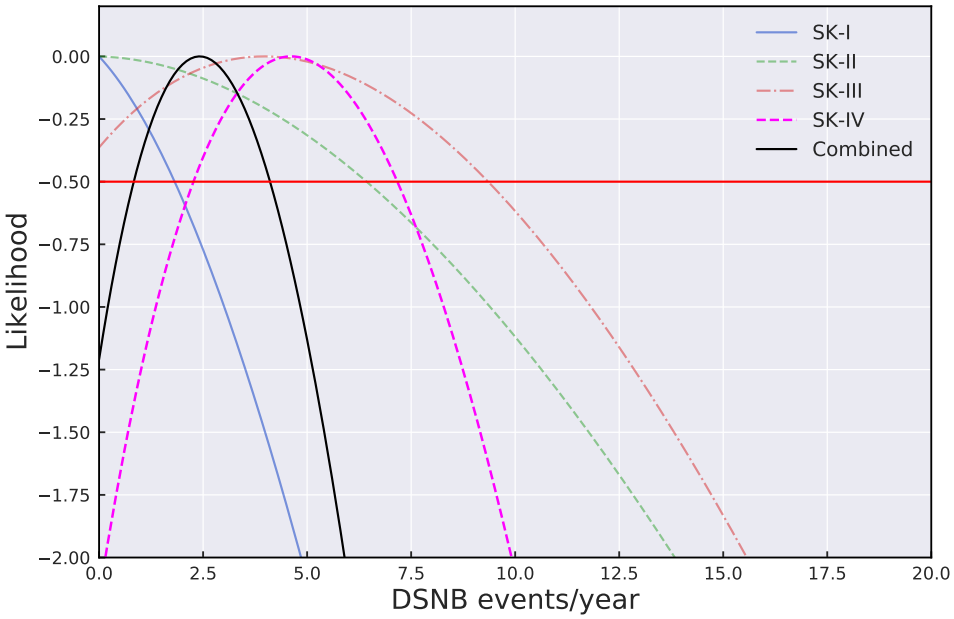

FIGURE 4.5: Likelihoods associated with phases I to IV of SK, as well as the combined likelihood. As mentioned in the main text, we can neglect correlations between background systematics when combining all phases of SK.

visualized in Fig. 4.3 (for the SK-IV fit) and Fig. 4.4 (for the SK-I, II, and III fits), where the Horiuchi+09 DSNB model [20] is used. The Cherenkov angle signal region, is, as expected, dominated by the decay electron spectrum, until about 60 MeV, where CC-ν_e backgrounds take over. The fit in the low and high Cherenkov angle sidebands also behave as expected, and are dominated, respectively, by μ/π and NCQE events. Due to the relatively low neutron tagging efficiency, a greater proportion of the best-fit signal rate is found to be in the non-IBD region. Nevertheless, the signal that makes into the IBD region also enjoys a lower background contamination. In total, the SK-IV fit using the Horiuchi+09 model found 170 decay electron events, 110 CC-interaction ν_e events, 88 NCQE interaction events, 28 μ/π events, and 9 spallation events, for the energy range considered. The best-fit rate for this model is then an excess of $4.6^{+2.5}_{-2.3}$ IBD/22.5 kton/year. The corresponding best-fit DSNB flux is $2.6^{+1.4}_{-1.3}$ $\overline{\nu}_e$/cm^2/year. Note that these values refer to the DSNB rate and flux above a neutrino energy $E_\nu > 17.3$ MeV, reflecting the energy threshold of the analysis. This choice also makes for a straightforward comparison with the results from Ref. [36], which uses the same convention.

While our re-analysis of the previous SK periods differs in some respects to the original analysis in Ref. [36] similar results are obtained, with discrepancies well within uncertainties. In particular, the most significant difference—the inclusion of possible contributions from spallation backgrounds—found no more than two events for any of the previous SK periods, therefore resulting in a limited impact on the best-fit DSNB flux. As with the previous analysis, the SK-I and SK-II periods—whose combined livetime approaches that of SK-IV—are witness to a strong down-fluctuation (assuming the presence of our signal), leading to a best-fit rate of 0 for these searches. The SK-III period sees an excess in its data samples, though its shorter livetime confers them a less important statistical weight. The combination of the four periods

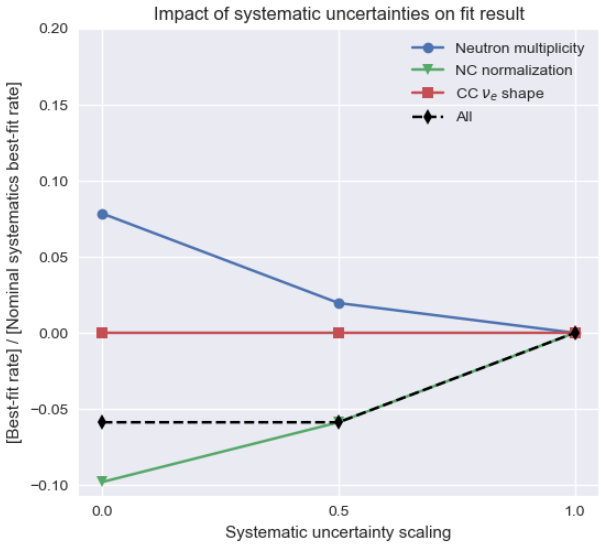

FIGURE 4.6: Relative variation of the best-fit flux for the Ando LMA model as a function of the relative variation of the atmospheric neutrino background spectral shape uncertainties.

therefore leads to a lower best-fit DSNB rate and flux, $2.4^{+1.6}_{-1.7}$ IBD/22.5 kton/year and $1.3^{+0.9}_{-0.9}$ $\overline{\nu}_e$/cm^2/year, respectively.

The marginal likelihood functions, for all SK periods and for the combined analysis, are plotted in Figure 4.5 as a function of the best-fit yearly DSNB interaction rate for the Horiuchi+09 model. The combination of the likelihoods relies on the assumption that the systematic uncertainties—which correlate across SK phases—are subdominant with respect to the statistical uncertainty of the data. To evaluate the impact of the systematic uncertainties on the result, the fit was repeated with a rescaling of most important of these uncertainties, namely the spectral shape uncertainties. The effects of modulating the atmospheric background shape uncertainties are shown in Figure 4.6. While the uncertainty on the CC-ν_e spectral shape is significant, its impact is very limited, as it is subdominant in the key energy region. While the NC and neutron multiplicity uncertainties have a measurable impact, their inclusion modifies the final best-fit only by about 6%. Inclusion of the spallation background systematics causes only a 3% shift in the results, for an overall impact of less than 7%. The effect of correlated systematic uncertainties is therefore neglected in the addition of the likelihoods.

The likelihood functions are then used to extract fit uncertainties and place upper limits on the DSNB interaction rates and fluxes, for each model considered. In particular, the 90% confidence level (C.L.) upper limit on the flux $\Phi^{90\%}$ is defined as

$$\frac{1}{N}\int_0^{\Phi^{90\%}} \frac{\mathcal{L}''_S(\Phi)}{\mathcal{L}''_S(\Phi^{\text{best}})}\, d\Phi = 0.9 \tag{4.23}$$

$$N = \int_0^{\infty} \frac{\mathcal{L}''_S(\Phi)}{\mathcal{L}''_S(\Phi^{\text{best}})}\, d\Phi. \tag{4.24}$$

For the SK-IV period, the 90% C.L. upper limits on the DSNB flux and IBD interaction rate are 4.6 $\overline{\nu}_e$/cm^2/year and 8.2 IBD/22.5 kton/year, respectively, once again for the Horiuchi+09 spectrum. This model is one of the most optimistic among the

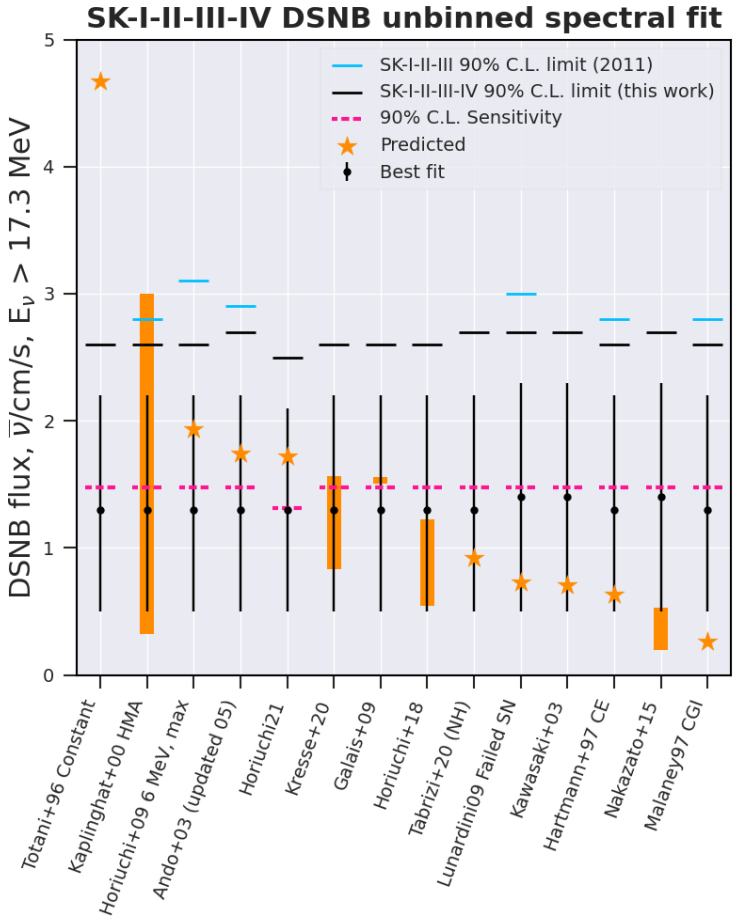

FIGURE 4.7: Spectral fit results for the models considered in this work. The black dots and vertical lines are the best-fit DSNB model fluxes and uncertainties, while the solid black horizontal lines are the 90% C.L. upper limits on the flux. The dashed magenta lines give the 90% C.L. sensitivity of the search for each model, here defined as 1.64 times the 1σ error on the fit. These values are compared to the SK-I-II-III results from the analysis in Ref. [36] (blue horizontal lines), as well as to the theoretical flux predictions of each model (orange stars and bars). A single star denotes a singular flux prediction, while a bar denotes a range of flux predictions. Since a larger number of models were considered with respect to Ref. [36], comparison of the limits is only possible for a few of them.

ones considered, predicting a flux of 1.94 $\overline{\nu}_e/\text{cm}^2/\text{year}$ for $E_\nu > 17.3$MeV. The upper limits on the flux and interaction rate obtained with the combined analysis are 2.6 $\overline{\nu}_e/\text{cm}^2/\text{year}$ and 4.8 IBD/22.5 kton/year.

The best-fit and 90% C.L. upper limit results for the models studied in this analysis are listed in Table C.1, along with the predicted integrated fluxes, and are visualized in Fig. 4.7. A similar plot for the corresponding best-fit event rates at SK is given in Fig. C.1. Overall, the results are fairly stable across models, demonstrating that the sensitivity to spectral features of the signal is limited. This is largely due to the statistical uncertainty, which, even with the substantial increase in exposure brought by SK-IV, is still a significant limitation. For all models, an excess is observed around a flux of $1.4^{+0.9}_{-0.9}$ $\overline{\nu}_e/\text{cm}^2/\text{year}$, though this excess has low statistical significance ($< 2\sigma$). The 90% C.L. exclusion of the optimistic predictions by the Totani+95 [73] and Kaplinghat+00 [74] models, already excluded in Ref. [36], is confirmed by this study. The upper flux limits placed in this analysis are tighter than those placed in Ref. [36], though this improvement is limited by the fact that the previous analysis

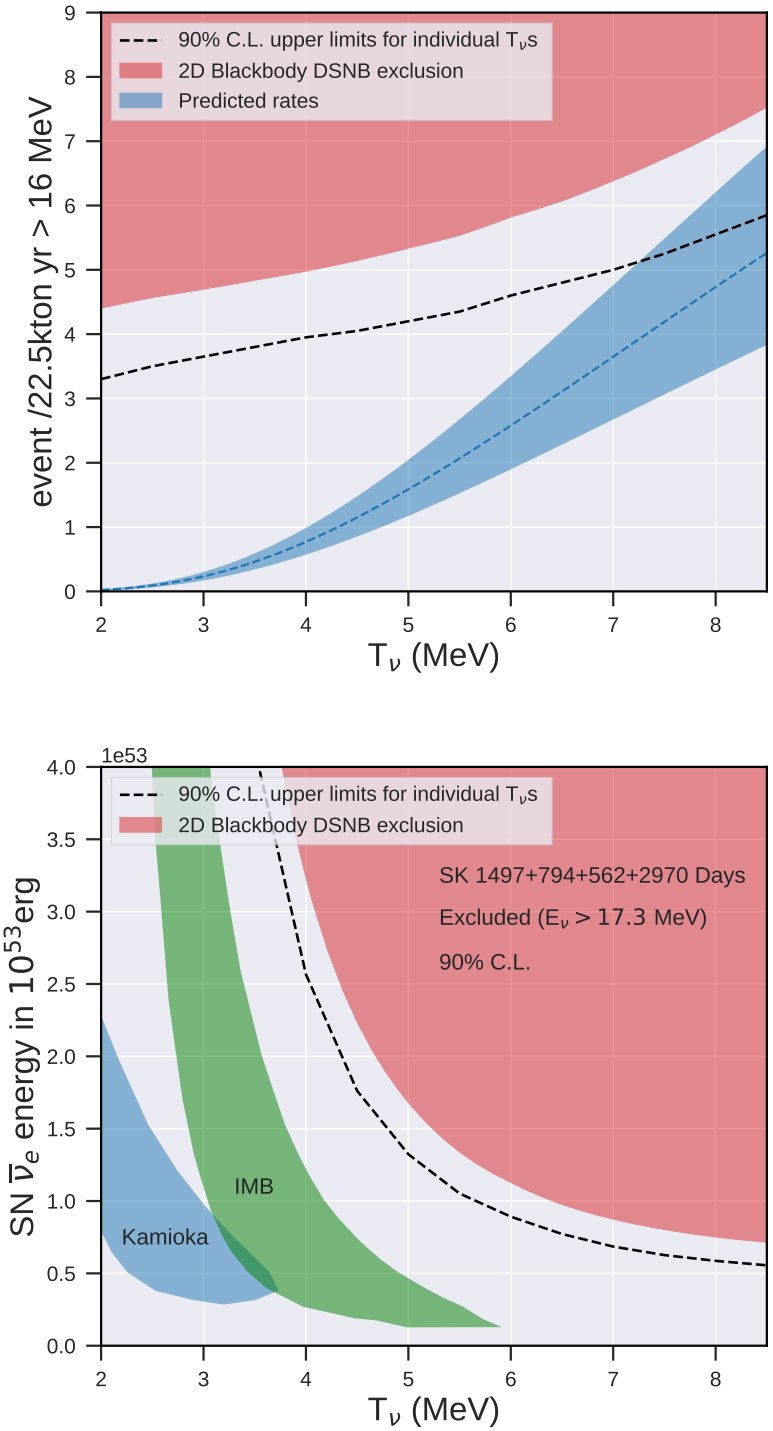

FIGURE 4.8: Results of spectral fits on a simple blackbody DSNB model parametrized by the neutrino temperature T_ν. Top: 90% C.L. excluded region for the DSNB rate for $E_\nu > 17.3$ MeV and the effective neutrino temperature, (red surface). To guide the eye, we also show the one-dimensional exclusion limit for individual neutrino temperatures. The blue band shows the predictions for the blackbody models, for the low, fiducial, and high star formation rates used in Ref. [20]. Bottom: 90% C.L. excluded region for the total neutrino energy emitted during a supernova and the effective neutrino temperature, (red surface). To guide the eye, we also show the one-dimensional exclusion limit for individual neutrino temperatures. We also show the 90% C.L. contours corresponding to the range of parameters allowed for SN1987A.

found a much lower best-fit signal, so our excess weakens our limits in comparison. Nevertheless, our sensitivity to the signal is, as expected, significantly improved, and is consistent with the roughly doubled exposure, as well as the higher sensitivity in neutron-tagging fit region, despite the inclusion of a new source of systematic uncertainty. It is, as of yet, the tightest limit on the DSNB flux. Further, we can now claim a 90% sensitivity—here defined by multiplying the 1σ uncertainty by 1.64—that is comparable to several modern DSNB flux predictions. This sensitivity to the DSNB is the best to date, and is comparable in particular to the Ando+03 [21], Horiuchi+09 ($T_\nu = 6$ MeV) [20], and Galais+10 [22] models, as well as to the high-end of the Kresse+21 [25] model. The relatively high flux predicted by these models is owed to the fact that they use the highest cosmic star formation rate that is allowed by modern observations.

In principle, we can go beyond the discrete DSNB models discussed here and perform fits on smoothly-varying parametric models, to derive conclusions about the spectral features of the DSNB signal and exclude regions of the parameter space. Indeed, this is one of the advantages of a model-dependent approach such as this one. In practice, however, the large statistical uncertainty on the fit is largely going to limit our ability to probe fine parameters in a DSNB model such as the proportion of black-hole-forming supernovas. In this work, we probe a simple Fermi-Dirac spectrum parametrized by the temperature T_ν and its luminosity. A range of neutrino temperatures, from 2 MeV to 8.5 MeV are tested. Each model is fitted as before, and the results are shown in Figure 4.8. Here, we extract a two-dimensional 90% C.L. exclusion on the neutrino temperature and on the DSNB rate, for the combined SK-I-II-III-IV spectral analysis. The one-dimensional limit for each temperature is also given, and is consistent with the rate limits found in the discrete models studied above. These limits are also compared with the expectation from a simple Fermi-Dirac DSNB model, parametrized by the neutrino temperature, as well as by the range of stellar formation rates given in Ref. [20]. The same limits can also be converted into limits on the equivalent total neutrino energy for a single supernova, which are also plotted in Figure 4.8. Note that since the DSNB is an integrated flux, this limit represents the total neutrino energy for the "average" supernova. In the plot, we compare this exclusion zone with the allowed parameter space obtained by the Kamiokande-II and IMB experiments, following the observation of SN1987A [75]. This particular supernova is thought to be remarkably low in temperature, and not representative of a typical supernova. Nevertheless, future DSNB measurements could begin probing the region allowed by IMB.

5
Prospects for the DSNB with SK-Gd

After placing limits on the DSNB with SK-IV, we turn our gaze to the future of the Super-Kamiokande experiment. What will the Gd era bring? To what extent will our sensitivity to the DSNB improve, and how likely is a positive DSNB detection in SK-Gd's lifetime? The aim of this chapter will be to begin answering these questions. In particular, the scope of this investigation will be to consider a DSNB analysis in the spirit of the one presented in the previous sections. Its future performance will be projected into the future mainly by estimating the impact of the detector's improved neutron tagging capabilities.

5.1 Signal modelling

Starting with SK-VI, we need to modify the coincidence signal picture described in Section 3.1 with the addition of Gd in the tank. The main change to the detection channel is, of course, the neutron capture signature. Depending on the Gd concentration, as discussed in Section 2.3, a percentage of the neutrons will be captured by the Gd, producing a photon cascade of roughly 8 MeV total energy. A summary of the isotopes present, with their neutron capture cross-sections, is given in Table 5.1. It is clear that two different Gadolinium isotopes—$^{155}\mathrm{Gd}$ and $^{157}\mathrm{Gd}$—boast enormous neutron capture cross-sections and represent the vast majority of the captures on Gd. For both of these isotopes, the cascade is the result of a transition from an excited state:

$$n + {}^{155}\mathrm{Gd} \rightarrow {}^{156}\mathrm{Gd}^* \rightarrow {}^{156}\mathrm{Gd} + N\gamma \text{ (8.536 MeV total)} \tag{5.1}$$

$$n + {}^{157}\mathrm{Gd} \rightarrow {}^{158}\mathrm{Gd}^* \rightarrow {}^{158}\mathrm{Gd} + N\gamma \text{ (7.937 MeV total)}, \tag{5.2}$$

where the resulting cascade is composed of a variable number N of photons, with an average multiplicity around four. An example of such a cascade is sketched in Figure 5.1. For cascades with several photons, the energy of a single photon in the cascade can fall below the Cherenkov threshold for detection of the associated Compton electron. The visible energy from the cascade will therefore be, on average, lower than the total photon energy. This effect is highly dependent on the photon emission model, so a precise account of this process is essential to develop a successful simulation.

To this end, a new simulation framework was deemed necessary. The SKDetSim package described in Section 3.1 is partially based on the Geant3 simulation framework, which is no longer being updated. In particular, the neutron interaction models are no longer up to date. While this shortcoming has little impact on a pure-water analysis, in which the neutron interaction and photon emission have long been well-understood, it could seriously affect a Gd-based simulation. For this reason, the *SKG4*

Isotope	Abundance (%)	Cross-section (barn)	γ energy (MeV)
^{1}H	99.99	0.333	2.224
^{16}O	99.76	0.0002	4.143
^{152}Gd	0.20	735	6.247
^{154}Gd	2.18	85	6.435
^{155}Gd	14.88	60900	8.536
^{156}Gd	20.47	1.8	6.360
^{157}Gd	15.65	254000	7.937
^{158}Gd	24.84	2.2	5.943
^{160}Gd	21.86	1.4	5.635
^{32}S	94.85	0.548	8.641
^{34}S	4.37	0.235	6.986
^{16}O	99.76	0.0002	4.143

TABLE 5.1: Natural abundances [76], thermal neutron capture cross-sections [77], and (total) photon energies emitted after neutron capture [78], for the main isotopes found in the SK-Gd tank. Top panel: main isotopes in pure water (H_2O). Bottom panel: main isotopes in the gadolinium sulfate compound dissolved in the tank ($Gd_2(SO_4)_3$). The relative contribution of each depends on the Gd concentration. The isotopes contributing to an appreciable degree to the neutron signal are highlighted.

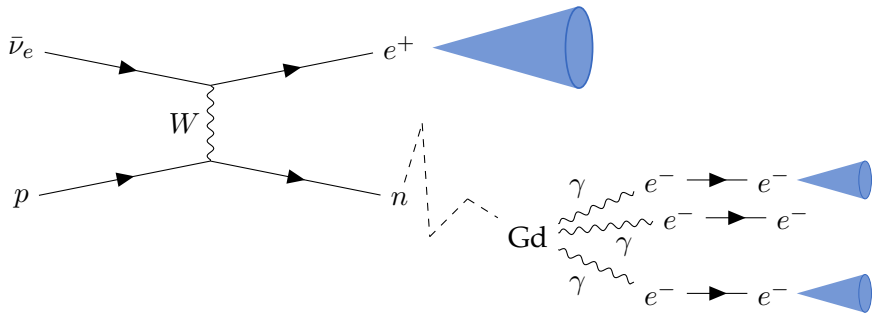

FIGURE 5.1: The main detection channel for DSNB events at Super-K after doping the water with Gadolinium, in the case of neutron capture by a Gd atom. This capture produces a cascade of photons, whose total energy sums up to ~ 8 MeV. The detectable neutron signal will then come from the recoil electrons from the Compton scattering of the photons above Cherenkov threshold. Here, the capture is followed by a three-photon cascade, of which only two are above threshold. The resulting Cherenkov signal, while still significantly weaker than the positron signal, is much stronger than the neutron signal in pure water. The characteristic capture timescale on Gd is also much shorter, $\tau \sim 30\mu$s, though still long enough to clearly separate the prompt and delayed events, on average.

package was developed by the SK collaboration [79], entirely based on Geant4, allowing for a more accurate account of the neutron scattering, thermalization, capture and other processes associated with neutrons in the Gd-doped water. This package was used to simulate the detector response of SK-Gd. Particular attention was given to the photon emission model used by SKG4, for which the ANNRI-Gd model [80, 81] was chosen. This model enjoys excellent experimental agreement, and was developed with the detection of neutrons in SK-Gd as one of the main physics goals.

The development of simulated signal samples proceeded in a similar fashion as for the simulation of IBD events for SK-IV, described in 3.1. The same assumptions about the IBD kinematics are made, and the same reweighting scheme is used to define the energy spectrum of the IBD sample. The initial states of the IBD events are used as input to SKG4. The parameters of SKG4 were tuned to match the condition of the detector during the SK-VI period. In particular, the Gd concentration in the simulation was set to 0.013 %, following the most recent measurements at the time of the sample's production. The inclusion of low-energy detector background also followed a similar strategy. Since a precise characterization of the low-energy background is crucial for neutron tagging and no complete simulation is available, it is taken directly from data, using randomly-triggered events taken from SK-VI data runs. Each signal event is overlapped with this background data, allowing us to model the accidental coincidences, and contaminating the true simulated neutron captures themselves.

5.2 Neutron tagging with Gadolinium

The spirit of the neutron tagging algorithm will still follow the general procedure of the SK-IV neutron tagging, making the same assumptions about the neutron capture vertex, calculating discriminating observables on $t-$TOF clusters, and applying a BDT algorithm to these observables. Of course, this is not the only possible approach. The cleaner Gd-capture signal certainly opens the door for a wider range of strategies. For example, direct triggering on the Gd neutron capture is now possible, namely via the SLE trigger, and this option is under study in the collaboration. Further, the neutron candidate preselection and the observables could be modified to take into account the differences between the 2.2 MeV H-capture photon and the Gd-capture cascade. Nevertheless, the algorithm was not drastically altered here. The first reason is simply that, as will be shown, such an algorithm already gave very good results, more than good enough to judge the scope of the improvement in performance. Secondly, although the neutron captures on Gd are the stars of the show in SK-Gd, a significant portion of the captures still occurs on Hydrogen (half of them, in SK-VI). Since H-captures are much harder to see, having an algorithm attuned to them is not such a disadvantage. The cleaner Gd-captures, on the other hand, are more forgiving, and do not require the same degree of algorithmic tuning. Incidentally, using a similar algorithm, this study will show the direct impact of the Gd doping, with little in the way of confounding factors.

5.2.1 Candidate preselection and characterization

Even in the context of a similar neutron tagging algorithm, the introduction of neutron captures on Gd meant that some basic changes to the candidate selection were in order. First, the higher PMT yield of the neutron capture signal requires a looser cut on the maximum number of hits in a 200-ns window N_{200}^{max}. This cut is used

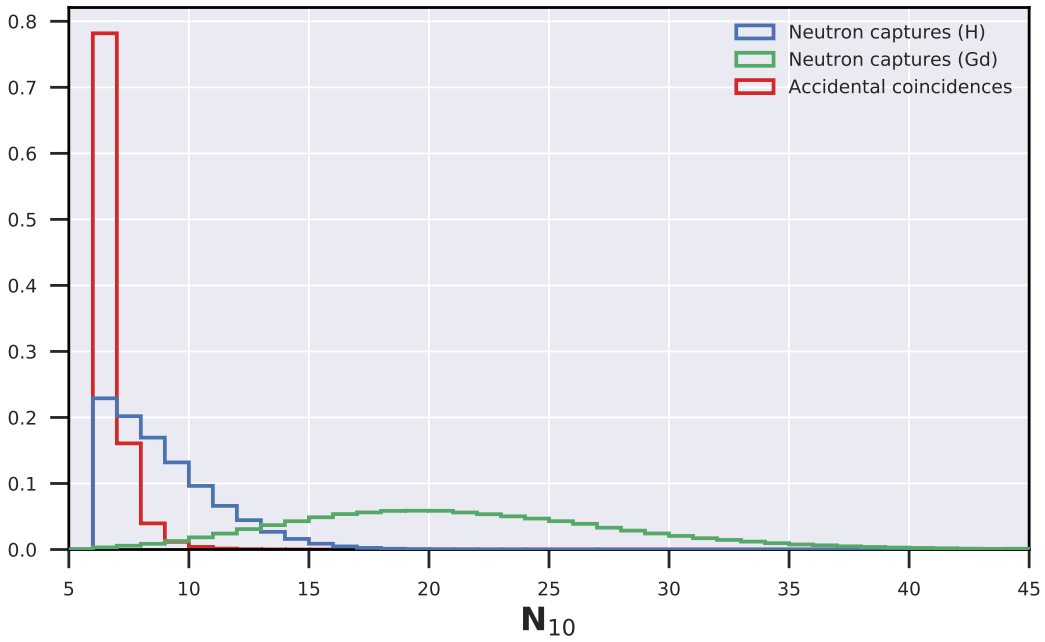

FIGURE 5.2: Probability density function of N_{10}, for accidental coincidences (in red) and true neutrons captures (in blue). The amount of background contamination rises sharply with decreasing N_{10}.

to remove unusually-high PMT activity in the neutron search window, and should have near-perfect efficiency. Since the typical PMT activity of the neutron signal is now sensibly higher, a higher cut is needed. In particular, a cut of $N_{200}^{\text{max}} > 100$ was used. Secondly, since neutrons get captured more quickly by Gd than by H, special care is needed when selecting the start time of the neutron search. Using the same starting time as with the search in pure water, 23 μs after the prompt event, would result in an unacceptable loss of Gd captures, occurring with characteristic time $\tau \sim 30$ μs. Therefore, the neutron search window was widened to start at just 1 μs after the prompt event, extending to the end of the AFT window, at 535 μs. In turn, this changed motivated the removal of the correlated dark noise cut, which would have required a restriction of this time window at the edges. The removal of this cut causes single-handedly a roughly three-fold increase in the rate of accidental coincidences at this stage. However, the brighter neutron captures in SK-Gd make this trade-off more than acceptable.

The procedure outlined in Section 3.5 is then carried out as before. Neutron capture candidates are searched for by leveraging their proximity to the prompt vertex. Clusters of hits are searched for by looking for clusters in a 10-ns window in time-TOF from the prompt event vertex. For an overview of the variables mentioned in this section, refer back to Table B.1. The N_{10} variable (displayed in Figure 5.2) is calculated in the same way as earlier, and a cut of $N_{10} > 5$ is used. Already at this stage, the difference between pure and Gd-doped water couldn't be more evident. While the preselection was one of the main sources of the low neutron tagging efficiency in SK-IV, we now observe an excellent efficiency for neutron captures occurring on Gd, whose mean value of N_{10} hovers around 20. In fact, the main cause of inefficiency for these neutron captures comes from captures occurring before the start of the search window 1 μs, leading to a 3.3% loss. For those captures that make it into the search window, the efficiency of the preselection is nearly perfect,

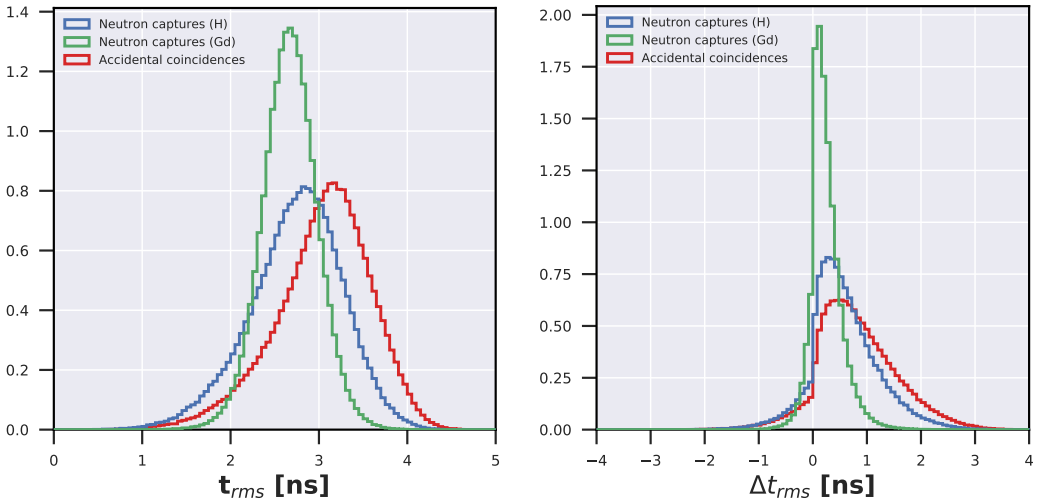

FIGURE 5.3: Probability density function of t_{RMS} (left) and Δt_{RMS} (right), in nanoseconds, for accidental coincidences (red) , true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). PMT hits from true neutron captures display lower time dispersion, and see this dispersion reduced to a greater extent when a better vertex is found. Note that, due to its definition, Δt_{RMS} corresponds to the *decrease* in t_{RMS} after the vertex improvement.

99.5%. The overall signal efficiency of the preselection scheme is then 96.2% for Gd neutron captures. On the other hand, it is 54.5% for H neutron captures. The total signal efficiency of the neutron capture candidate preselection for the whole sample is then 75.3%. After preselection, the expected number of accidental coincidences in the neutron search window of a prompt event is a whopping 54.0. This is a notable increase—roughly five-fold—from the rate observed with the SK-IV algorithm. As mentioned above, the removal of the continuous dark noise cut accounts for only a three-fold increase. The remaining increase is consistent with a 1.5-fold rise in the dark noise rate in the PMTs, which was measured starting with SK-V. The exact causes and mechanism for this rise are still under investigation by the collaboration. Nevertheless, since this effect was found in SK-V—before any Gd sulfate was added—the Gd itself was ruled out as a potential cause. Though further study is required to pinpoint the issue, this unknown does not affect our strategy. As our account of the accidental coincidences comes from SK-VI data, the effect of a higher dark noise rate is inherently taken into account.

Neutron capture vertex determination

The time-of-flight subtraction procedure, when applied to Gd neutron captures, yields clusters of hits that are more tightly grouped together in t-TOF. This effect can be observed in the distribution of t_{RMS} and Δt_{RMS}, plotted in Figure 5.3, and can be explained by the fact that hit clusters from true captures can be contaminated with background hits, uncorrelated from the prompt vertex, thereby increasing the t-TOF dispersion. This effect is significantly diminished in Gd captures, where the higher hit count means that any background hits within a signal cluster will have a smaller impact.

The neutron capture vertex found by searching for the point minimizing the t-TOF dispersion will therefore be closer to the prompt vertex, as is visible in the distribution of the distance (plotted Figure 5.4, along with ΔN_{10}). A Gd capture is found by our

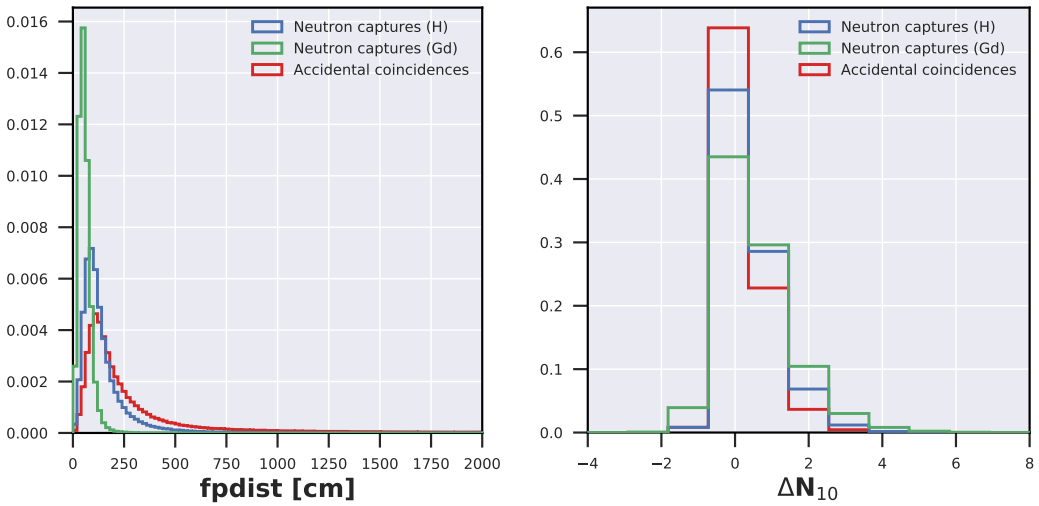

FIGURE 5.4: Probability density function of $\mathrm{fp}_{\mathrm{dist}}$, in cm (left) and ΔN_{10} (right), for accidental coincidences (red) and true neutrons captures (blue). As expected, true neutron candidates are reconstructed close to the prompt vertex, while accidental coincidences can be reconstructed at larger distances. Here, the neutron vertex reconstruction is extended to the entire detector to take advantage of the long tail of the background distribution.

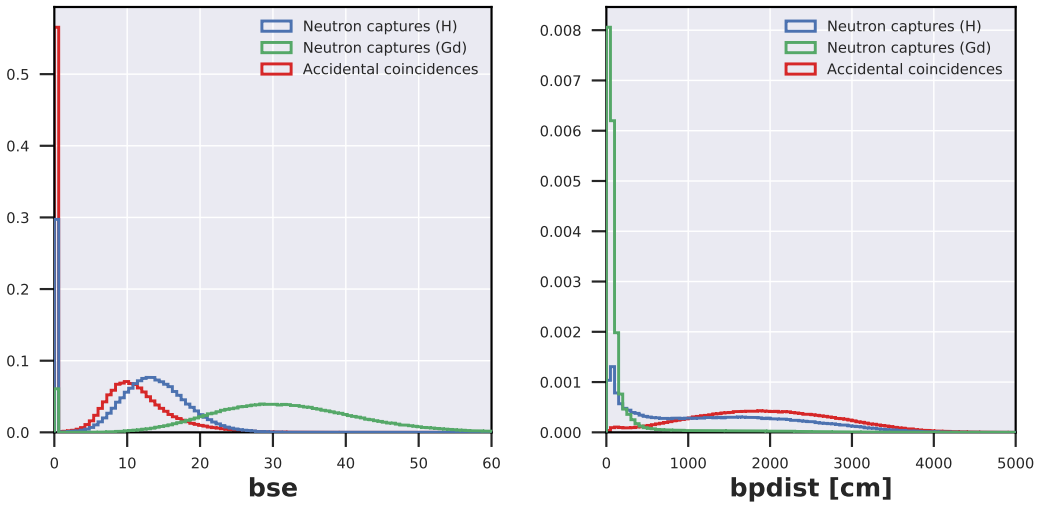

FIGURE 5.5: Probability density function of `bse` (left) and $\mathrm{bp}_{\mathrm{dist}}$, in cm (right), for accidental coincidences (red), true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). Values of 0 for `bse` indicate that the fit could not converge and are dominated by background. For vertex fits that converged successfully, background candidates display a flatter `bse` distribution. Here, $\mathrm{bp}_{\mathrm{dist}}$ is shown for candidates with non-zero `bse`. As expected, true neutrons display much greater consistency between the two vertex fits. In contrast, for accidental coincidences the two fits often differ by several tens of meters.

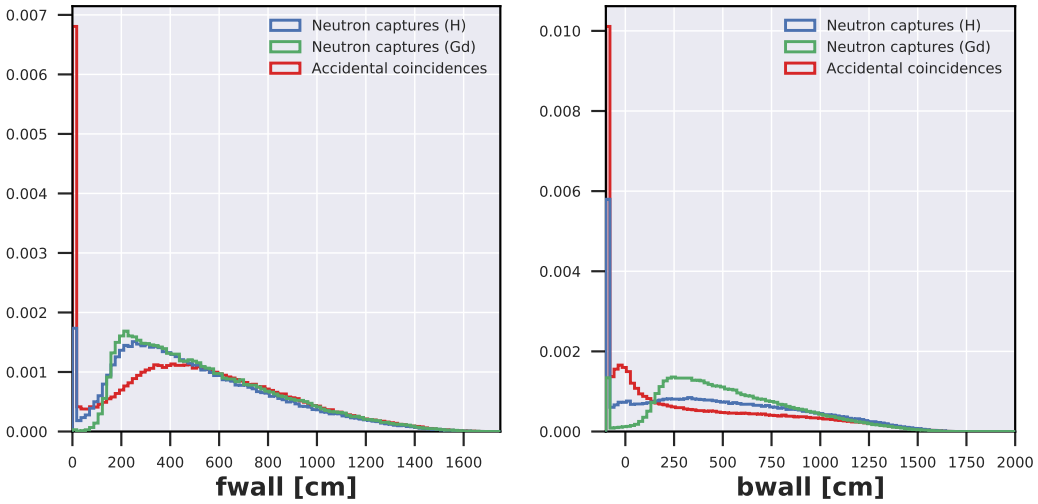

FIGURE 5.6: Probability density function of f_{wall} (left) and b_{wall} (right), both in cm, for accidental coincidences (red), true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). A large number of background candidates are reconstructed with their vertex near the detector walls. The BONSAI fitter can reconstruct the vertex outside the detector itself, unlike the t_{RMS}-based fit, leading to different behavior at the lower end of the distribution.

algorithm to occur ~ 0.5 m away from the prompt vertex, on average, compared to the average ~ 1.5 m distance travelled by neutrons before being captured by Hydrogen. Also notable is the difference in the spatial resolution between Gd and H neutron captures. If we approximate the spatial resolution as the full width at half maximum of the travel distance peak, the resolution on the Gd capture is about twice as good as that of the H captures, roughly 0.5 m and 1 m, respectively. The change in the number of hits when adjusting the neutron vertex from the prompt vertex to the newly-fit vertex, ΔN_{10}, is highly correlated with N_{10} and is higher for Gd captures.

In parallel, the BONSAI fitter is also used to independently search for the neutron vertex. Here, too, the differences between types of neutron captures are striking. First, the fitter fails to converge about one fifth as frequently as in the case of Hydrogen captures. As expected, the number of effective hits `bse` obtained in the fit is higher: ~ 13 hits for Hydrogen, ~ 30 hits for Gd, on average. This is consistent with a neutron capture signal with an average visible Cherenkov energy of ~ 5 MeV for Gd captures. The high end of the distribution is also consistent with the 8.536 MeV maximum visible energy from the capture on ^{155}Gd, when accounting for resolution effects. The distance bpdist, between the BONSAI vertex and the previously-fitted neutron vertex, is now much more tightly clustered around 0, indicating closer agreement between the two fits. Further, the bpdist distribution of Gd captures does not feature the tail present in the H capture distribution, which spans several tens of meters and indicates that the fitter converged to the wrong vertex. The `bse` and bpdist observables are shown in Figure 5.5. The distance between the fitted vertex and the detector wall, for both fits (f_{wall} and b_{wall}) are plotted in Figure 5.6. For Gd captures, the t_{RMS}-based fit virtually never converges to vertices on the wall, which indicate that the fit was driven by radioactivity or clustered noise hits. Instead, the only visible peak in the distribution, is at 2 m, corresponding to the 2-m fiducial volume cut applied to the prompt event. A similar behavior is observed for the distance from the wall of the BONSAI-fitted vertices that reached convergence.

Cherenkov event topology

The Cherenkov signal produced by a Gd capture looks quite different from that of a Hydrogen capture. Indeed, it isn't simply a more energetic Cherenkov cone. Rather, it is formed by multiple Cherenkov cones of variable number and direction. Reconstructing a single Cherenkov ring will therefore result in a wider, *fuzzier* ring. This translates to larger θ_{mean} and θ_{RMS}, plotted in Figure 5.7, together with ϕ_{RMS}. The mean opening angle of a Gd capture measures about $57°$, with a typical dispersion of $30°$ within the same event. Interestingly, the azimuthal symmetry of the Gd captures is markedly higher, as can be inferred from the ϕ_{RMS} distribution. This effect could be explained by the fact that reconstructing a single ring starting from a multi-ring event results in a reconstructed ring center that is relatively far away from the true centers of the multiple rings. At this higher distance, the angular separation between hits in the same ring will be lower. The lower impact of the background hits in the signal hit cluster also likely plays a role.

The observables dedicated to quantifying the number of hits in unlikely positions are more finely tuned for Hydrogen captures, so they are less meaningful for characterizing a Gd capture. The number of backward hits N_{back} is not as relevant since the Gd photon emission is not limited to a preferred direction. While it enjoys fairly high discrimination power on its own even for Gd captures, this is highly correlated with the discrimination from the overall number of hits N_{10}. Similarly, N_{low} and $N_{\text{low}\theta}$ see their background discrimination power diluted for Gd captures. Nevertheless, these observables can still be valuable to distinguish the two different neutron capture topologies from each other.

Noise characterization

Since Gd captures are less noise-like than Hydrogen captures, the distribution of the mean charge deposition per PMT hit, Q_{mean} is more strongly peaked at 1. The Q_{mean} and Q_{RMS} are plotted in Figure 5.8. As previously discussed, observables related to hit counts will be highly correlated with N_{10}. This is the case for N_{highQ}, N_{clus}, and N_{300}. In particular, the usefulness of these variable for distinguishing Gd captures from accidental background is either diminished, or highly correlated with the overall hit count. An interesting feature of the N_{clus} distribution is the multiple peak structure displayed for Gd captures. In Hydrogen captures we typically only have a maximum a single set of closely-clustered hits, though usually none are found. In Gd captures, these clusters are both more frequent and more numerous than either Hydrogen captures or noise hits. Since a cluster is defined as having three or more hits, this behavior is observed in the N_{clus} peaks at 3, 6, and 9 hits, indicating that single, double, and triple clusters are fairly frequent. This can be caused firstly by the overall hit number being higher, but also by the multi-ring nature of the signal, resulting in discrete clusters from different photons in the Gd cascade, or from the overlap of two rings. Finally, the $t^{(3)}_{\text{RMS}}$ and $t^{(6)}_{\text{RMS}}$ distributions, shown in Figure 5.9, show a more marked sharpening of the time dispersion peak from Gd captures when restricting the number of hits to 3 or 6. Indeed, the Gd signal is already purer that the Hydrogen signal, and by taking the best 3 or 6 hits we are most often left with a perfectly pure set of hits.

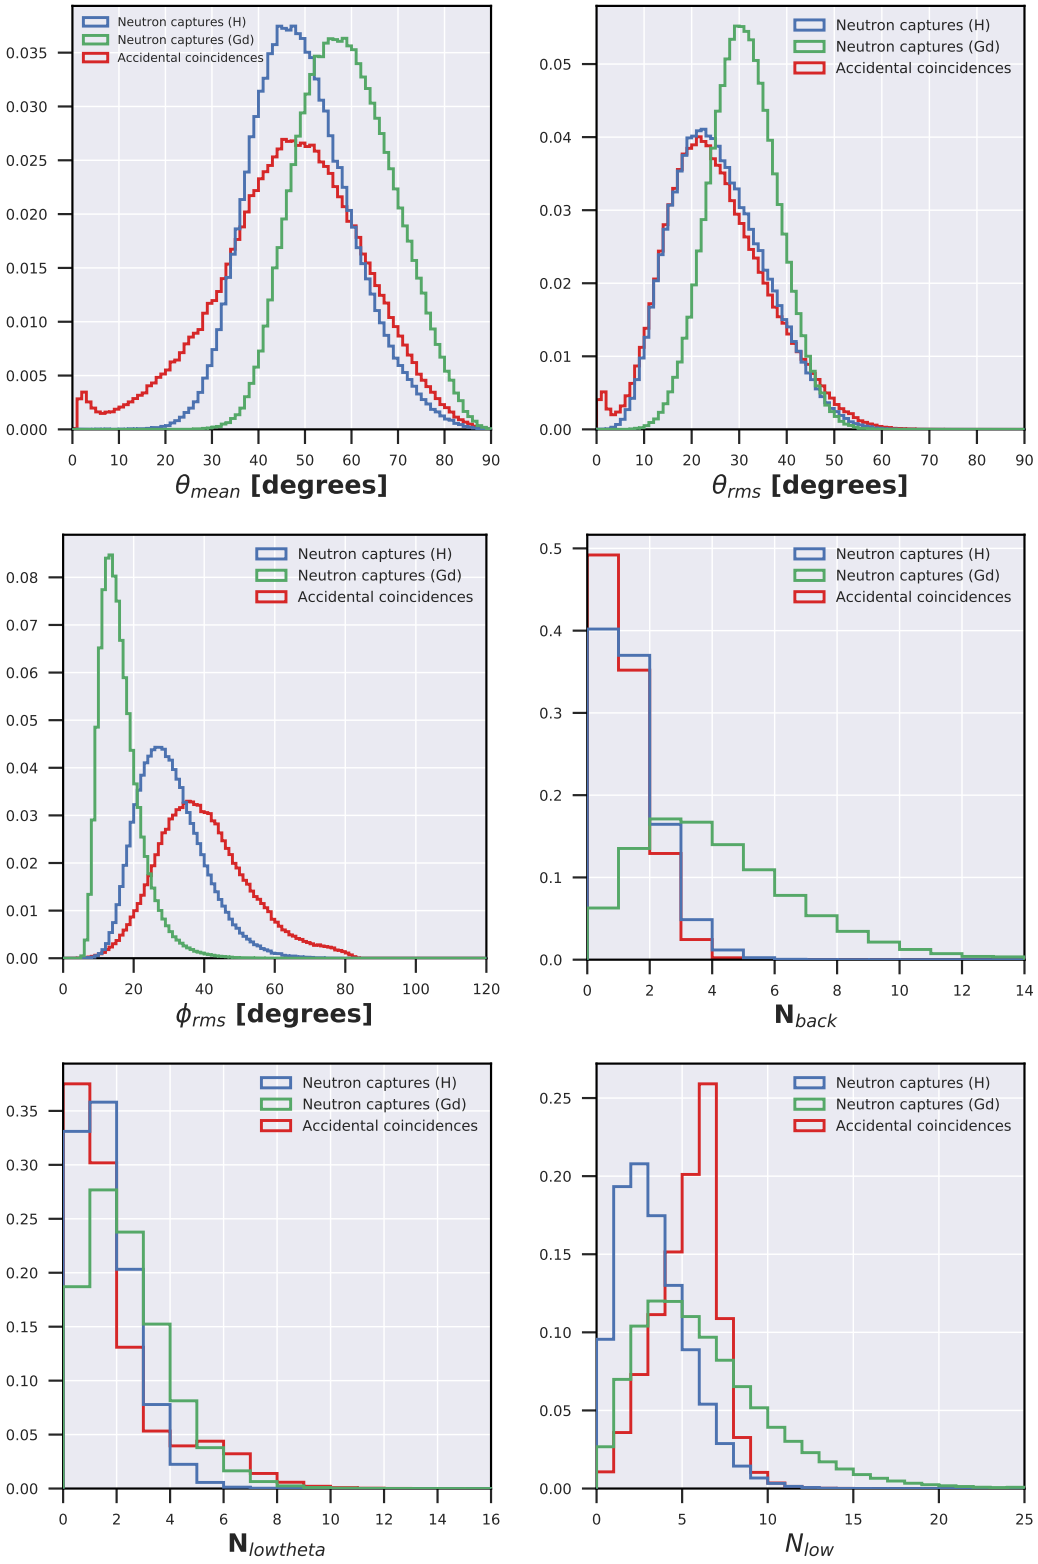

FIGURE 5.7: Probability density function of θ_{mean} (top left), θ_{RMS} (top right), and ϕ_{RMS} (middle left), N_{back} (middle right), $N_{\text{low}\theta}$ (bottom left), and N_{low} (bottom right) for accidental coincidences (red), true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). While our resolution does not allow us to precisely determine the Cherenkov angle of the neutron candidate, we can still clearly see the θ_{mean} peak at about $42°$ for the H-capture's single gamma, and a wider mean angle for Gd. Highly-localized backgrounds will be reconstructed with θ_{mean} and θ_{RMS} near 0.

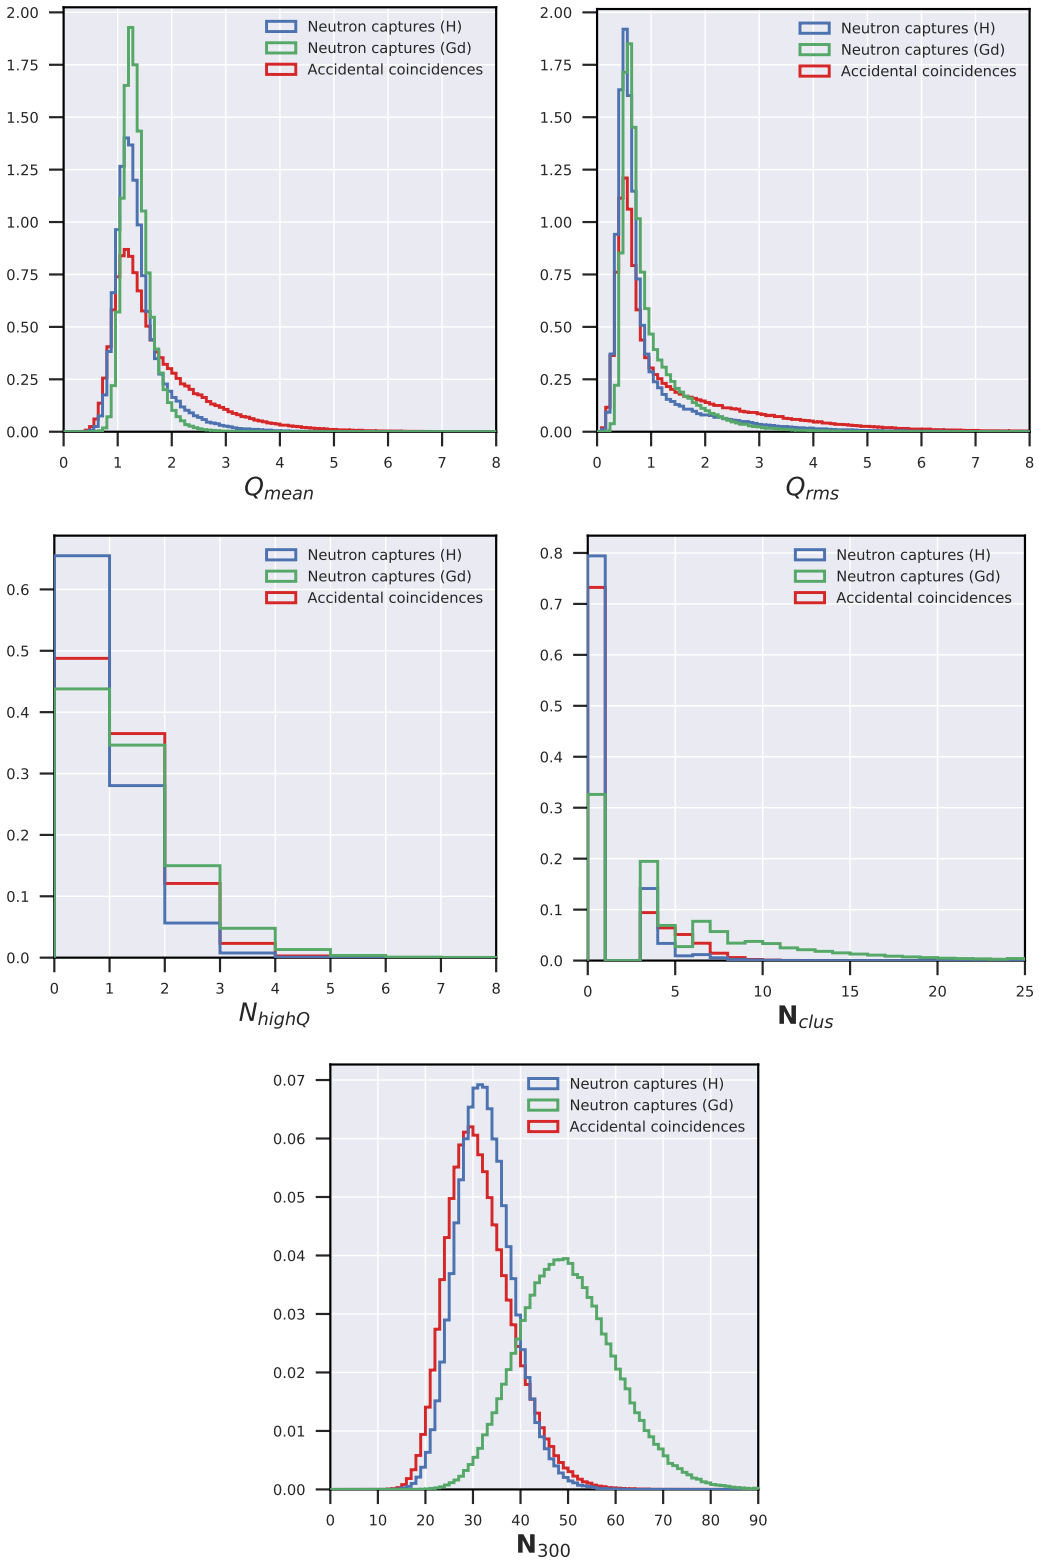

FIGURE 5.8: Probability density function of Q_{mean} (upper left), Q_{RMS} (upper right), both in photoelectrons/PMT, N_{highQ} (middle left), N_{clus} (middle right), and N_{300} (bottom), for accidental coincidences (red), true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). The discontinuity in the N_{clus} distribution is due to the definition of a cluster as containing a minimum of three hits.

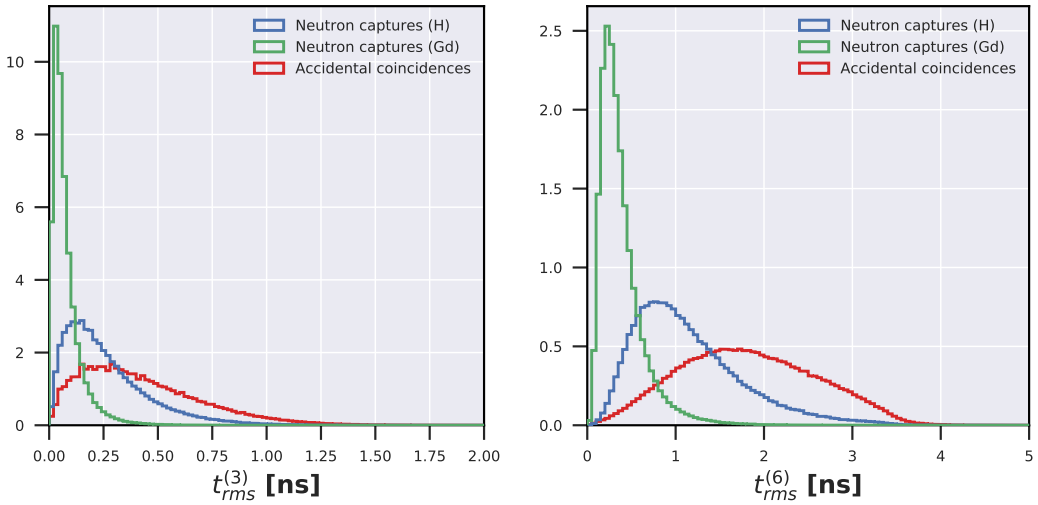

FIGURE 5.9: Probability density function of $t^{(3)}_{\text{RMS}}$ (left) and $t^{(6)}_{\text{RMS}}$ (right), in nanoseconds, for accidental coincidences (red), true neutrons captures on Hydrogen (blue), and true neutron captures on Gadolinium (green). Restricting the number of hits results sharpens the time dispersion peak of the signal to a greater extent than the background.

5.2.2 Boosted Decision Tree selection

As before, we entrusted the job of best utilizing these observables to a BDT algorithm. The same blueprint outlined in Section 3.5 was followed. A sample of 1 M simulated neutron captures and 10 M accidental coincidences was used to construct the model. A 50% fraction of this sample was used for the training itself to update the model parameters. Of the remaining half, 25% was used for model validation, and 25% was set aside for model testing. The `xgboost` package was used and the set of hyperparameters detailed in Table D.1 was chosen. The model was set to keep adding trees until its performance—defined as the area under the ROC curve calculated for the validation set—did not increase for more than 50 iterations. While a maximum of 1500 trees was allowed during training, the model converged much more quickly, reaching its best performance after 200 iterations.

After the training finished, the final performance of the model was calculated on the test dataset. Before discussing its performance, it is instructive to look at the relative importance of each variable in the BDT's decision-making algorithm, shown in Figure 5.10. The score of the N_{10} variable clearly stands out above the rest, demonstrating not only its obvious discriminating power, but also its correlation to variables that would have otherwise been very powerful on their own, but are instead relegated to the lowest ranks. It is perhaps an unsurprising result, but it confirms that, with the inclusion of captures on Gd, neutron tagging moves away from being a careful consideration of hidden correlations, becoming more of a straightforward, uncomplicated selection.

As usual, the performance of the neutron tagging algorithm is summarized by the ROC curve, plotted in Figure 5.11. Separating the performance of the BDT into Hydrogen and Gd neutron captures allows us to appreciate the wide gulf that is the difference in performance between the two. Indeed, the difference in the number of captures retained before we even apply the BDT accounts for much of this separation. This is visible in Figure 5.11 as the different maximum efficiencies reached by each ROC curve, with the Gd curve starting out with a two-fold advantage over the H curve. But this fact does not account for all of the difference: the shape of the curves

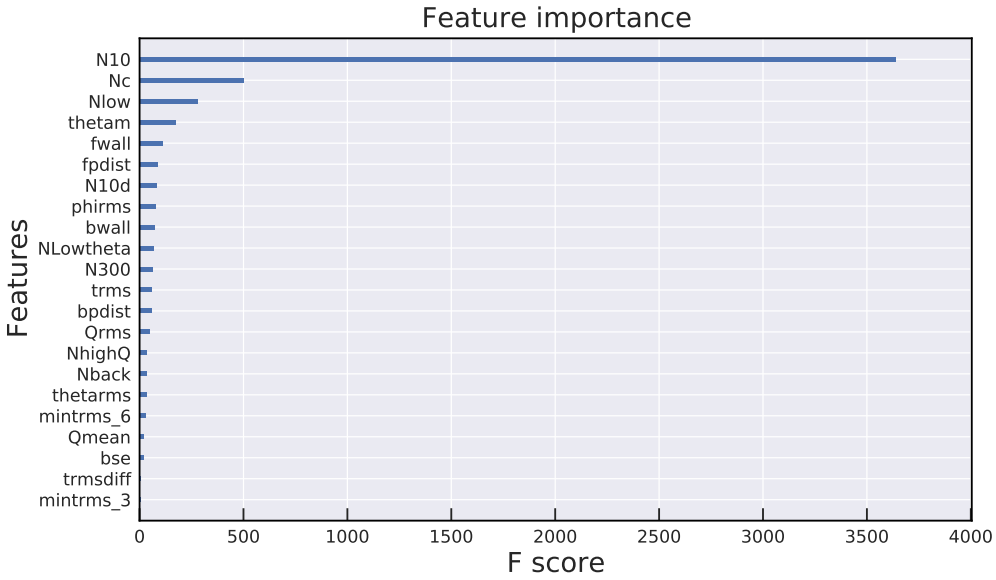

FIGURE 5.10: Relative discriminating feature importance in the BDT algorithm, as parametrized by the F-Score. Features scoring higher are utilized more often and have greater discriminating power when they are used.

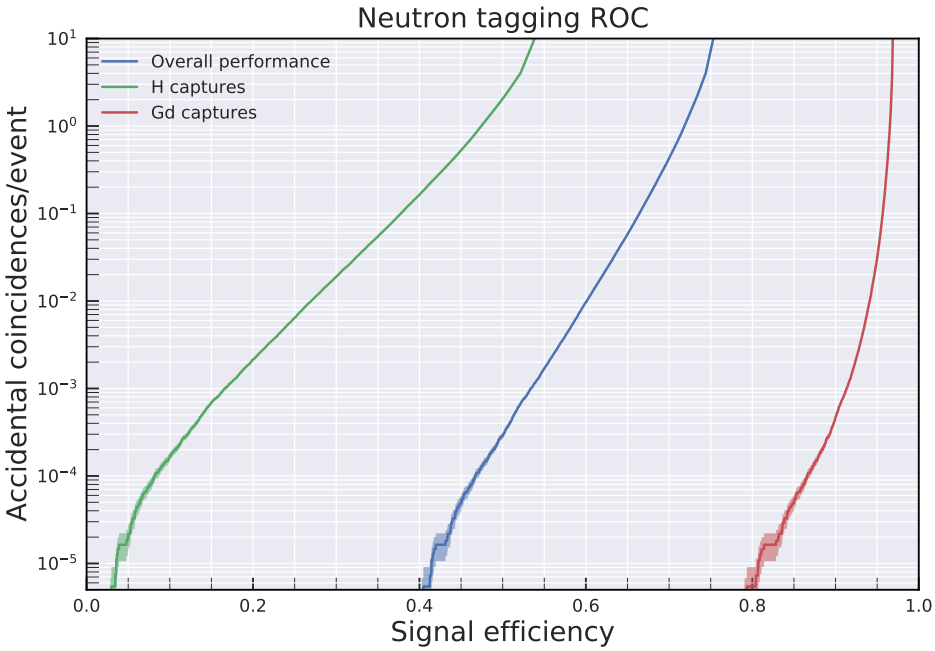

FIGURE 5.11: ROC curve for the SK-VI neutron tagging BDT model, shown for the entire test dataset (blue line), for the Hydrogen neutron captures only (green line) and for the Gd neutron captures only (red line).

inform us of the fact that—perhaps unsurprisingly—Gd captures are also much easier to distinguish even past the preselection stage. This can be observed as a steeper ROC curve for Gd, meaning that, compared to the pure-water case, the performance of the BDT in identifying a given capture is also better. Because of this steepness, the proportional difference in signal efficiency gets better the more we tighten our BDT cut. For example, if we take a looser cut, we could expect signal efficiencies of 34% and 95% for H and Gd captures, respectively, an improvement by a factor of about 3, much of which is obtained at preselection. If we now take a much tighter cut, leaving us with just 11% of the H captures, we still retain 88% of the Gd captures, an eight-fold improvement.

5.3 Prospects of future SK-Gd analyses

Armed with knowledge about SK's low-energy neutron tagging abilities with Gadolinium, we can start making predictions about the future of the DSNB search. Of course, making claims about the sensitivity of future analyses is fraught with uncertainties. In the interest of clarity, then, we should spell out the scope of the following study, and the assumptions on which it rests. First, we will mainly deal with characterizing a spectral analysis of the type presented in the previous sections. Different statistical treatments are of course possible and will present different challenges and systematic uncertainties. The performance of the prompt event selection cuts are assumed to be the same as in SK-IV. This particular assumption might eventually prove too conservative. For example, the spallation reduction cuts could benefit from the newly-doped SK water. Nevertheless, such effects would require new studies and are not considered here. To model the DSNB flux over the SK data-taking, we use the range of flux models plotted in Chapter 1.2.1, Figure 1.4. In particular, the range of DSNB interaction rates is defined by taking the Nakazato [26] and Horiuchi [20] models as minimum and maximum predictions, respectively. The Totani [73] and Kaplinghat [74] models are excluded from the study, due to their relative age, and the fact that their high flux predictions are disfavored by SK. The timeline below is chosen for the Gd era of the SK experiment.

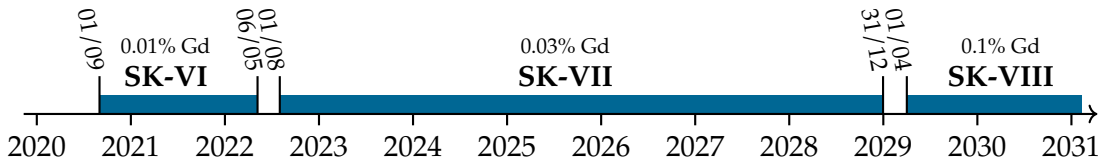

In particular, thee SK periods are considered: SK-VI, SK-VII (underway at the time of writing), and SK-VIII, with corresponding Gd concentrations of 0.01%, 0.03%, and the ultimate SK target of 0.1%. Between each data period, a three-month hiatus is added. This gap accounts for the detector downtime needed to introduce the Gd sulfate in the tank, as well as the online time needed to let the Gd evenly disperse, and to allow detector conditions to stabilize before good-quality data is available. The end date of SK-VII and the eventual start of SK-VIII is the subject of much discussion. While the nominal SK target is to reach a 0.1% Gd concentration, logistical and budgetary challenges make this proposition less certain. At present, achieving nominal concentration is not a short-term goal of the collaboration. Current SK funding extends until the end of the 2027 the Japanese fiscal year, i.e. until early 2028. This date corresponds to the expected start of the Hyper-Kamiokande data-taking,

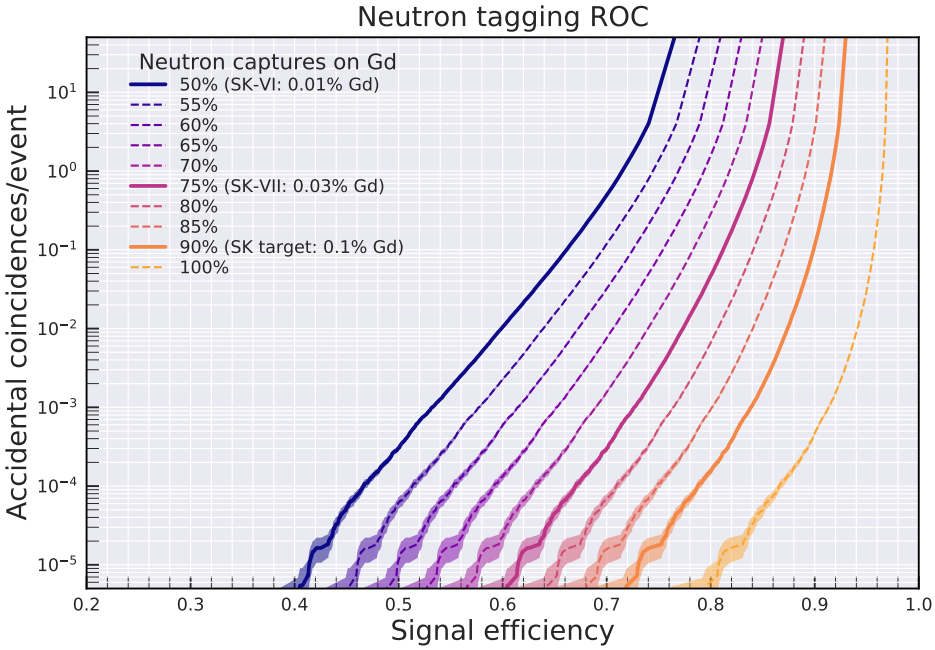

FIGURE 5.12: ROC curve for the SK-VI neutron tagging BDT model, shown for different fractions of neutron captures occurring on Gd. The solid lines correspond to the Gd concentrations of actual SK periods, including the final planned Gd concentration of 0.1% (with 90% of captures occurring on Gd).

beyond which point the future of SK remains less certain. For the purposes of this study, SK-VII was considered to run for more than six years, until 2029.

While the neutron tagging algorithm presented in the previous section was developed with a 0.01% Gd simulation, it can be extended to different concentrations to evaluate its performance at different SK periods. Crucially, increasing the Gd concentration in the detector should not have a significant impact on key detector conditions such as the water transparency or the dark noise in the PMTs, as studied with the EGADS experiment [82]. As mentioned above, while an increased dark noise rate was observed starting with SK-V, this increase was not associated with the introduction of Gd into the tank, so we don't have a reason *a priori* to expect further changes in the accidental backgrounds due to further doping of the water. We can then approximate the detector response as a function of the Gd concentration, by simply varying the proportion of neutron captures that occur on Gadolinium. To evaluate the effect of the BDT at different Gd concentration, we can estimate its efficiency by separating its performance on H captures from Gd captures, and reweighting each one accordingly. The resulting ROC curves, plotted for a range of Gd concentrations, are shown in Figure 5.12. Due to the nonlinear nature of the Gd capture fraction as a function of concentration, the SK-VII neutron tagging performance is not far from the expected performance of the detector at nominal 0.1% concentration, despite having reached less than a third of this concentration.

Atmospheric neutrino background estimation

To estimate the background contribution, the main source taken into account are atmospheric neutrino backgrounds, which will dominate the energy range. Solar

neutrino and spallation backgrounds are assumed to be reduced to the point of being negligible or largely subdominant, similarly to SK-IV, and are neglected. The reactor neutrino background will also be neglected, though it will define the energy threshold of the search. When this study was conducted, no simulation of the atmospheric backgrounds was yet available for the SK-VI period. For this reason, the atmospheric background simulation for SK-IV (as described in Section 3.2.5) is used. The features of the prompt events from atmospheric neutrino interactions are not expected to change significantly from the SK-IV, for the purpose of this study. On the other hand, atmospheric neutrino events can often be accompanied by the production of secondary neutrons. The fact that we can see neutron captures more clearly with SK-Gd will therefore also affect these backgrounds, so we cannot directly apply the SK-Gd BDT to this simulation. Rather, the impact of the neutron tagging is modeled as a convolution of two effects. First, the effect of the algorithm's efficiency on IBD-like neutrons ε, as evaluated in the previous section. For an event with a neutron multiplicity n, this translates to a probability $P_{\text{ntag}=1}$ of passing the neutron tagging cut with exactly one neutron being tagged:

$$P_{\text{ntag}=1}(n) = P^{\text{neut}}_{\text{ntag}=1}(n) + P^{\text{acc}}_{\text{ntag}=1}(n) \tag{5.3}$$

$$P^{\text{neut}}_{\text{ntag}=1}(n) = \varepsilon(1-\varepsilon)^{n-1} \times (1-\varepsilon_{\text{acc}})^{N_{\text{acc}}} \tag{5.4}$$

$$P^{\text{acc}}_{\text{ntag}=1}(n) = (1-\varepsilon)^{n} \times \varepsilon_{\text{acc}}(1-\varepsilon_{\text{acc}})^{N_{\text{acc}}-1}, \tag{5.5}$$

where $P^{\text{neut}}_{\text{ntag}=1}(n)$ is the probability of tagging exactly one true neutron, $P^{\text{acc}}_{\text{ntag}=1}(n)$ is the probability of tagging exactly one accidental coincidence, ε_{acc} is the accidental background acceptance, and N_{acc} is the average number of accidental coincidences after preselection. A more naïve approach would be to simply rescale the background expectation according to the neutron tagging efficiency in SK-Gd. However, the performance of the tagging algorithm presented above is only valid for identifying neutron captures from IBD interactions. The coincidence signature from the secondary neutrons in atmospheric interactions can differ significantly from the IBD coincidence, yielding a different neutron tagging performance for these backgrounds. Typically, this translates to a greater rejection of these neutrons, and an overall decrease in the atmospheric backgrounds in the signal region. Therefore, a second correction factor accounts for this difference. This effect is estimated for the SK-IV simulation, by comparing the rates predicted by Equations 5.3-5.5 with the BDT trained for SK-IV BDT, to the actual rates obtained by directly applying the same BDT to the atmospheric simulation, $N^{\text{ntag, SK-IV}}$. Since different types of atmospheric backgrounds produce secondary neutrons in different ways and with different topologies, a correction factor c_i is calculated for each background category i:

$$c_i = \frac{N_i^{\text{ntag, SK-IV}}}{\sum_{j=1}^{N_i} P^{\text{SK-IV}}_{\text{ntag}=1}(n_{ij})} \tag{5.6}$$

where, in the denominator, we apply the neutron tagging probability event-by-event to the N_i atmospheric events of type i, to yield the expectation value of the rate. The atmospheric rate estimate $N_i^{\text{ntag, SK-Gd}}$ for each background i, after neutron tagging, for SK-Gd, is then

$$N_i^{\text{ntag, SK-Gd}} = c_i \sum_{j=1}^{N_i} P^{\text{SK-Gd}}_{\text{ntag}=1}(n_{ij}). \tag{5.7}$$

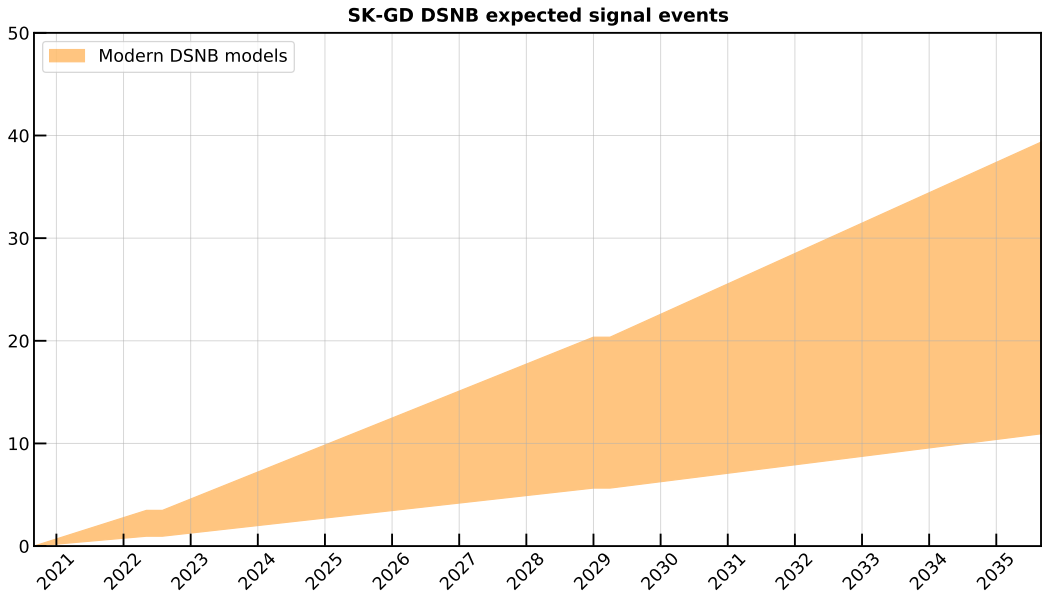

FIGURE 5.13: Cumulative number of DSNB signal events over the lifetime of SK-Gd, for events reconstructed in a 12-30 MeV energy window that pass the neutron tagging cut. The yellow band represents the variation due to the different flux predictions of the DSNB models considered. The flat sections correspond to breaks in the data-taking.

The values of $P_{\text{ntag}=1}^{\text{SK-Gd}}$ and $N_i^{\text{ntag, SK-Gd}}$ will vary depending on the Gd concentration. The background categories i used in this estimate are the same as the ones defined in Section 4. This categorization, which was initially chosen according to the spectral shape of each one, is also useful to distinguish between different neutron capture topologies. Take for example the decay-electron category, in which the prompt event is triggered by the electron from the decay of an invisible muon or pion. One feature of these events is that the neutron capture will occur, on average, further away from the prompt vertex than is typical of an IBD event. Since the algorithm is tuned for neutrons close to the prompt vertex, it will achieve a lower efficiency for these distant neutrons.

5.3.1 A simplified picture

With an idea of the attainable neutron tagging efficiencies and an estimate of the atmospheric background contribution, we can start estimating the expected DSNB signal significance of future searches. Before carrying out a more sophisticated study, we can obtain a first estimate of the significance for a simple counting experiment with a single bin. This approach will give us a first estimate, and has the advantage of being transparent and easily reproducible.

We first define a signal region between 12 and 30 MeV. We also require that exactly one neutron be tagged. To estimate the number of signal events, the DSNB flux expectations are convolved with an energy resolution function to mimic the effects of energy reconstruction. To account for the neutron tagging efficiency, an accidental coincidence rate of < 0.01 per prompt event is required (i.e., achieving similar levels of background rejection as SK-IV), yielding efficiencies of roughly 60%, 77%, and 87% for the neutron tagging cut with SK-VI, SK-VII, and SK-VIII, respectively. The cumulative signal expectation in this signal region is plotted in Figure 5.13. Note that the DSNB flux models considered are not uniformly distributed

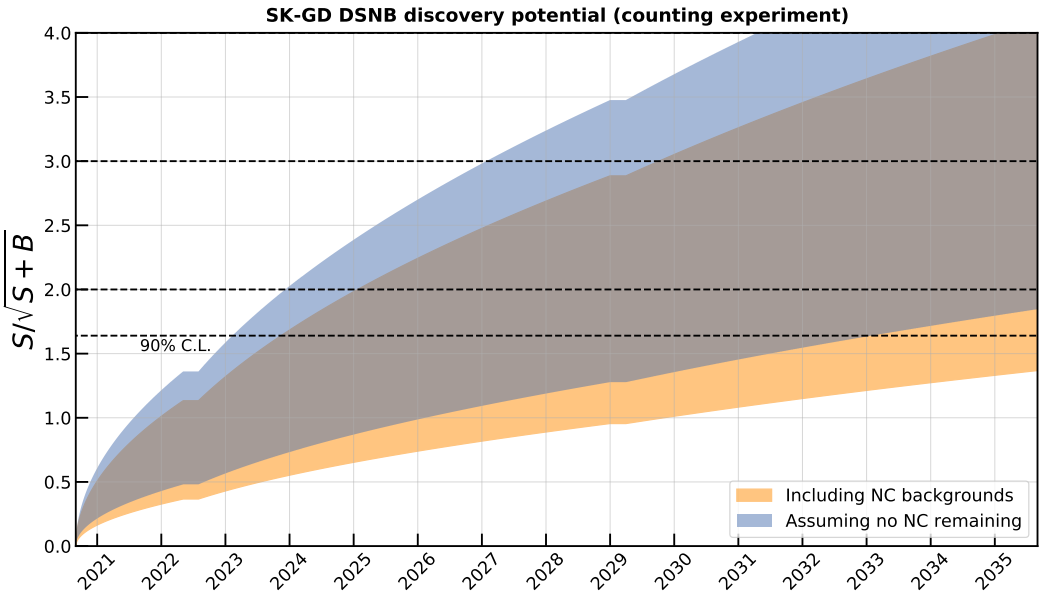

FIGURE 5.14: DSNB signal significance events over the lifetime of SK-Gd, for events reconstructed in a 12-30 MeV energy window that pass the neutron tagging cut. The bands represent the variation due to the different flux predictions of the DSNB models considered. The flat sections correspond to breaks in the data-taking. The yellow band uses the nominal atmospheric background prediction, while the blue band assumes complete removal of the NC backgrounds.

across this range of expectations. In particular, the upper half of the range is more populated than the lower half, in which the Nakazato model predicts a particularly low flux.

The signal significance as a function of time is then calculated as the figure of merit $S/\sqrt{S+B}$. This simplified model only accounts for the statistical uncertainty of the measurement, neglecting systematic uncertainties. To model the impact of a potential improvement in the rejection of the atmospheric backgrounds, the study is repeated with the removal of the neutral current contribution. This choice was not arbitrary: a recent study, external to SK, put forward the possibility of reducing this type of background at SK by more than 95% using machine learning techniques, with excellent signal efficiency [83]. Indeed, NC backgrounds are not irreducible, and it is conceivable that we could better exploit their signature than we currently do. These backgrounds are also peaked in a similar energy region as the DSNB and are therefore of particular interest. The extent to which this study can be reproduced with official SK simulation, and how it compares to the existing reduction cuts, is currently under investigation within the collaboration.

The resulting projections are showed in Figure 5.14. The projection for the significance without NC background can be interpreted as the margin for improvement that is possible through better removal of this background. It is clear that due to the wide range of flux models, the projected potential for a DSNB detection is highly variable. For the lowest flux predictions, detection prospects are not bright, with the Nakazato model not reaching a 2σ level within the foreseeable SK-Gd running time. On the other hand, more the more optimistic models, a 2σ level could be reached by 2025, with potential analysis improvements allowing for this goal to be reached as early as 2024.

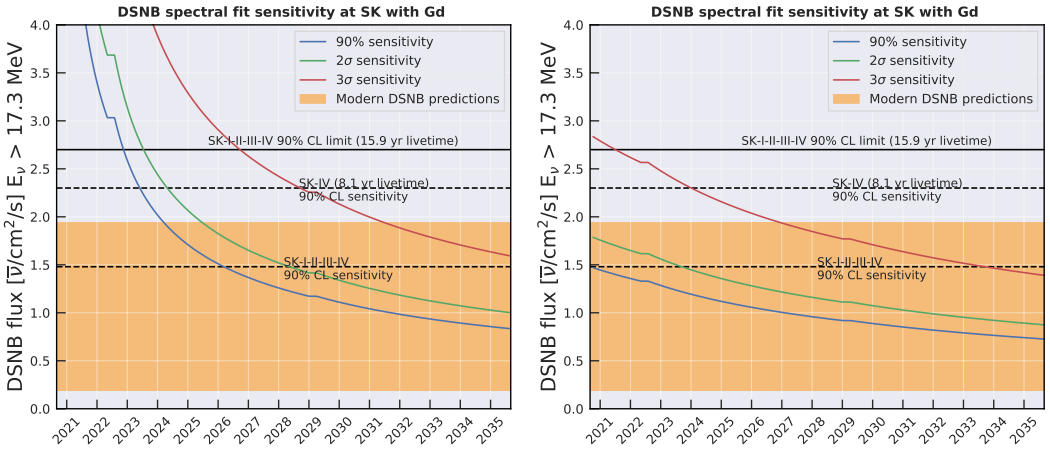

FIGURE 5.15: DSNB flux exclusion sensitivity of a spectral analysis conducted at SK-Gd. The sensitivity is shown for the 90% (blue line), 2σ (green line), and 3σ confidence levels. The yellow band represents the range of predicted fluxes. The solid black line is placed at the combined SK-I-II-III-IV 90% C.L. spectral analysis upper limit. The dashed lines are placed at the 90% C.L. exclusion sensitivity of SK-IV (top line) and of the combined SK-I-II-III-IV (bottom line). Note that the fluxes are given above a neutrino energy of 17.3 MeV to ease the comparison to previous analyses, which does not reflect the energy threshold of the projected SK-Gd search. On the right, the expected sensitivities are shown for an analysis combining the SK-Gd results with the SK-I-II-III-IV results.

5.3.2 Spectral analysis simulation

For a more realistic projection, we now turn to modelling the performance of a search like the SK-IV spectral analysis for SK-Gd. This will include several effects not accounted for in the simpler model in the previous section. First, the spectral shape features of the DSNB and of the backgrounds will be taken into account. All else being equal, this should allow for better performance compared to a single-bin analysis. Second, the full systematic uncertainties of the spectral analysis will be included, which will make our projections significantly more conservative. Finally, in this study we will also analyze the events that do not pass the neutron tagging cut, as in SK-IV. While in principle this will improve the power of our search, in practice its effect will likely be limited. Unlike in SK-IV, starting with SK-VI most of the DSNB candidates will pass neutron tagging—especially starting with SK-VII, when we can already expect efficiencies approaching 80%. Note that while we will be analyzing both the events that pass and those that fail the neutron tagging cut, the spectral fit will be more sensitive as more of the signal falls in the $N_{\text{ntag}} = 1$, which is less contaminated with background. In addition, we can attain a lower energy threshold in the $N_{\text{ntag}} = 1$ region, meaning that both signal purity and quantity increase with higher Gd concentrations.

A frequentist approach is taken for this study: the strategy will be to generate a large sample of artificial datasets according to a set of starting assumptions, apply a spectral analysis on each one, and assess the expected performance of our search. We first investigate the background-only hypothesis H_0, with the objective to evaluate the analysis' sensitivity to model exclusion. Each dataset is a set of energies $\{E\}$, divided into the six spectral analysis regions (two neutron tagging regions and three Cherenkov angle regions, as in Section 4). The energies are sampled from the background spectra $P_B^i(E; \theta_C, N_{\text{ntag}})$, for background category i in Cherenkov angle

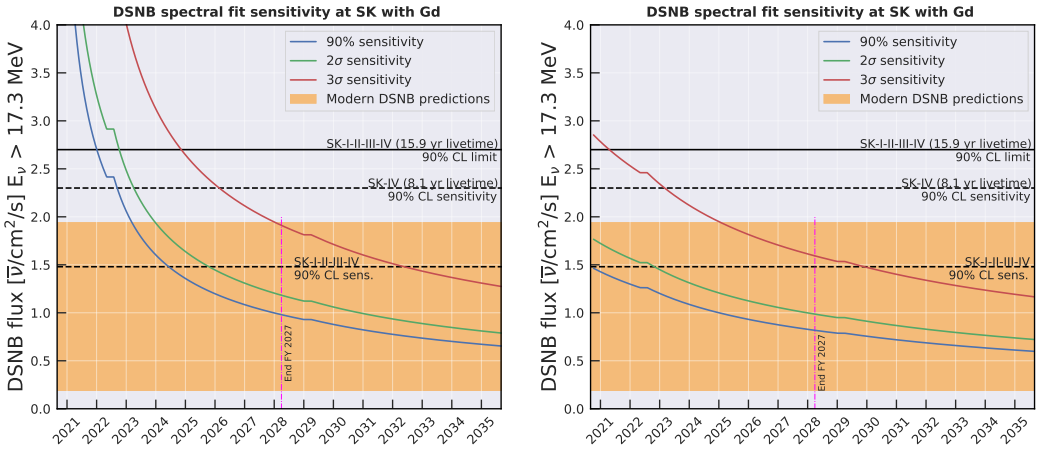

FIGURE 5.16: DSNB flux exclusion sensitivity of a spectral analysis conducted at SK-Gd, assuming an $\mathcal{O}(50\%)$ NC spectral shape systematic uncertainty. On the right, the expected sensitivities are shown for an analysis combining the SK-Gd results with the SK-I-II-III-IV results.

region θ_C and neutron tagging region N_{ntag}:

$$P(E|H_0) \equiv P_B(E;\theta_C, N_{\text{ntag}}) = \sum_i P_B^i(E;\theta_C, N_{\text{ntag}}) \tag{5.8}$$

$$E = \{E_1, ..., E_N\}, \tag{5.9}$$

where the total number of events N follows a Poisson distribution with the expected rate N_B predicted by the simulation. Here, the Gd concentration will affect the proportion of each P_B^i occupying the two neutron tagging regions. The systematic uncertainty on the atmospheric background shapes enters in the generation of the datasets by stochastically modifying the nominal energy spectra $P_{B,0}^i(E)$, of each background category i, with a spectral distortion function $S_i(E;\epsilon_i)$, parametrized by ϵ_i:

$$P_B^i(E) = S_i(E;\epsilon_i)\ P_{B,0}^i(E), \tag{5.10}$$

where ϵ_i is distributed according to prior distribution $\text{Pr}_i(\epsilon_i)$. The spectral distortion functions parametrizations and the prior distributions follow the same scheme as those introduced in Section 4.

After generating the datasets, we apply our spectral analysis. Here we fit a chosen DSNB model, whose spectrum after reduction cuts has been evaluated according to the performance of the neutron tagging algorithm for the appropriate SK period. The background expectations, which are fit as nuisance parameters, are also adapted according to the SK period, as before. The systematic uncertainties on the spectral shapes of the backgrounds enter a second time during the fit. The energy threshold is set at 12 MeV for the $N_{\text{ntag}} = 1$ region, and at 16 MeV for the $N_{\text{ntag}} \neq 1$ region. The upper flux limit is extracted for each model as in Section 4. The exclusion sensitivity of the analysis, at a given confidence level, is then defined as the median upper limit at the chosen confidence level, given the background hypothesis H_0. The resulting sensitivities are plotted in Figure 5.15. In this projection, the SK-IV sensitivity (corresponding to a livetime of 8.1 years) could be reached by SK-Gd in less than three years, in mid-2023, whereas the combined SK-I-II-III-IV sensitivity (a 15.9-year livetime) could be reached in early 2026, in less than six years. The effect of

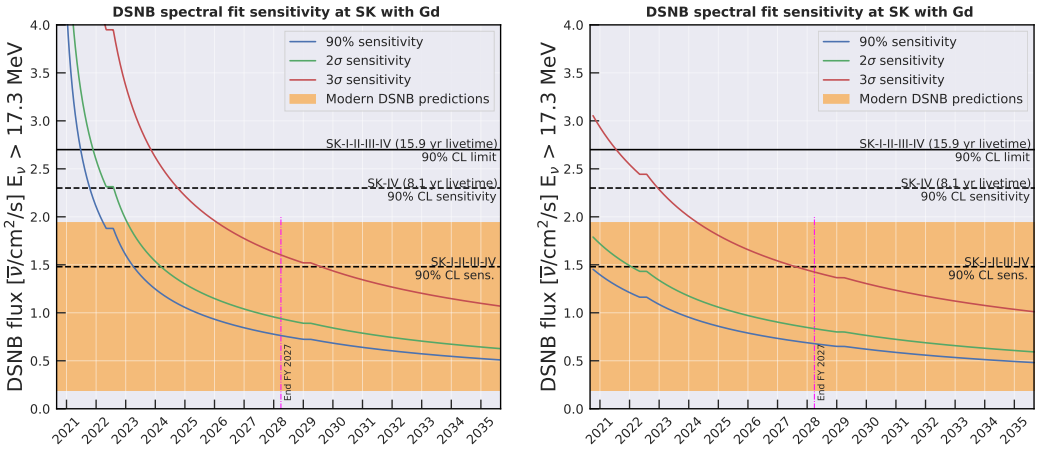

FIGURE 5.17: DSNB flux exclusion sensitivity of a spectral analysis conducted at SK-Gd, assuming a complete removal of the NC backgrounds. On the right, the expected sensitivities are shown for an analysis combining the SK-Gd results with the SK-I-II-III-IV results.

combining the SK-Gd limits with the SK-I-II-III-IV limits was also tested. Here, we see that within a little over six years, the combined 90% C.L. sensitivity could already cover half of the range of predicted fluxes. Within 10 years, a combined spectral analysis could exclude a significant portion of this range at the 3σ level.

As already mentioned, one of the hampers on this analysis' sensitivity is the impact of atmospheric neutral-current backgrounds and the associated systematic uncertainty in the Cherenkov angle cut efficiency. Since this background could represent a veritable bottleneck for our search, it must be given special attention. In the spectral analysis, a conservative $\mathcal{O}(100\%)$ systematic uncertainty was assigned to the amount of NC backgrounds passing the Cherenkov angle cut. The overall impact of this assumption was limited by the fact that the analysis had a relatively high energy threshold of 16 MeV, such that the NC contribution was still modest in the Cherenkov angle signal region. As we lower the energy threshold to make full use of the Gd-doped detector, the NC contribution becomes dominant in the lower energy range. An overly-conservative estimate could therefore artificially weaken the sensitivity of the analysis. Recent measurements [84] have observed a discrepancy from simulation in the Cherenkov angle distribution, such that the variation in our signal region would be of $\mathcal{O}(30\%)$ within 1σ. This measurement is limited by low statistics, and further work is needed to characterize and understand this discrepancy. For the sake of this study, the sensitivity projection was repeated with a systematic uncertainty on the NC spectral shape reduced to $\mathcal{O}(50\%)$. This is likely a fairer way to capture our knowledge of these backgrounds. The modification of the uncertainty impacts both the prior probability of the generated NC spectra and the uncertainty in the spectral fit. The resulting projections are shown in Figure 5.16. This adjusted projection tells us that we could have already reached the SK-IV sensitivity, shortly after the start of SK-VII in mid-2022, and that we might reach the combined SK-I-II-III-IV sensitivity less than two years thereafter, in mid-2024. By the end of the 2027 Japanese Fiscal year, until which time SK is guaranteed to take data, the 90% C.L. sensitivity of the SK-Gd analysis considered here will cover about half of the modern flux predictions. With another five years, the exclusion sensitivity of a combined analysis could cover three quarters of the predicted fluxes at 90% C.L., and about 40% at the 3σ C.L.

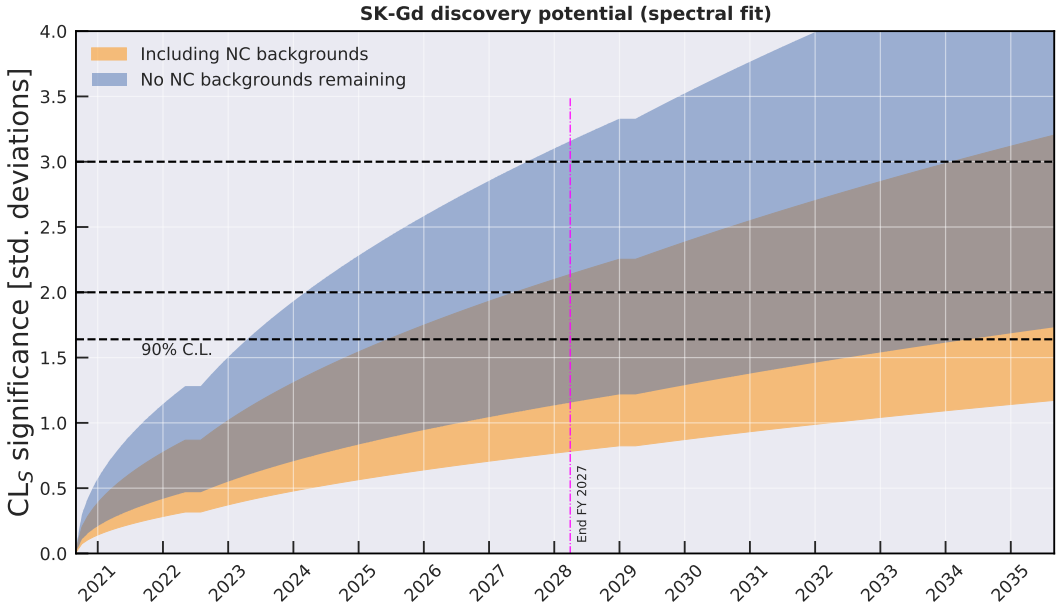

FIGURE 5.18: Evolution of the expected DSNB detection significance at SK with a spectral analysis using a 12 MeV energy threshold, for a range of DSNB flux models. Two hypotheses are shown: a projection assuming an $\mathcal{O}(100\%)$ systematic uncertainty associated with the NC backgrounds (yellow band) and a projection assuming efficient removal of the same backgrounds (blue band). The variation within each band is due to the variation of the DSNB models. The standard deviation is obtained from the CL_S value assuming a Gaussian distribution.

Once again, we also test the effect of the efficient removal of the NC background from the signal region. In this case, the dominant background contribution will be from the electrons produced in the decay of invisible muons with production of a secondary neutron. The resulting projections are plotted in Figure 5.17. Compared to the previous projection, such an assumption reduces the time needed to reach similar sensitivities by about a third.

Finally, we investigate the discovery potential of the spectral analysis. The signal hypothesis H_1 is constructed by adding a model-dependent DSNB flux to the expected energy spectrum:

$$P(E|H_1) = P_{\text{DSNB}}(E;\theta_{\text{C}},N_{\text{ntag}}) + P_B(E;\theta_{\text{C}},N_{\text{ntag}}), \tag{5.11}$$

after which the generation of the dataset proceeds as before. A pool of datasets is generated for various DSNB flux models. For each model considered, the spectral fit is carried out by fitting the same DSNB model that was generated. For each dataset $\{E\}$, we then define the likelihood ratio LR:

$$\text{LR}(\{E\}) = \frac{\mathcal{L}(\Phi_{\text{DSNB}} = \Phi_{\text{pred}}|\{E\})}{\mathcal{L}(\Phi_{\text{DSNB}} = 0|\{E\})}, \tag{5.12}$$

using the likelihood function defined in Section 4, and using the flux Φ_{pred} predicted by the DSNB model in question. This likelihood ratio is also calculated for the background-only datasets. For the H_0 hypothesis and for each H_1 hypothesis, the probability density functions $P(\text{LR}|H_0)$ and $P(\text{LR}|H_1)$ are inferred. The discovery

significance calculated for dataset $\{E\}$ is then defined by the CL_S statistic:

$$CL_S(\{E\}) = 1 - \frac{p_B}{p_{S+B}} = \frac{\int_{l=\mathrm{LR}(\{E\})}^{\infty} P(\mathrm{LR} = l\,|H_0)\,dl}{\int_{l=\mathrm{LR}(\{E\})}^{\infty} P(\mathrm{LR} = l\,|H_1)\,dl}, \tag{5.13}$$

where $l = \mathrm{LR}(\{E\})$ is the observed value of LR. Using this statistic as our discovery criterion allows us to be more conservative in our estimate. The CL_S statistic can be thought of as a rescaling of the p-value p_B, that prevents us from making excessive claims in cases of low sensitivity. Given the H_1 hypothesis, the expected value of the LR observation is extracted, and the corresponding CL_S expectation is computed. The study is repeated for the considered model range. The impact of the NC backgrounds is also tested for. The results are shown in Figure 5.18. It is not particularly surprising that the projected potential for discovery using this method should be more pessimistic than the simplified counting experiment showed above. Indeed, the inclusion of the systematic uncertainties in the analysis causes a remarkable rift between the nominal analysis and the analysis assuming the removal of the NC backgrounds. Comparing the performance with the simplified case also indicates that this effect mostly impacts the nominal analysis. The performance of the analysis with the NC backgrounds removed is comparable between the simplified projection and the spectral analysis projection. These observations confirm not only that atmospheric spectral uncertainties can have a significant impact at low energy thresholds, but also that the NC systematic uncertainty is the main systematic uncertainty affecting our sensitivity to the DSNB. It should therefore be one of the priorities of future DSNB searches to further characterize these backgrounds. With a bit of optimism, we can interpret the shift between the two bands in Figure 5.18 as the significant potential for improvement for new studies in this domain.

6
Search for the DSNB with SK-VI

In the final section of this thesis work, the machinery presented so far is applied to search for the DSNB using the latest SK data, from the SK-VI data-taking period. This is the first SK period after the addition of gadolinium in the tank, and thus represents an important turning point in the collaboration's DSNB search efforts. Note that for the purposes of this work, we will gloss over the SK-V period. This data-taking period represented the transition between SK-IV and the SK-Gd era, and included much of the preparatory work necessary for the doping of the water. Due to its relative shortness, and the fact that it still uses pure water, analyzing the data from this period was not an immediate priority for the collaboration at the time of writing.

In the analysis that follows, the overall structure of the SK-IV DSNB search, as presented in Chapter 3, will be preserved. As such, the focus of this chapter will be on the changes made to the previous analysis, as well as any new measurements and results obtained with the SK-VI data samples. In addition, some aspects of the analysis didn't change significantly from the SK-IV analysis, but were only part of this thesis work for this new analysis. These aspects will therefore be treated in more detail in this chapter than in Chapter 3.

The reader should also note that the results obtained from this search should still be considered preliminary. As Super Kamiokande is still in the early stages of its Gd journey, the strategies and techniques to make the most of it are still in a state of flux. Further improvements in one or more aspects of the analysis presented here are therefore possible, if not expected. Nonetheless, this chapter will attempt to give a first look of the new potential of the detector, and to lay the groundwork for the next analyzers.

The logical progression of the chapter will follow the model of Chapters 3 and 4. In this analysis, the inverse beta decay remains the detection channel of this search. In addition, the changes in its signature due to the addition of the gadolinium sulfate have already been addressed in Chapter 5, so no further discussion about the detection channel will be necessary. The background sources in this search are also unchanged from SK-IV, so no additional description of the backgrounds is provided. In Section 6.1, the SK-VI data samples, and the additional simulation sampled required by the analysis are discussed. The preliminary data reduction step, as well as the results of the IBD positron identification, are presented in Section 6.2. Once again, the development and tuning of the cuts aimed at the removal of the spallation backgrounds was not part of this work, but their form follows the strategy in Appendix A. The neutron tagging cut is then discussed in Section 6.3. Finally, the statistical treatment of the final sample and the results of the analysis are given in Section 6.4.

6.1 Data and simulation samples

Data sample

The data sample chosen for this search will comprise the SK data runs between runs 85220 and 87220. This corresponds to the period between August 28th, 2020, and June 1st, 2022. The start of the data sample is chosen such that the detector conditions have had a chance stabilize after the addition of the Gd into the water. After this point, all the available SK data until the end of SK-VI is used. After accounting for calibration runs, poor quality runs, or otherwise unusable data, we are left with 520.2 days of livetime for the SK-VI period. As in the previous analysis, the trigger used to define our event will be the SHE+AFT trigger pair, allowing for the search of the prompt and delayed signals in coincidence. For the majority of the livetime, the threshold for the SHE trigger was set to 60 hits in 200 ns. In principle, each SHE trigger should be immediately followed by an AFT trigger. Unfortunately, for a significant fraction of the events in SK-VI, this process was compromised. Indeed, starting with SK-VI, an AFT trigger is activated after every OD trigger, and not just after an SHE trigger. This change was made to aid in the detection of the secondary neutrons in the hadronic shower accompanying cosmic muon events. However, this had the unforeseen consequence of producing more AFT events than the DAQ software was prepared to handle, due to the high cosmic muon rate. This was corrected on March 3rd, 2022, when the maximum number of recorded AFT events in 20 ms was raised from one to three. For the data until this point, we therefore suffer from an efficiency of $\sim 70\%$ due to this effect. An overview of the changes in data sample over the SK-VI period is given in Table 6.1.

SHE threshold	AFT duration	E_{thr}	livetime	max AFT/20 ms
60 hits	500 μs	8 MeV	305.6 days	1
52 hits	500 μs	7 MeV	11.0 days	1
60 hits	500 μs	8 MeV	157.5 days	1
60 hits	500 μs	8 MeV	78.2 days	3

TABLE 6.1: Variation of the SHE trigger threshold and AFT event duration over time in the SK-VI data sample. The maximum number of AFT triggers in a 20 ms window is also indicated.

IBD simulation

The strategy for the simulation of IBD interactions follows the same blueprint shown in Section 3.1 and updated for the Gd-doped water in Section 5.1. Namely, it is an SKG4-based simulation in which we input the final-state positron and neutron from IBD interactions in the desired energy range. Since the study in Chapter 5, the simulation has been tuned with up-to-date parameters from recent measurements to better reflect the state of the detector during the SK-VI period. These tunings include parameters such as the water transparency, detector asymmetries, and the quantum efficiency of the PMTs. The noise is, as usual, taken directly from data for an accurate model of the accidental coincidences. The noise is taken from randomly-triggered data samples in the $[85220, 87220]$ range. Since the objective will be to detect the neutron captures on Gd, which occur on a shorter timescale than those on H, the neutron search will have to start very soon after the prompt event trigger. Therefore, the entire simulated event is contaminated with noise from data. This is in contrast

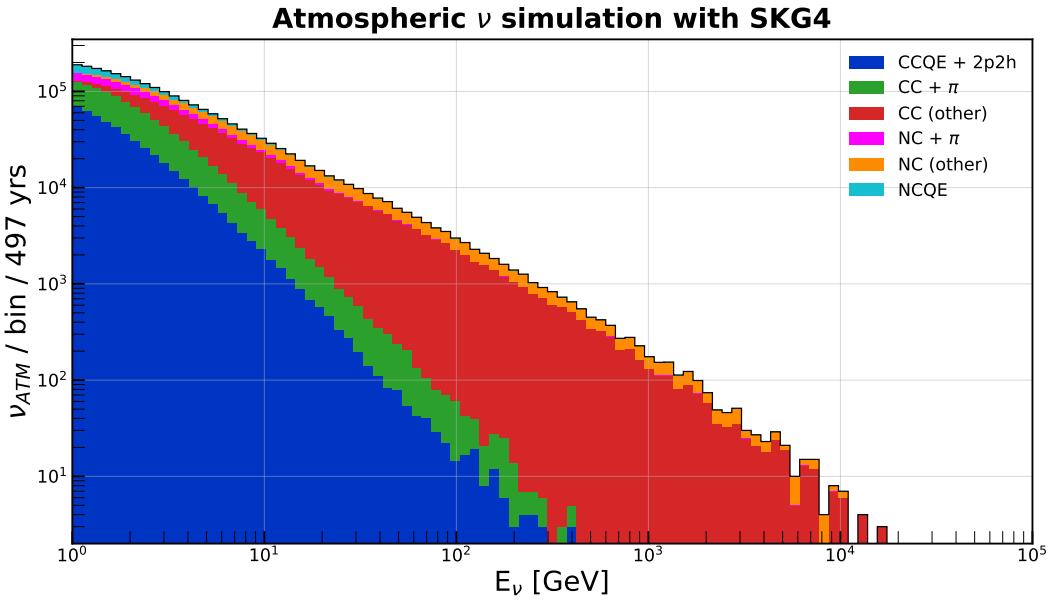

FIGURE 6.1: Simulated atmospheric neutrino events at SK, separated into the different types of interactions undergone in the water, corresponding to 497 years of simulated livetime.

to SK-IV, in which a relatively large portion of the event (before the neutron search window) was contaminated with simulated noise.

Atmospheric neutrino simulation

A new simulation of atmospheric neutrino interactions was also produced for SK-VI. For the SK-IV analysis, we were able to use the official atmospheric neutrino simulation produced by the collaboration. However, this was not possible for SK-VI. This is because the official atmospheric simulation uses simulated noise to contaminate the prompt event, lasting over 10 μs. While this was not a problem in the SK-IV analysis, it was incompatible with the fully-data-driven approach of the SK-VI simulation. Therefore, an atmospheric neutrino simulation was produced to only use noise from data. As before, the Honda atmospheric neutrino flux [62], calculated for the location of SK, was used. The resulting neutrino interactions in the detector are calculated by the NEUT package [63]. The flux of atmospheric neutrinos interacting in the tank is plotted in Figure 6.1. The detector response was simulated with SKG4. The entire event is then contaminated with noise, mirroring the IBD simulation. Neutrino energies up to several tens of TeVs are simulated. While the energy range of our analysis won't extend beyond 100 MeV, highly-energetic atmospheric neutrinos can cause prompt events that are reconstructed at low energies, for example through the reconstruction of a secondary particle. This is particularly relevant for the simulation of neutral current backgrounds, whose signature is driven by secondary particle emission.

Solar neutrino simulation

A new simulation was also used to model the behavior of solar neutrino backgrounds at SK-VI. These samples were provided by the solar neutrino working group at SK. They model the spectrum of ^{8}B and *hep* solar neutrinos, using the Winter [85]

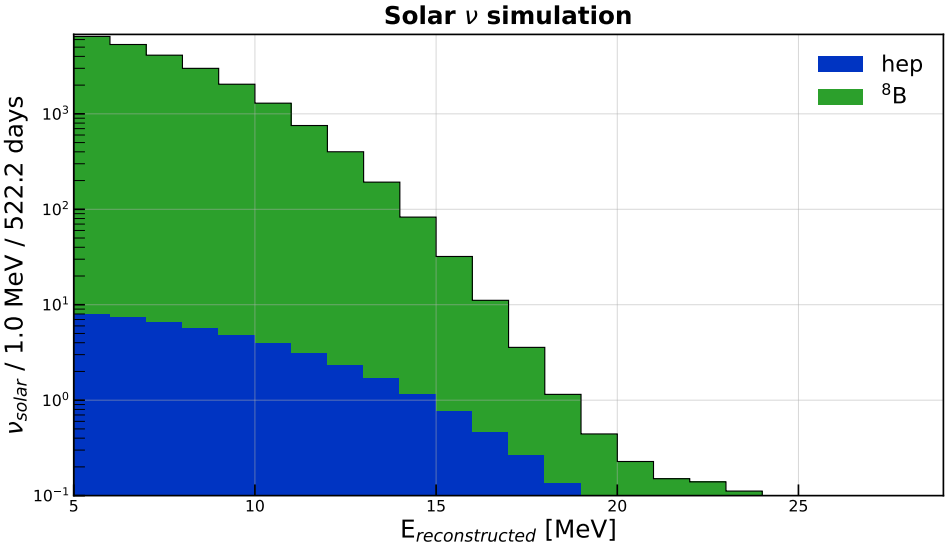

FIGURE 6.2: Reconstructed energy spectrum of simulated solar ^{8}B and *hep* neutrinos for SK-VI. About ~ 15 solar neutrino events are expected above 16 MeV.

and Bahcall [57] theoretical predictions, respectively. Neutrinos from these two production processes will be the only ones producing a visible background in our search. The simulated spectrum of solar neutrinos, after event reconstruction, is shown in Figure 6.2.

6.2 Prompt event selection

The BONSAI reconstruction algorithm is applied to the data and simulation samples, reconstructing the events' energy and interaction vertex. The first step of the data reduction procedure applied after reconstruction, as in Section 3.3. Non-physical events and events associated with muon vetoes from the OD are removed. Events with high total photoelectrons or high average charge deposition per PMT are also rejected. Poorly-reconstructed events are removed, as well as events outside the fiducial volume, once again defined by taking events with a minimal distance from the wall of 2 m. This first set of cuts is summarized in Table 6.2. At this stage, cosmic muon events before and after each DSNB candidate event are also independently fitted. This allows for the rejection of events consistent with muon-induced radioactive spallation backgrounds. Events occurring within 50 μs of a muon are rejected. The determination of the cuts to remove these backgrounds was not part of this work, but the strategy is as described in Appendix A and Ref. [49]. The resulting signal and spallation efficiencies in the energy range covered by this search are listed in Table 6.3.

6.2.1 Positron characterization

The third set of reduction cuts is aimed at the characterization of the IBD positron event. First, these cuts remove events that are inconsistent with a single positron (or electron) Cherenkov ring. Here, the relevant backgrounds are from atmospheric neutrino interactions. The Cherenkov opening angle is a key observable here, allowing us to separate electron-like events heavier particles (muons and pions) and from the signature of neutral-current interactions. The Cherenkov opening angle for the

Cut	Value	Apply to simulation?
Physics event selection	/	
OD trigger	veto	
Total photoelectrons	< 1000	
Time after muon	$> 50\ \mu$s	
$N_{Q<0.5}/N_{tot}$	< 0.55	
d_{wall}	> 2 m	✓
BONSAI goodness	> 0.5	✓
E_{rec}	$[8, 100]$ MeV	✓

TABLE 6.2: First set of cuts applied to the data and simulated samples. Some cuts only address backgrounds that are not present in simulation, and have negligible impact on the signal, so they are only applied to the data sample.

Energy bin [MeV]	16-18	18-20	20-24
Signal efficiency	73.3%	78.3%	85.6%
Spallation rate	0.013	<0.001	<0.001

TABLE 6.3: Signal efficiencies and rate of remaining spallation backgrounds associated with the spallation cuts in the energy region of this search. Above 24 MeV, the spallation background contribution is negligible.

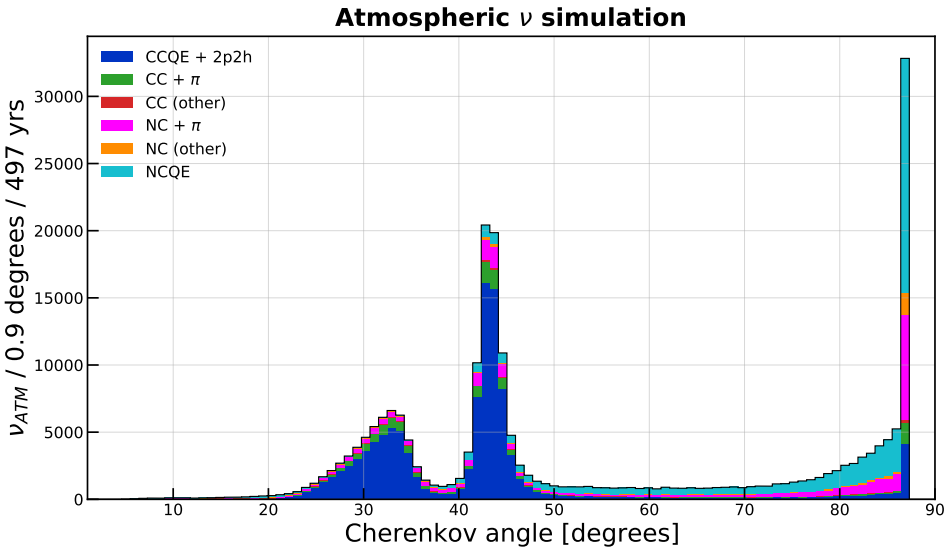

FIGURE 6.3: Cherenkov opening angle for the simulated atmospheric neutrino sample in SK-VI, after goodness cuts and mean charge deposition cut. Here the three-peak structure of the distribution is clearly visible, dominated by muons/pions, electrons/positrons, and NC interactions at low, medium, and high Cherenkov angles, respectively.

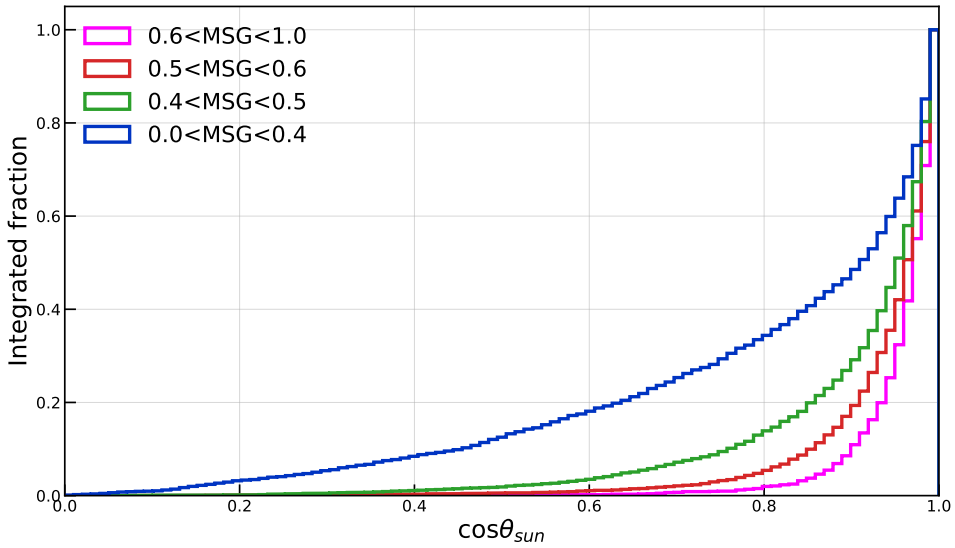

FIGURE 6.4: Cumulative distribution function of $\cos\theta_{\text{sun}}$ in the SK-VI solar neutrino simulation sample, divided into the MSG bins used for the solar neutrino cut. Solar neutrino events are highly peaked around $\cos\theta_{\text{sun}} = 0$, but events with lower MSG values have poorer angular resolution. Here, only events with $E_{\text{rec}} > 16$ MeV are considered. At lower energies, electrons scatter more, leading to a lower resolution. True DSNB events are uncorrelated with the direction of the sun and are therefore uniformly distributed in $\cos\theta_{\text{sun}} \in [-1, 1]$.

SK-VI atmospheric neutrino simulation is shown in Figure 6.3. Further rejection of high-momentum particles is achieved with a cut on the ring clearness, R_{pion}. Events with unusually-high PMT activity are also removed. First, this entails a cut on the mean charge deposition per PMT in 50-ns window, q_{50}/n_{50}. Then, PMT activity is looked for before and after the main activity peak of the event, to remove events with a double-peak structure. This translates to a cut on the maximum number of PMT hits in a 15-ns window preceding the event, $N_{\text{pre}}^{\text{max}}$, and in the number of electron-like rings following the event, $N_{\text{decay-e}}$. These cuts are tuned by using the atmospheric neutrino simulation. Finally, two directionality cuts are implemented. First, the d_{eff} cut removes events that originate from the detector wall, by measuring the distance to the wall along the reconstructed track direction. This cut is aimed at the removal of radioactivity from the walls and the rock surrounding the detectors, which can still be present even after the fiducial volume cut. It is tuned with data, in the same way as described in Section 3.4. The same energy-dependent cut criterion that was used in SK-IV was found to be suitable.

The direction of the DSNB candidate is also used to reject possible solar neutrino events. The visible solar neutrino flux will be made up of ν_e events. At SK, solar neutrino interactions therefore produce mainly electrons that are strongly forward-peaked. We can then use the angle θ_{sun}, between the DSNB candidate and the current position of the Sun, to identify possible solar neutrinos. The angular resolution of our measurement of θ_{sun} is highly dependent on the events' reconstruction quality. In particular, low-energy electrons tend to undergo multiple scattering as they emit Cherenkov radiation, leading to less clear rings and higher uncertainty in the event direction. The Multiple Scattering Goodness (MSG) variable was devised for the search of solar neutrinos at SK, and parametrizes this effect. DSNB candidates with higher MSG will therefore have undergone less scattering, and have lower uncertainty in their directionality. The solar neutrino cut is then implemented in bins

		Maximum value of $\cos\theta_{\text{sun}}$			
Energy bin [MeV]		$16-17$	$17-18$	$18-19$	$19-20$
MSG	$0.0-0.4$	0.05	0.35	0.45	0.93
	$0.4-0.5$	0.39	0.61	0.77	0.93
	$0.5-0.6$	0.59	0.73	0.81	0.93
	$0.6-1.0$	0.73	0.79	0.91	0.93

TABLE 6.4: Solar neutrino cut on $\cos\theta_{\text{sun}}$. A maximum value for $\cos\theta_{\text{sun}}$ is used depending on the reconstructed energy of the event, and on the value of MSG.

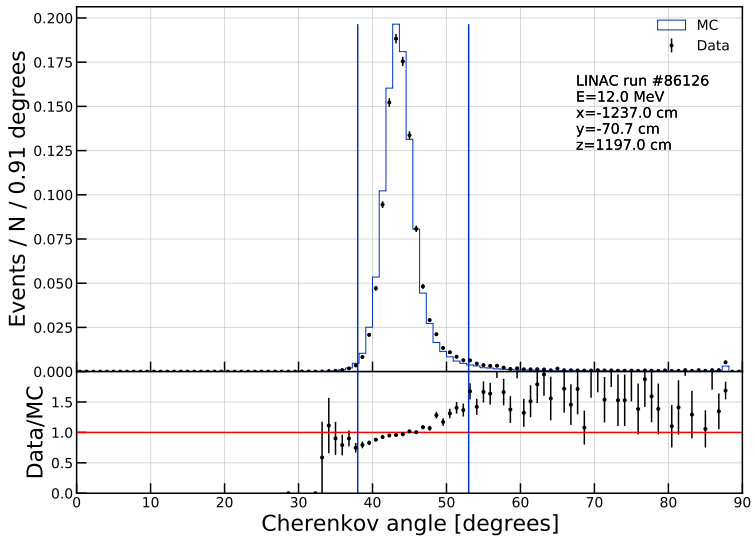

FIGURE 6.5: Distribution of the Cherenkov angle in the LINAC calibration samples, for simulation and data, using run 86126. The two-sided cut on the Cherenkov angle is displayed as the two vertical bars.

of reconstructed energy and MSG. For each bin, a cut is placed on the value of $\cos\theta_{\text{sun}}$. The distribution of $\cos\theta_{\text{sun}}$ for the MSG bins considered are shown in Figure 6.4, and the chosen cut values are displayed in Table 6.4. This cut has an efficiency of 93.4% above 16 MeV, leaving < 0.4 solar neutron events in this region over the SK-VI livetime. Solar neutrinos will also be efficiently removed by requiring the detection of a neutron in coincidence with the prompt event. Therefore, the solar neutrino cut is only applied to events not passing the neutron tagging cut, which will also be analyzed in this search.

6.2.2 Systematic uncertainty on efficiency

The systematic uncertainty associated with the modelling of the IBD positron's Cherenkov signature can be evaluated using calibration samples. Specifically here we use data samples from the LINAC calibration in SK. These calibration samples are produced by injecting mono-energetic electrons in the detector using a linear accelerator. The electrons are injected at different locations in the SK tank, and at different beam energies. A dedicated simulation sample is used for each energy and position, modelling the detector response with SKG4. The samples are run through

Cut	Value	Signal eff.
Preliminary reduction	Various	> 99 (1) %
θ_C	$\in [38°, 53°]$	95.0 (1.3) %
R_{pion}	< 0.37	98.7 (0.6) %
q_{50}/n_{50}	< 2.0	97.9 (0.3) %
$N_{\text{pre}}^{\text{max}}$	< 12	> 99 %
$N_{\text{decay-e}}$	< 1	> 99 %
d_{eff}	Energy-dependent	98.4 (0.1) %
θ_{sun}	MSG-dependent	93.4 (1.2) %

TABLE 6.5: Preliminary reduction and positron characterization cuts, along with signal and background efficiencies. Note that the solar neutrino cut is only applied to events not passing the neutron tagging cut.

the same particle reconstruction and positron characterization steps as above. The systematic uncertainty on the efficiency for a given cut is then defined as the maximum discrepancy between data and simulation among the samples. The distribution of the Cherenkov angle for one of the LINAC samples is shown in Figure 6.5. Here, a discrepancy was observed, driven at least in part by an excess of events in data for high Cherenkov angles. These high angles could be explained by electrons escaping the LINAC beam pipe before the endcap. Nevertheless, the cut is loose enough that this effect has a limited impact on the uncertainty estimation. The chosen cut points for the prompt event selection, except for the spallation cuts, are summarized in Table 6.5, along with their efficiencies and the associated systematic uncertainties.

6.3 Neutron tagging

The final step in the event selection is the detection of the IBD neutron in coincidence. This is, of course, the crucial step differentiating the SK-VI search from previous searches, gaining additional efficiency compared to SK-IV thanks to the Gd-doped water. As before, neutron capture candidates are defined as groups of PMT hits clustered in time-TOF from the prompt vertex. A number of observables is then defined to characterize the neutron capture's vertex, its Cherenkov signature, and its noise-like qualities. A BDT is then trained to combine these observables into a single output, ranging from 0 to 1, parametrizing the likelihood for a given hit cluster to be associated with a true neutron capture. A different neutron tagging cut on the BDT output is chosen according to the energy of the prompt event, obtaining the signal efficiencies shown in table 6.6. The details of the neutron tagging algorithm have already been discussed in Section 3.5. In addition, the modifications to the algorithm for SK-VI and its performance on SK-VI simulation was also already shown in Section 5.2. Therefore, in this section we will limit ourselves to discussing the systematic uncertainty on the algorithm's efficiency, which is evaluated using the new SK-VI calibration data samples.

6.3.1 Systematic uncertainty on efficiency

The general strategy outlined in Section 3.5.4 is followed to evaluate the performance of the neutron tagging algorithm on real data. In particular, the efficiency of the algorithm is subject to systematic uncertainties associated with the modelling of the neutron capture signal on both hydrogen and gadolinium. As in the SK-IV analysis,

Energy [MeV]	16-18	18-20	20-22	22-24	24-26	>26
Signal eff.	60.4	60.6	59.8	59.8	60.8	60.0
Acc. rate	0.020	0.027	0.032	0.032 4	0.023	0.020

TABLE 6.6: Signal efficiencies and accidental coincidence rate for the neutron tagging cut in different energy bins. The given signal efficiencies are obtained by requiring exactly 1 neutron to be tagged.

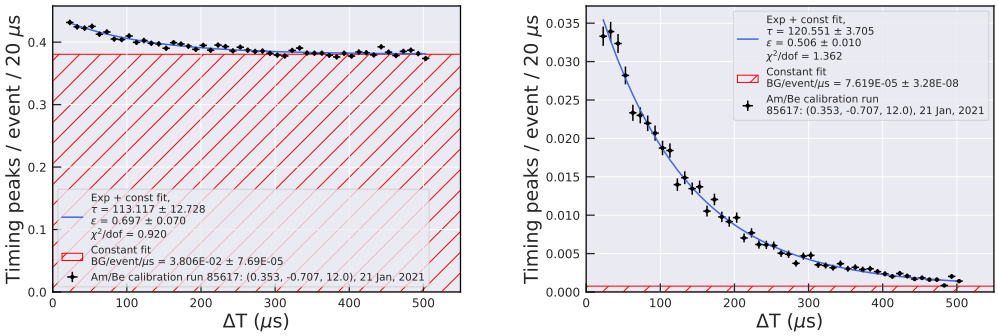

FIGURE 6.6: Time distribution of neutron candidates in an Am/Be calibration data sample, from run 85617 in SK-VI. Left: after neutron candidate preselection. Right: after a BDT cut of 0.9. The data is fitted with an exponential (blue line) on top of a constant background (red hatched area).

a radioactive source, containing $^{241}\mathrm{Am}$ and $^{9}\mathrm{Be}$ and enclosed in a BGO scintillator, is placed at different points in the detector. For this study, we use the data runs 85613, 85615, and 85617, which were taken with the Am/Be source at $z = -12$ m (near the bottom of the tank), at the center of the tank, and at $z = 12$ m (near the top of the tank).

The decay chain of the Am/Be source, as described in Section 3.5.4, produces a scintillation prompt event, activating the SHE trigger, which is followed by a neutron capture signal. Here, the prompt event is selected by requiring an SHE trigger in which the number of detected photoelectrons (p.e.) verifies the condition $750 < p.e. < 1050$. To ensure well-separated events, 1.5 ms are required to separate each trigger. In addition, to prevent any scintillation light from contaminating the neutron search window, we reject events in which more than 70 PMT hits are found in a 200-ns window over the entire event. The neutron tagging procedure detailed in Section 5.2 is applied, with the application of the BDT on the tagging observables. The number of neutrons in data can be estimated by using the distribution of the time difference ΔT between neutron candidates and the respective prompt event. This distribution is fitted with an exponential + constant function, as in Equation 3.50, to extract the remaining neutron events after a given cut. Figure 6.6 gives an example of such a fit. In this fit, we can also observe an exponential time constant (which is allowed to vary in the fit), measured as 120.6 ± 3.7 μs. This measurement is consistent with half the neutron captures occurring on Gd ($\tau_{\mathrm{ncap}} \sim 30$ μs) and half on H ($\tau_{\mathrm{ncap}} \sim 200$ μs). With the remaining number of neutrons obtained from the fit, we can then compute the efficiency of neutron tagging in data.

Dedicated simulation samples are also produced with SKG4, to appropriately compare the simulated detector response to the signal. Here, the scintillation light itself was not simulated, as in the previous analysis. Rather, IBD events are generated at the source positions of each Am/Be data run, matching the detector conditions

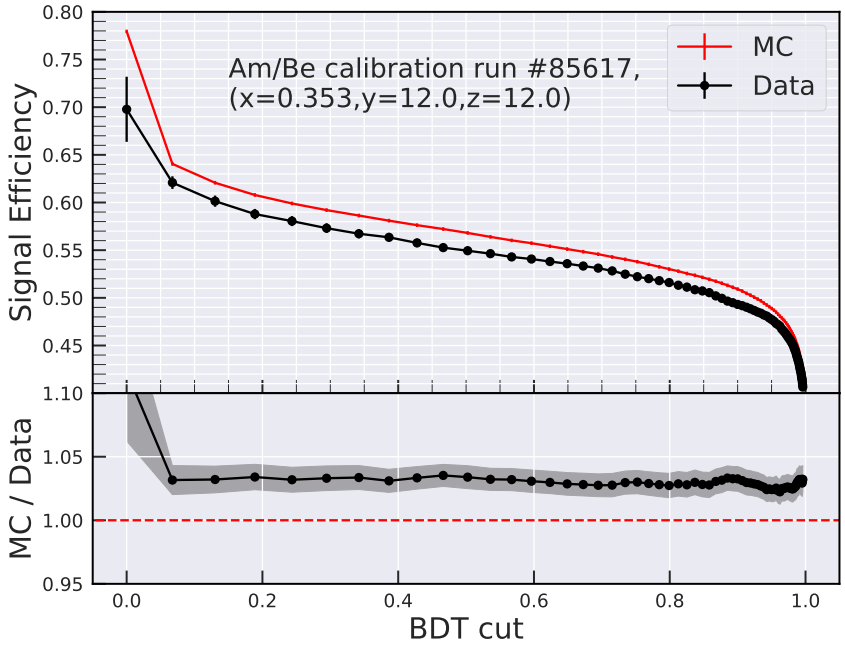

FIGURE 6.7: Signal efficiency of the neutron tagging algorithm as a function of the cut on the BDT output, for the Am/Be calibration data sample from run 85617 in SK-VI. The efficiency is calculated for simulation (red) and data (black). The bottom panel shows the simulation/data ratio used to estimate the systematic uncertainty on the efficiency. The shaded area indicates the statistical uncertainty on the ratio. A BDT cut of 0 indicates the efficiency at preselection. At this point, the uncertainty associated with the exponential fit is high, inflating the difference between data and simulation.

at the time of data-taking. The 70-hit maximum condition applied to data is also applied to the simulation to avoid biasing our comparison. The neutron tagging algorithm is applied as above, and the efficiency is calculated directly from the truth information from the simulation. The signal efficiency can then be compared between data and simulation for a range of possible cuts on the neutron tagging BDT output. An example of this comparison is plotted in Figure 6.7. This comparison is repeated for the Am/Be samples used in this study, and the results are used to estimate the systematic uncertainty on the efficiency. By taking the data run with the maximum observed discrepancy, a systematic uncertainty of 5% is assigned to the neutron tagging efficiency.

6.4 Spectral analysis

The statistical approach described presented in Chapter 4 is used for this analysis. The reader can refer back to this chapter for a more detailed treatment. The spectral fit is applied in the same way to the final data sample, for events passing all the reduction steps: preliminary reduction, spallation cuts, prompt event characterization, and neutron tagging. As before, the sample is subdivided into six fitting regions: three Cherenkov angle regions, and two neutron tagging regions according to the number of neutrons N_{ntag} identified by the BDT. We will then perform an unbinned, maximum likelihood fit on the data based on the expected spectra function of the remaining backgrounds and the DSNB signal. In this analysis, we concern ourselves with the energy region above 16 MeV. At these energies, the energy spectrum is dominated

		Cherenkov angle		
		$20° - 38°$	$38° - 53°$	$78° - 90°$
N_{ntag}	0 or >1	μ/π	e^+/e^-	NC
	1	$\mu/\pi + 1n$	$e^+/e^- + 1n$	NC+$1n$

TABLE 6.7: Overview of the fitting regions used in this analysis. We subdivide the parameter space according to the Cherenkov angle and the number of tagged neutrons N_{ntag}. Regions with small and large Cherenkov angles are dominated by events with visible muons and pions, and by NC atmospheric events respectively.

by atmospheric neutrino backgrounds. The low Cherenkov angle region contains events with smaller Cherenkov cones, from high-momentum particles, and will be dominated by visible (above Cherenkov threshold) muons and pions produced in atmospheric neutrino interactions. The high Cherenkov angle region will be dominated by neutral current atmospheric interactions. The signal Cherenkov angle region contains events whose Cherenkov rings are consistent with electrons and positrons, and will be dominated by charged-current atmospheric neutrino interactions. In particular, the dominating background in this region is from invisible muons created in charged-current ν_μ interactions, decaying into a visible electron. Atmospheric backgrounds accompanied by a single secondary neutron will contaminate the signal neutron tagging region, in which a single neutron is found.

For the purposes of the fit, six spectral categories are defined in the fitting energy region. The dominant category contains electrons from the decay of muons and pions (otherwise known as *decay electrons*). Its background shape is obtained directly from data samples of decay electrons at SK. Simulation is only used to estimate the proportion of this background occupying each fitting region. Subdominant in the Cherenkov angle signal region, we find the CC ν_e background category. Muons and pions with no associated visible decay electrons are assigned another category. Events from atmospheric neutral current events make up the last category of atmospheric backgrounds. The spectral shapes of the CC ν_e, μ/π, and NC categories are estimated with simulation. The expected contribution of the atmospheric backgrounds after cuts is shown in Figure 6.8. A piecewise polynomial smoothing is applied to obtain a smooth PDF for each atmospheric background spectrum. A further category, is defined, in the same way as in Section 4.1 and Ref. [49], to capture the spectral shape of the remaining spallation in the sample. These backgrounds are made up of the electrons or positrons from the β decay of ^{8}B, ^{8}Li, and ^{9}C. Finally, the DSNB spectral shape is determined from the IBD simulation, and depends on the flux model considered. Each PDF is normalized to 1 over the six fitting regions.

After defining the PDFs, f_j, capturing our understanding of the spectral categories j, we proceed with the maximum likelihood fit. We perform a simultaneous fit on all six parameter space regions, on the data sample $\vec{E}$. The number of events N_j in each category if fitted, treating the five background categories as nuisance parameters. We therefore aim to find the value of $\vec{N} = \{N_j\}$ that maximizes the following likelihood function:

$$\mathcal{L}(\vec{N};\vec{E}) = e^{-\sum_j N_j} \prod_{i=1}^{N_{\text{data}}} \sum_j N_j \, f_j(E^i;\theta^i_{\text{C}}, N^i_{\text{ntag}}). \tag{6.1}$$

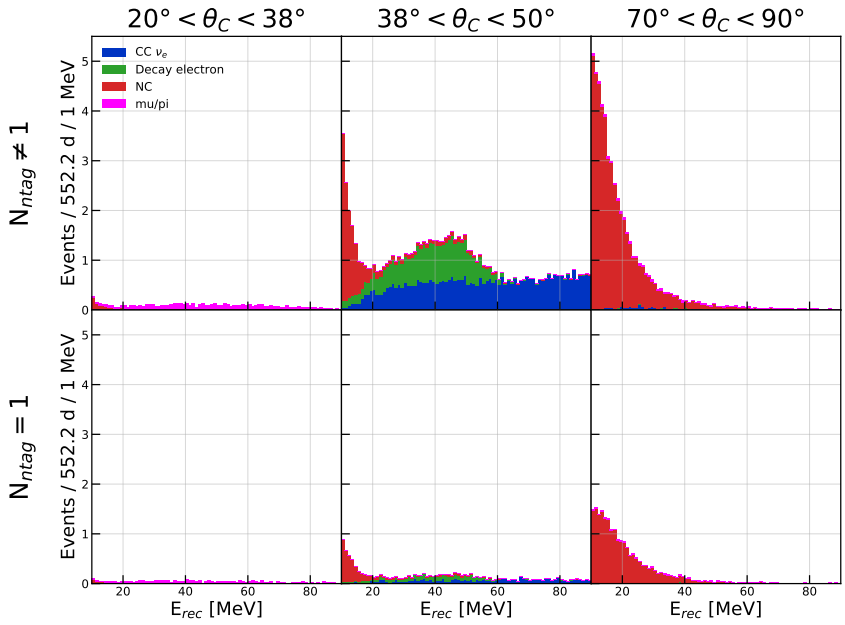

FIGURE 6.8: Atmospheric neutrino backgrounds to the DSNB search, after cuts, in SK-VI, here, we subdivide the backgrounds into the four spectral categories used in the spectral fit. Note that here the minimum energy in this plot is 10 MeV, below the threshold of this analysis.

The systematic uncertainty associated with the spectral form of the signal and backgrounds are accounted for as spectral distortions of the PDFs. This includes distortions associated with energy scale and resolution uncertainties, the spectral shapes of the CC ν_e, NC, and spectral categories, as well as uncertainty on the neutron multiplicity. Energy-independent systematic uncertainties enter as uncertainties in the signal efficiency, and affect the conversion from the number of observed events to the event rate and flux. In particular, these are the uncertainties associated with the fiducial volume cut, the livetime computation, and in the theoretical IBD cross-section, amounting respectively to 1.5%, 0.1%, and 1%. More details about the systematic uncertainties are found in Section 4.3. Note that in this analysis, the magnitudes of the systematic uncertainties were left unchanged from the previous search.

6.4.1 Results

The spectral fit is applied to 552.2 days of the SK-VI dataset, with a lower energy threshold of 16 MeV, for a range of DSNB flux models. The combination of the results from this analysis with the previous one will bring the total exposure up to 22.5×6375 kton·day. In the final data sample, we find 79 events, of which 63 in the Cherenkov angle signal region. Of these 63 events, 12 are found to contain a single neutron. Only three events in the Cherenkov angle signal region were found to contain 2 neutrons, while no neutrons were found for the remaining events. An example of a spectral fit on the SK-VI data sample is plotted in Figure 6.9. As expected, the decay electron represents the main contribution in the Cherenkov signal region. At high Cherenkov angles, NC backgrounds are found to dominate. A crucial difference from the previous analysis, however, is the distribution of the signal across the two

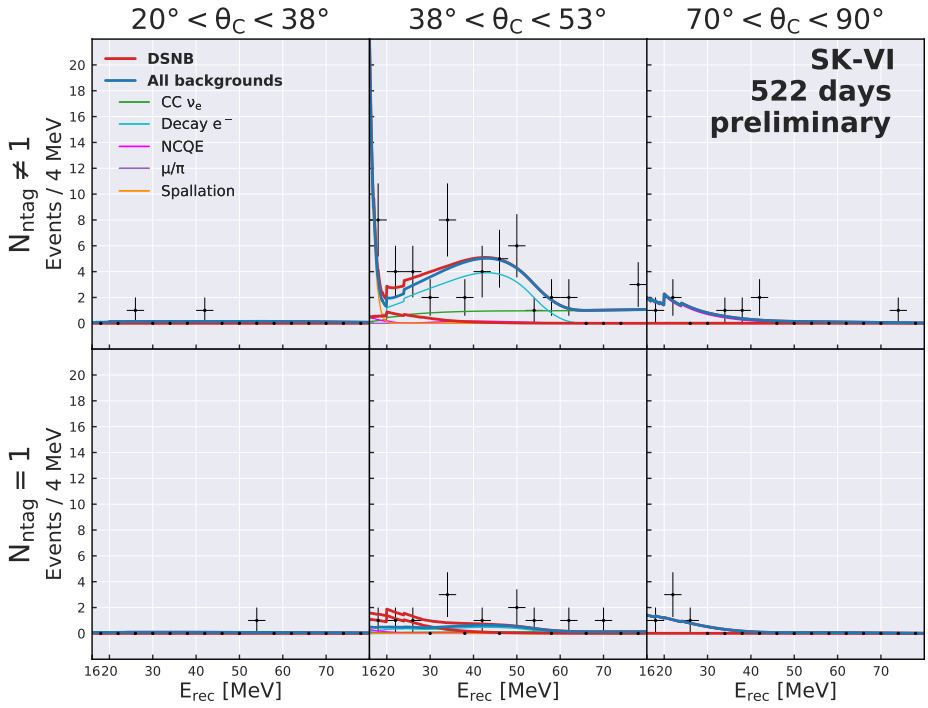

FIGURE 6.9: Best-fit signal and background spectra (solid lines) for the SK-VI period, assuming the DSNB flux predicted by the Horiuchi+09 model [20], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

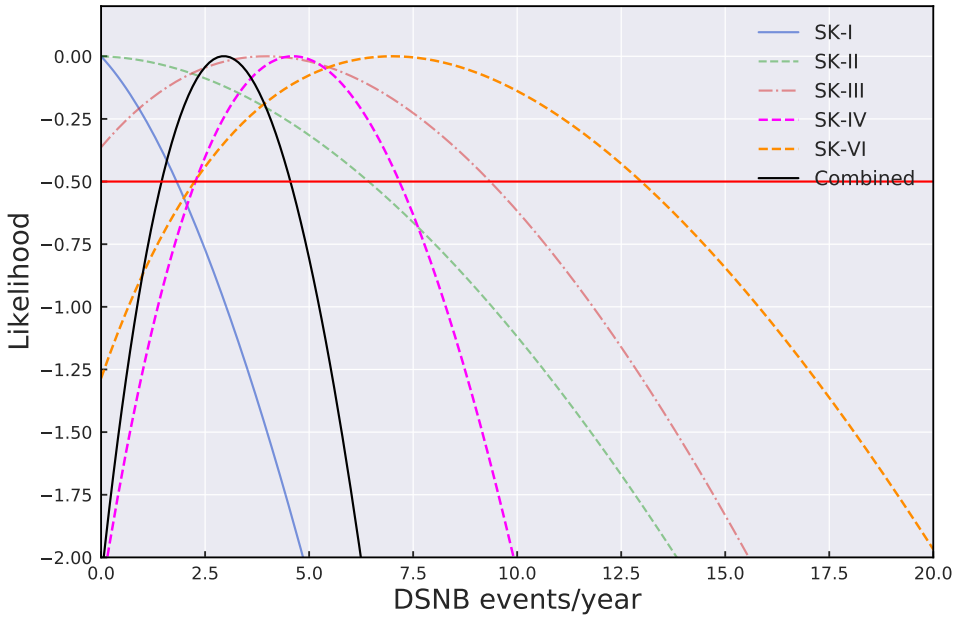

FIGURE 6.10: Likelihoods associated with phases I, II, III, IV, and VI of SK, as well as the combined likelihood, for the fit using the Horiuchi+09 model [20].

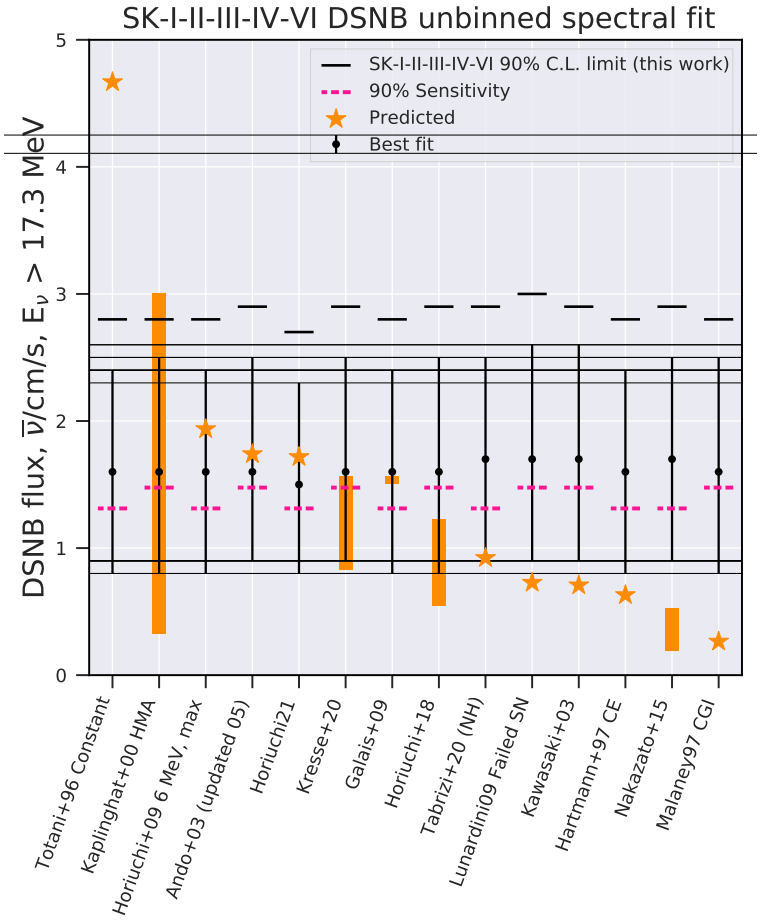

FIGURE 6.11: Spectral fit results for the models considered in this work. The black dots and vertical lines are the best-fit DSNB model fluxes and uncertainties, while the solid black horizontal lines are the 90% C.L. upper limits on the flux. The dashed magenta lines give the 90% C.L. sensitivity of the search for each model. These values are compared to the theoretical flux predictions of each model (orange stars and bars). A single star denotes a singular flux prediction, while a bar denotes a range of flux predictions.

neutron tagging regions. Thanks to the higher neutron tagging efficiency afforded by SK-VI, we now expect more than half of the DSNB events to end up in the $N_{ntag} = 1$ region, which suffers much lower background contamination. Overall, the SK-VI fit that uses the Horiuchi+09 model finds 32 decay electron events, 19 CC ν_e events, 14 NCQE events, 3 μ/π events, and 4 spallation events. The data sample is found to be particularly signal like, leading, for this model, to a best-fit rate of $7.0^{+5.9}_{-4.7}$ IBD/22.5 kton/year. The corresponding best-fit flux is $3.9^{+3.3}_{-2.6}$ $\overline{\nu}_e/\text{cm}^2/\text{year}$. As before, these values refer to the DSNB rate and flux above a neutrino energy $E_\nu > 17.3$ MeV, reflecting the energy threshold of the analysis, and allowing for straightforward comparison with the previous searches. Note that the fit is dominated by a large statistical uncertainty due to the limited extend of the livetime. In particular, the nominal 552.2 days of livetime does not reflect the inefficiency of the trigger scheme mentioned in Section 6.1. In the energy range of the DSNB, this results in a $\sim 70\%$ efficiency from the trigger scheme alone, leading to an effective livetime of a little over a year.

The SK-VI result is then combined with the search presented in Chapter 4 by summing together the marginal likelihood functions from Eq. 4.21, as in Eq. 4.22.

The marginal likelihood functions, as well as the combined likelihood of the entire SK dataset, are plotted in Fig. 6.10. Because of the excess observed in the SK-VI data sample, the combined best-fit increases when compared to the previous combined analysis. Specifically, this leads to a combined best-fit DSNB rate of $2.9^{+1.6}_{-1.4}$ IBD/22.5 kton/year, and a combined best-fit DSNB flux of $1.6^{+0.8}_{-0.8}\,\overline{\nu}_e/\mathrm{cm}^2/\mathrm{year}$. This represents a 2σ excess in the data. For the same reason, the observed upper limits of the combined analysis are also higher. Using the formula in Eq. 4.24 to compute the limits, we observe a 90% upper limit of 5.1 IBD/22.5 kton/year on the DSNB interaction rate, and of $2.9\,\overline{\nu}_e/\mathrm{cm}^2/\mathrm{year}$ on the DSNB flux.

The best-fit and 90% C.L. upper limit results for the models studied in this analysis are listed in Table E.1, along with the predicted integrated fluxes, and are visualized in Fig. 6.11. A similar plot for the corresponding best-fit event rates at SK is given in Fig. E.1. The flux excess found by the combined observation is of $1.6\,\overline{\nu}_e/\mathrm{cm}^2/\mathrm{year}$ for most models. While the combined 90% C.L. upper limits got higher due to this excess, the same models that were excluded by the previous analysis are also excluded here. Despite these looser limits, this preliminary search demonstrates improved sensitivity compared to the previous searches. Focusing on the Horiuchi+09 model, we observe a 90% C.L. sensitivity on the flux—here defined by multiplying the 1σ uncertainty by 1.64—of $5.4\,\overline{\nu}_e/\mathrm{cm}^2/\mathrm{year}$ for the SK-VI period. In comparison, the SK-IV search, with 7.2 times the (effective) livetime, has a sensitivity that is only 2.3 times better. To reach the same sensitivity with the SK-VI analysis, we would therefore need $\sim 2.3^2 \approx 5.3$ times the effective livetime of SK-VI. This preliminary analysis therefore requires roughly 30% less exposure to reach a similar sensitivity as an SK search before Gd loading.

Further improvements to this analysis are possible. First among these would be to attempt lowering the energy threshold of this search. The difficulty in lowering the threshold lies in the treatment of the spallation backgrounds. Below 16 MeV, spallation backgrounds quickly ramp up and begin largely dominating the energy spectrum. Due to our limited ability to model these backgrounds, a spectral analysis would require a very strong background suppression. With SK-Gd, it should be possible to reach the necessary levels of spallation background suppression (through neutron tagging) while also being able to retain a meaningful portion of signal. Additional study of the spallation backgrounds at these energies is necessary, so it was not attempted in this work. The lowering of the energy threshold would fully realize the potential of SK-Gd's DSNB search capabilities. Additionally, the collaboration is currently investigating several sources of systematic uncertainty. In particular, studies about the reduction of the NC backgrounds and the associated spectral uncertainty—which would be particularly important if we do lower the threshold—are underway. The neutron multiplicity of atmospheric backgrounds is also under study. Nevertheless, these first results from SK-Gd, showing improved sensitivity and the hint of a signal, provide us with a good outlook for the future of the experiment.

Conclusion and future directions

The essential object of the work presented in this thesis was to further the efforts towards the possible future detection of the Diffuse Supernova Neutrino Background with the Super-Kamiokande detector. This goal was achieved in several ways. First, an analysis searching for the DSNB using Super-Kamiokande was presented. Using the inverse beta decay interaction channel, DSNB neutrinos were looked for in the detector's 22.5 kton fiducial volume of water using the SK-IV data period. Within this analysis, the strategy for identifying the IBD positron and, in particular, the coincidence neutron signal was detailed. Neutron tagging, possible for the first time in SK starting with SK-IV, is a hallmark of this search. This work used a BDT model to maximize our ability to detect the faint 2.2 MeV neutron capture signal on Hydrogen. The neutron tagging efficiency obtained for neutron captures in SK-IV's pure water was between $20-30\%$, for an accidental background reduction between 3 and 4 orders of magnitude. A dedicated statistical treatment, leveraging an unbinned spectral fit of the data, was developed to take advantage of the neutron signal, while also being able to produce a combined measurement using the entire SK dataset until SK-IV. A range of modern theoretical DSNB flux model were tested with a total exposure of 22.5×5823 kton·day. An excess of around $4.6^{+2.5}_{-2.3}$ IBD events/year (depending on the exact DSNB model) was found in the SK-IV dataset. Combining this observation with spectral fits on previous SK periods yielded upper limits on the DSNB flux for $E_\nu > 17.3$ MeV of around $2.6\ \overline{\nu}_e/\text{cm}^2/\text{year}$, at a 90% C.L. At the time of writing, this figure represents the most stringent limit yet on the DSNB flux. What's more, the sensitivity of the combined fit was found to be comparable to several modern, realistic flux models, offering an optimistic outlook for the future of the DSNB signal.

Beyond its immediate results, the significance of the SK-IV DSNB analysis with neutron tagging lies in relation to the future of the detector. With the advent of the Gd era of SK data-taking, the neutron capture signal was significantly enhanced by doping the detector water with Gadolinium, starting with a 0.01% concentration in SK-VI and reaching a 0.03% concentration in SK-VII, currently underway. SK is therefore poised to significantly increase its sensitivity to neutron signals, and in particular to the DSNB. The SK-IV analysis effectively lays the groundwork for the sensitive SK-Gd analyses, providing a blueprint for the strategy of the searches to come. The potential of the analysis in light of the Gd doping was directly assessed with new sensitivity studies. The improved ability of SK to detect neutron captures was tested, observing an up to eight-fold improvement in the signal efficiency for neutron captures on Gd, compared to captures on H, for realistic neutron tagging cuts. The way this efficiency gain maps onto different SK eras was studied, projecting the performance of the analysis onto the future of the experiment. This undertaking was the first study evaluating the sensitivity of SK-Gd to the DSNB that for a realistic analysis, fully accounting for the impact of systematic uncertainties. The findings showed that SK could reach the combined sensitivity of the SK-IV spectral analysis could be reached in just three years of running time, by mid-2023. By the end of the SK-Gd project, sensitivity of a combined SK analysis will likely span a significant portion of the parameter space, more than just the most optimistic flux models.

Realistic advances in our treatment of the systematic uncertainties will likely further improve upon these prospects. A first look at the early SK-Gd data was given with a preliminary analysis searching for the DSNB in the SK-VI data period. While the livetime was limited, an increased sensitivity was demonstrated, and the significance of the previously-observed excess was reinforced.

This thesis therefore hopes to prove as impulse for future efforts and studies on the road to a first detection of the DSNB. In SK, the avenues for further exploration are numerous. Already mentioned was the improvement of key systematic uncertainties in the analysis. In particular, lowering the systematic uncertainty associated with atmospheric NCQE backgrounds will significantly improve our sensitivity to the lower-energy range of the DSNB flux and allow us to fully capitalize on the higher neutron tagging efficiency of SK-Gd. Another important statistical uncertainty to master will also be that on the neutron multiplicity of the atmospheric backgrounds, which can limit our understanding of how these backgrounds responds to the neutron tagging cut. Starting with SK-Gd, the structure and strategy of the DSNB analysis could be revisited to further exploit the neutron signal. The performance of spallation cuts will likely improve along with the improved visibility of the secondary neutrons associated with the spallation muon's energy deposition. SK-Gd also opens the door for a dedicated treatment of the atmospheric backgrounds to attempt direct discrimination of the secondary neutrons originating from these backgrounds, instead of only discriminating IBD-like neutrons from accidental coincidences. The newly-available spallation simulation for SK will also drastically improve our ability to understand, model, and deal with spallation, and will be key to realizing the full potential of the experiment. Finally, more sophisticated algorithms such as deep-learning models could be applied to the search for a variety of goals, such as additional improvements in neutron tagging performance, better discrimination of NCQE backgrounds, and a combined reconstruction of the positron and neutron in the full IBD event. Many of these investigations are already underway within the collaboration. The prospects for the future of the DSNB search at Super-Kamiokande are therefore quite bright. This thesis is then one of the first bricks on a way ahead rich with possibilities.

A
Spallation modelling and cuts

The analysis in this thesis relies in no small part on our ability to remove spallation backgrounds from the data sample. In this section, a brief overview of the spallation reduction strategy is given. For a more detailed treatment, turn to Ref. [49]. As mentioned in the text, at the time the analysis was carried out, no simulation was available to directly model these backgrounds. This has since changed, notably thanks to the recent simulation work in Ref. [50]. Nevertheless, due to relative freshness of these results, they were not taken into account for this work. Updating the way we treat spallation in the DSNB analysis, in light of these results, will be the focus of dedicated studies within the SK collaboration. In this thesis, the modelling and reduction of the spallation backgrounds is therefore entirely data-driven.

The main principle guiding our strategy is the definition of two samples from the muon data. In Section 3.3, we described the fitting of muon tracks preceding and following the prompt event trigger of a DSNB candidate. To determine whether a DSNB candidate is consistent with a spallation event, the name of the game is to pair each muon with the associated DSNB candidate and analyze each pair. To develop appropriate cuts, the *pre* sample, defined as the sample of muon-candidate pairs for muon preceding the SHE trigger, is compared to the *post* sample, containing pairs with muons after the SHE trigger. The former sample contains the correlated spallation-causing muons and the respective spallation-induced DSNB candidate, as well as muons that are totally uncorrelated to the prompt event. The latter sample will instead only include uncorrelated muons. Comparing these two samples is what allows us to have an expectation for the behavior of the spallation backgrounds, in the absence of simulation. In general for a given spallation cut, the signal and spallation background efficiencies can be estimated from data as

$$\epsilon_{\text{signal}} = N^{\text{post}}_{\text{after cut}} / N^{\text{post}}_{\text{before cut}} \tag{A.1}$$

$$\epsilon_{\text{spall}} = (N^{\text{pre}} - N^{\text{pre}}_{\text{non-spall}})_{\text{after cut}} / (N^{\text{pre}} - N^{\text{pre}}_{\text{non-spall}})_{\text{before cut}}, \tag{A.2}$$

where N^{pre} and N^{post} are the number of events in the pre-sample and post-sample, respectively, and N^{pre}non-spall is the expected contribution of non-spallation backgrounds to the pre-sample. One of the limitations of this data-driven approach is that it suffer from significant statistical uncertainty at higher energies, where spallation backgrounds are rarer and the available data sample is smaller. Therefore, special care is required when estimating the performance of these cuts. The details and assumptions of the strategy are available in Ref. [49].

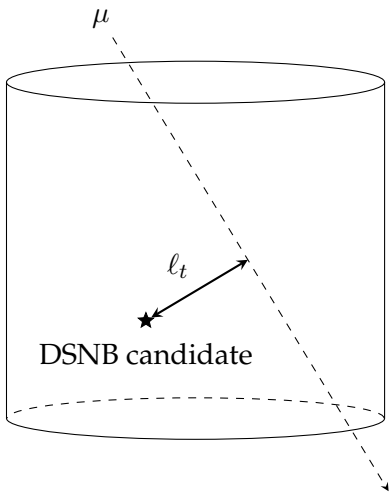

FIGURE A.1: Example of a muon-DSNB candidate pair, showing the calculation of the transverse track distance ℓ_t

A.1 Preselection

Two sets of preselection cuts are applied to remove spallation backgrounds from the most energetic muons, with little impact on the signal efficiency. First, a cut aimed at the removal of multiple spallation events from the same parent muon, since highly-energetic muons can induce more than one spallation event. To this end, DSNB candidates that are within 60 seconds and 4.9 m of another candidate below 25 MeV. This particular cut already removes 45% of the spallation background while retaining 98% of the signal, in the SK-IV data sample.

Secondly, DSNB candidates associated with so-called *neutron clouds* are removed. Spallation events are often accompanied by hadronic showers and the production of several secondary neutrons. Following the strategy described in Ref. [86], we search for neutron captures occurring within 500 μs after each muon event, using the WIT trigger [87]. This trigger (short for *Wideband Intelligent Trigger*) was developed for the detection of very low-energy events (< 3 MeV), and applies reconstruction software in real time to deal with the enormous event rate at these energies. If more than one neutron capture is found within 5 m of the muon track, a neutron cloud is defined, and DSNB candidates close in space and time to this cloud are removed. The WIT trigger was introduced during the livetime of SK-IV, so this cut was not applicable to the SK-IV entire dataset. Where possible, this cut retains virtually all signal while removing roughly 40% of the spallation background. Overall, accounting for the correlation between the two above cuts, preselection removes around 55% of the spallation background.

A.2 Rectangular cuts

Above 16 MeV, the spallation backgrounds are dominated short-lived isotopes. In this energy range, removal of spallation relies on a series of rectangular cuts applied to the time difference dt and the transversal distance ℓ_t between the DSNB candidate and each potential muon track (see Figure A.1). The dt distribution of spallation events reflects the half-life of the associated radioactive isotope. The transverse distance

associated with a spallation event depends on the path lengths of the secondary particles in the event. A typical spallation event will have an ℓ_t within a few meters.

The guiding principle behind these cuts is to remove spallation events by exploiting the correlation between the half-lives of the spallation isotopes and the energy of the event. For the purposes of the analysis in Chapters 3 and 4, these cuts are optimized together with the neutron tagging cut described in Section 3.5. In particular, the cuts are chosen such that the combined action of the spallation and neutron tagging cuts reduce the expected spallation contribution down to negligible levels. Two sets of cuts are chosen, one for each of the N_{ntag} regions of the spectral analysis presented in Chapter 4. The rectangular spallation cuts for the $N_{\text{ntag}} = 1$ and $N_{\text{ntag}} \neq 1$ are given in Tables A.1 and A.2, respectively. More stringent spallation cuts are necessary in the $N_{\text{ntag}} \neq 1$ to achieve levels of background removal comparable to those achieved in combination with neutron tagging.

Muon type	Energy range (MeV)	dt (sec)	ℓ_t (cm)
All	< 24	> 0.001 > 0.1	- > 400
Misfit	$[16, 20]$	> 1.5	-
Stopping		> 0.05	-
All multiple		> 0.05	-
Well-fitted single through-going	$[16, 18]$	> 7	> 150
Poorly-fitted stopping		> 6	-

TABLE A.1: Summary on the rectangular spallation cuts for events with exactly one tagged neutron. When a given cut has requirements on both dt and ℓ_t, at least one should be fulfilled.

Muon type	Energy range (MeV)	dt (sec)	ℓ_t (cm)
All	< 24	> 0.001 > 0.1	- > 400
Well-fitted single through-going	$[18, 20]$	> 7	> 100
Misfit	$[16, 20]$	> 1.5	-
Stopping		> 0.05	-
All multiple		> 0.05	-
Well-fitted multiple		> 0.1 > 1	- > 400
Well-fitted single through-going	$[16, 18]$	> 7	> 150

TABLE A.2: Summary on the rectangular spallation cuts for events with tagged neutrons different from one. When a given cut has requirements on both dt and ℓ_t, at least one should be fulfilled. Here tighter rectangular cuts are applied to the muon categories generating most of the spallation —single through-going and multiple muons— to remove contributions from short-lived isotopes.

B
SK-IV neutron tagging BDT

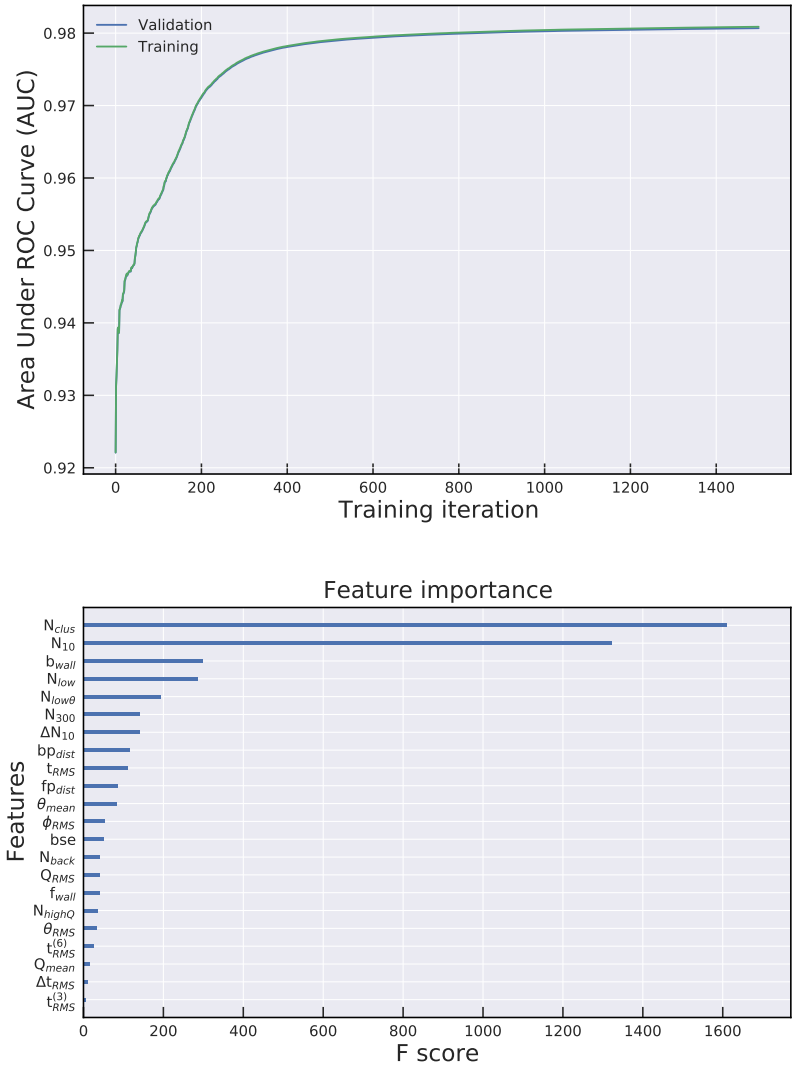

FIGURE B.1: Top: evolution of the model performance metric $(1 - \text{AUC})$ as function of the training iteration. The model improves sharply in the first 200 iterations, and begins plateauing after around 500 iterations. The performance on the training set (green line) and validation set (blue line) are both plotted. In practice, the two lines coincide almost completely due to the limited amount of overtraining. Bottom: relative importance ranking of each of the observables used by the BDT. Higher importance implies that the observable is used more often, and to make more powerful domain splits.

Category	Variable	Description
Neutron capture vertex	N_{10}	Number of PMT hits in a 10-ns window, comprising the neutron candidate.
	ΔN_{10}	Change in N_{10} after refining the neutron vertex estimate from the initial guess (i.e. the prompt event vertex) to $\mathbf{x}_{\text{fit}}$ to decrease the dispersion of PMT hit times.
	t_{RMS}	RMS deviation of PMT hit times.
	Δt_{RMS}	Change in t_{RMS} after refining the neutron vertex estimate to $\mathbf{x}_{\text{fit}}$.
	fp_{dist}	Distance between the reconstructed prompt vertex and neutron capture vertex estimate $\mathbf{x}_{\text{fit}}$.
	bp_{dist}	Distance between the $\mathbf{x}_{\text{fit}}$ and $\mathbf{x}_{\text{BONSAI}}$ neutron capture vertex estimates.
	`bse`	Number of emitted Cherenkov photons at the neutron capture vertex, reconstructed with BONSAI.
	f_{wall}	Distance of neutron capture vertex estimate $\mathbf{x}_{\text{fit}}$ from the detector wall.
	b_{wall}	Distance of the BONSAI neutron capture vertex estimate $\mathbf{x}_{\text{BONSAI}}$ from the detector wall.
Cherenkov event	θ_{mean}	Mean PMT hit opening angle, calculated from the average PMT hit direction.
	θ_{RMS}	RMS deviation of the PMT hit opening angles.
	ϕ_{RMS}	RMS deviation of the azimuthal separation between PMT hits.
	N_{back}	Number of PMT hits found in the backward hemisphere from the average PMT hit direction.
	$N_{\text{low}\theta}$	Number of hits at an opening angle $< 20°$.
	N_{clus}	Number of PMT hits inside a hit cluster, defined as a group of at least 3 hits with a max. separation of $14.1°$.
	N_{low}	Number of hits on low-probability PMTs, based on the neutron capture vertex.
Noise event	Q_{mean}	Mean charge deposition per PMT hit.
	Q_{RMS}	RMS deviation of PMT hit charge deposition.
	N_{highQ}	Number of PMT hits with high charge deposition.
	N_{300}	Number of PMT hits in a 300-ns window around the original 10-ns neutron candidate window.
	$t^{(3)}_{\text{RMS}}$	Minimum RMS deviation of hit times in 3-hit subset.
	$t^{(6)}_{\text{RMS}}$	Minimum RMS deviation of hit times in 6-hit subset.

TABLE B.1: Summary of the discriminating observables calculated for each neutron capture candidate, divided into three broad categories according to their function.

C

SK-IV spectral fits

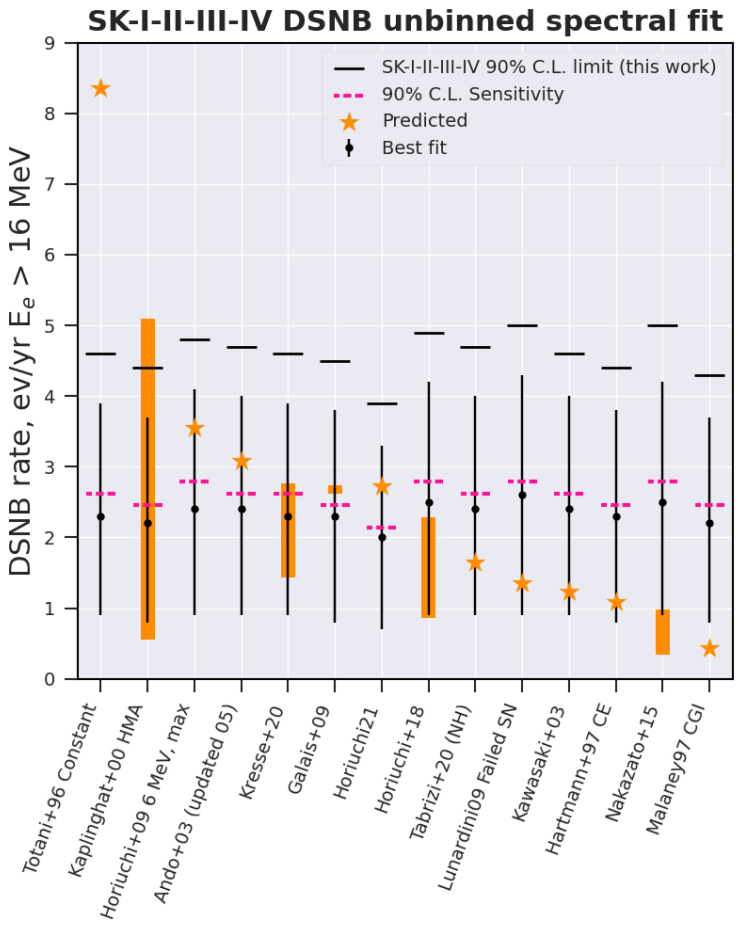

FIGURE C.1: Spectral fit results for the models considered in this work. The black dots and vertical lines are the best-fit DSNB event rates at SK with the associated uncertainties, while the solid black horizontal lines are the 90% C.L. upper limits on the event rate. The dashed magenta lines give the 90% C.L. sensitivity of the search for each model, here defined as 1.64 times the 1σ error on the fit. These values are compared to the theoretical event rate predictions of each model (orange stars and bars). A single star denotes a singular rate prediction, while a bar denotes a range of rate predictions.

	Best fit flux		90% CL limit					Pred.
Model	**SK4**	**All**	**SK1**	**SK2**	**SK3**	**SK4**	**All**	
Totani+95 Constant	$2.5^{+1.4}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.3	7.0	4.5	2.6	4.67
Kaplinghat+00 HMA (max)	$2.6^{+1.5}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.7	7.1	4.7	2.6	3.00
Horiuchi+09 6 MeV, max	$2.6^{+1.4}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.4	6.0	7.0	4.6	2.6	1.94
Ando+03 (updated 05)	$2.7^{+1.5}_{-1.4}$	$1.4^{+0.9}_{-0.9}$	2.3	6.6	7.2	4.7	2.7	1.74
Horiuchi+21	$2.1^{+1.3}_{-1.2}$	$1.2^{+0.9}_{-0.9}$	3.4	4.3	5.9	3.9	2.5	1.72
Kresse+21 (High, NO)	$2.7^{+1.5}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.3	6.7	7.2	4.7	2.7	1.57
Galais+09 (NO)	$2.5^{+1.4}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.3	7.0	4.5	2.6	1.56
Galais+09 (IO)	$2.6^{+1.4}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.4	7.0	4.5	2.6	1.50
Horiuchi+18 $\xi_{2.5} = 0.1$	$2.6^{+1.4}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.4	6.1	7.1	4.6	2.7	1.23
Kresse+21 (High, IO)	$2.7^{+1.5}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.3	6.7	7.1	4.7	2.7	1.21
Kresse+21 (Fid, NO)	$2.7^{+1.5}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.3	6.8	7.2	4.7	2.7	1.20
Kresse+21 (Fid, IO)	$2.7^{+1.5}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.3	6.8	7.2	4.7	2.7	1.02
Kresse+21 (Low, NO)	$2.7^{+1.5}_{-1.4}$	$1.4^{+0.9}_{-0.9}$	2.3	6.8	7.2	4.8	2.7	0.96
Tabrizi+21 (NO)	$2.7^{+1.5}_{-1.3}$	$1.4^{+0.9}_{-0.9}$	2.4	6.6	7.1	4.7	2.7	0.92
Kresse+21 (Low, IO)	$2.7^{+1.5}_{-1.4}$	$1.4^{+0.9}_{-0.9}$	2.3	6.8	7.2	4.8	2.7	0.84
Lunardini09 Failed SN	$2.8^{+1.5}_{-1.4}$	$1.4^{+0.9}_{-0.9}$	2.4	6.8	7.3	4.8	2.8	0.73
Hartmann+97 CE	$2.6^{+1.4}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.5	7.1	4.6	2.6	0.63
Nakazato+15 (max, IO)	$2.7^{+1.5}_{-1.4}$	$1.4^{+1.0}_{-0.9}$	2.4	6.5	7.2	4.8	2.7	0.53
Horiuchi+18 $\xi_{2.5} = 0.5$	$2.7^{+1.5}_{-1.4}$	$1.3^{+0.9}_{-0.9}$	2.2	7.1	7.1	4.8	2.6	0.55
Malaney97 CGI	$2.7^{+1.5}_{-1.3}$	$1.3^{+0.9}_{-0.9}$	2.3	6.8	7.1	4.7	2.6	0.26
Nakazato+15 (min, NO)	$2.8^{+1.5}_{-1.4}$	$1.4^{+1.0}_{-0.9}$	2.3	6.8	7.2	4.8	2.7	0.19

	Best fit rate		90% CL limit					Pred.
Model	**SK4**	**All**	**SK1**	**SK2**	**SK3**	**SK4**	**All**	
Totani+95 Constant	$4.5^{+2.5}_{-2.2}$	$2.3^{+1.6}_{-1.5}$	4.2	11.2	12.2	7.9	4.6	8.35
Kaplinghat+00 HMA (max)	$4.4^{+2.4}_{-2.2}$	$2.2^{+1.5}_{-1.4}$	3.9	11.2	11.7	7.7	4.4	5.09
Horiuchi+09 6 MeV, max	$4.6^{+2.5}_{-2.3}$	$2.4^{+1.7}_{-1.6}$	4.4	11.7	12.5	8.2	4.8	3.54
Ando+03 (updated 05)	$4.7^{+2.5}_{-2.4}$	$2.4^{+1.6}_{-1.6}$	4.2	11.8	12.4	8.2	4.7	3.09
Kresse+21 (High, NO)	$4.5^{+2.5}_{-2.3}$	$2.4^{+1.6}_{-1.5}$	4.1	11.7	12.2	8.1	4.6	2.76
Galais+09 (NO)	$4.4^{+2.4}_{-2.2}$	$2.2^{+1.6}_{-1.4}$	4.0	11.0	11.9	7.8	4.5	2.74
Horiuchi+21	$3.7^{+2.2}_{-2.1}$	$2.0^{+1.6}_{-1.6}$	6.0	7.5	10.2	6.8	4.2	2.72
Galais+09 (IO)	$4.4^{+2.4}_{-2.2}$	$2.2^{+1.6}_{-1.4}$	4.0	11.0	11.9	7.7	4.5	2.62
Horiuchi+18 $\xi_{2.5} = 0.1$	$4.8^{+2.6}_{-2.4}$	$2.5^{+1.7}_{-1.6}$	4.5	12.1	12.8	8.4	4.9	2.29
Kresse+21 (High, IO)	$4.6^{+2.5}_{-2.3}$	$2.4^{+1.6}_{-1.5}$	4.1	11.7	12.2	8.1	4.6	2.14
Kresse+21 (Fid, NO)	$4.5^{+2.5}_{-2.2}$	$2.3^{+1.5}_{-1.5}$	4.0	11.7	11.9	7.9	4.5	2.06
Kresse+21 (Fid, IO)	$4.5^{+2.5}_{-2.2}$	$2.3^{+1.5}_{-1.5}$	4.0	11.6	11.9	7.9	4.5	1.75
Kresse+21 (Low, NO)	$4.5^{+2.5}_{-2.2}$	$2.2^{+1.6}_{-1.4}$	3.9	11.6	11.8	7.9	4.5	1.65
Tabrizi+21 (NO)	$4.6^{+2.5}_{-2.3}$	$2.4^{+1.6}_{-1.5}$	4.2	11.8	12.3	8.2	4.7	1.64
Kresse+21 (Low, IO)	$4.5^{+2.5}_{-2.2}$	$2.2^{+1.6}_{-1.4}$	4.0	11.6	11.9	7.9	4.5	1.43
Lunardini09 Failed SN	$5.0^{+2.7}_{-2.5}$	$2.6^{+1.7}_{-1.7}$	4.5	12.7	13.1	8.8	5.1	1.36
Hartmann+97 CE	$4.4^{+2.4}_{-2.2}$	$2.2^{+1.6}_{-1.4}$	4.0	11.2	11.9	7.7	4.5	1.09
Nakazato+15 (max, IO)	$5.0^{+2.7}_{-2.5}$	$2.5^{+1.8}_{-1.6}$	4.5	12.1	13.0	8.7	5.0	0.99
Horiuchi+18 $\xi_{2.5} = 0.5$	$4.2^{+2.3}_{-2.1}$	$2.1^{+1.5}_{-1.4}$	3.5	11.4	10.9	7.4	4.1	0.87
Malaney97 CGI	$4.3^{+2.4}_{-2.1}$	$2.2^{+1.5}_{-1.4}$	3.8	11.2	11.3	7.6	4.3	0.44
Nakazato+15 (min, NO)	$4.7^{+2.5}_{-2.4}$	$2.4^{+1.6}_{-1.6}$	4.0	11.8	12.1	8.2	4.6	0.34

TABLE C.1: Top: best-fit values and 90% C.L. upper limits on the DSNB fluxes (in $\mathrm{cm}^{-2}\cdot\mathrm{sec}^{-1}$). Bottom: best-fit values and the 90% C.L. upper limits on the DSNB rates (in events·year^{-1}). The fit results are shown for a variety of theoretical models, for phases SK-I to IV as well as for the combined analysis. Here the upper limits are given for $E_\nu > 17.3$ MeV.

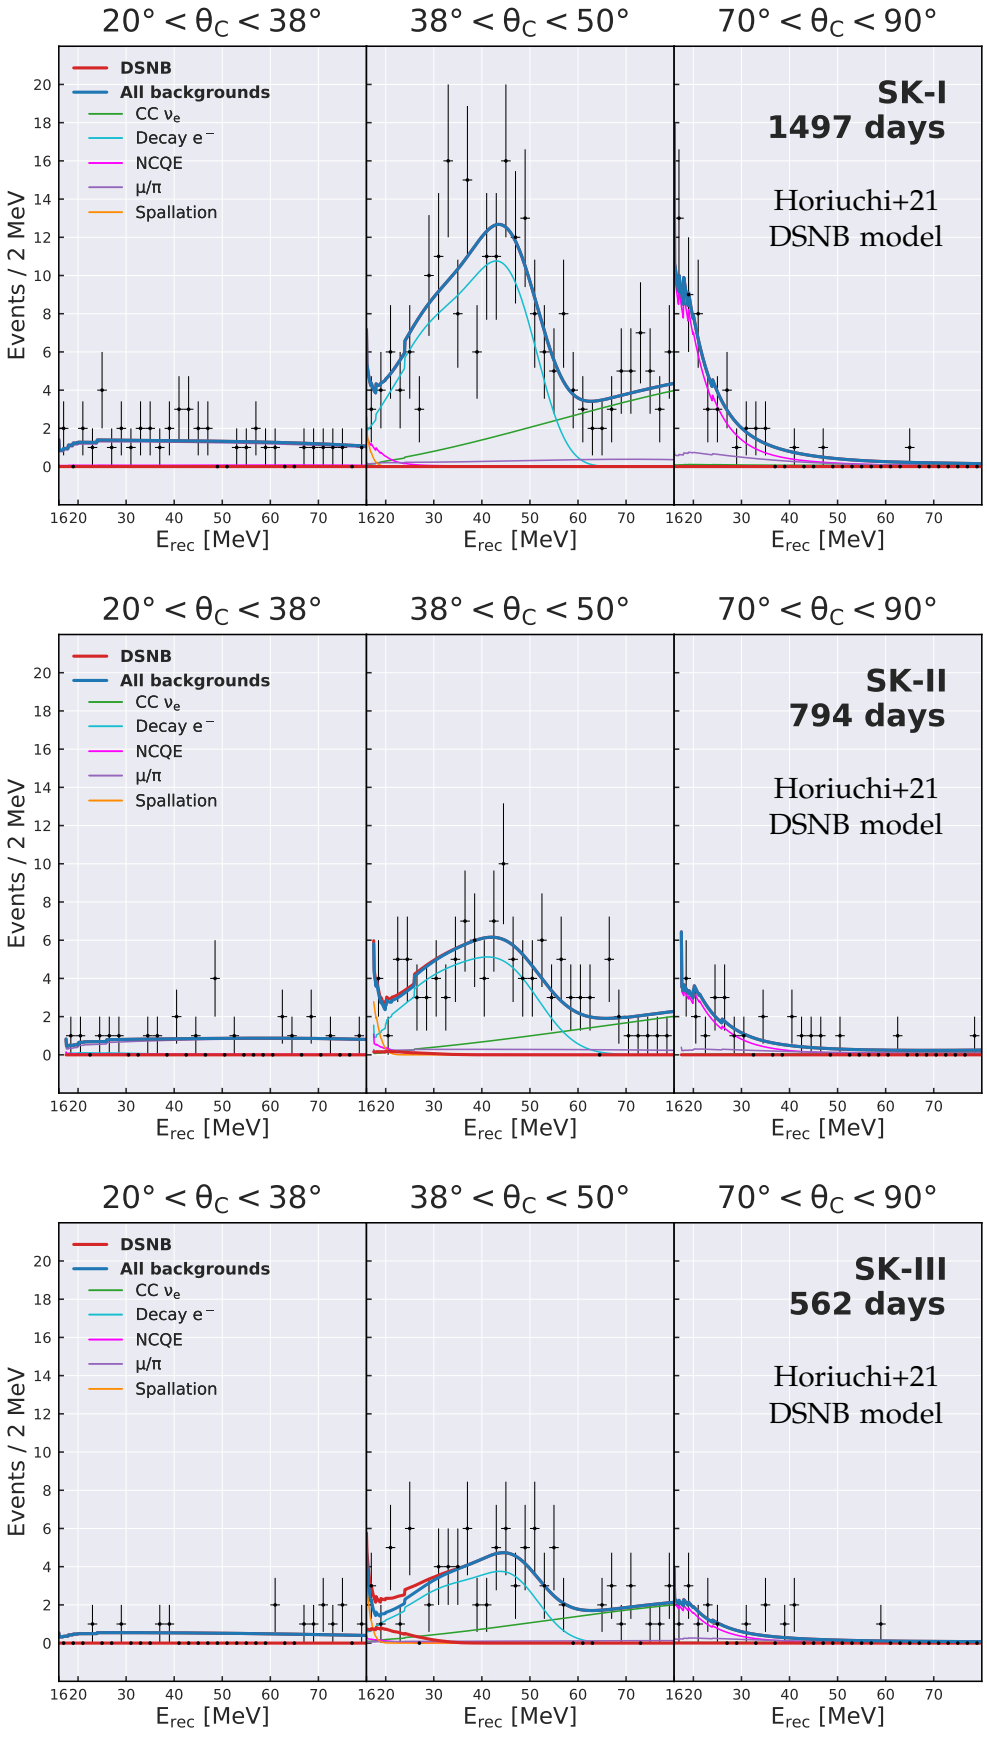

FIGURE C.2: Best-fit signal and background spectra (solid lines) for the SK-I (top), SK-II (middle), and SK-III (bottom) periods, assuming the DSNB flux predicted by the Horiuchi+21 model [19], along with the energy spectrum of the remaining events in data after reduction. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

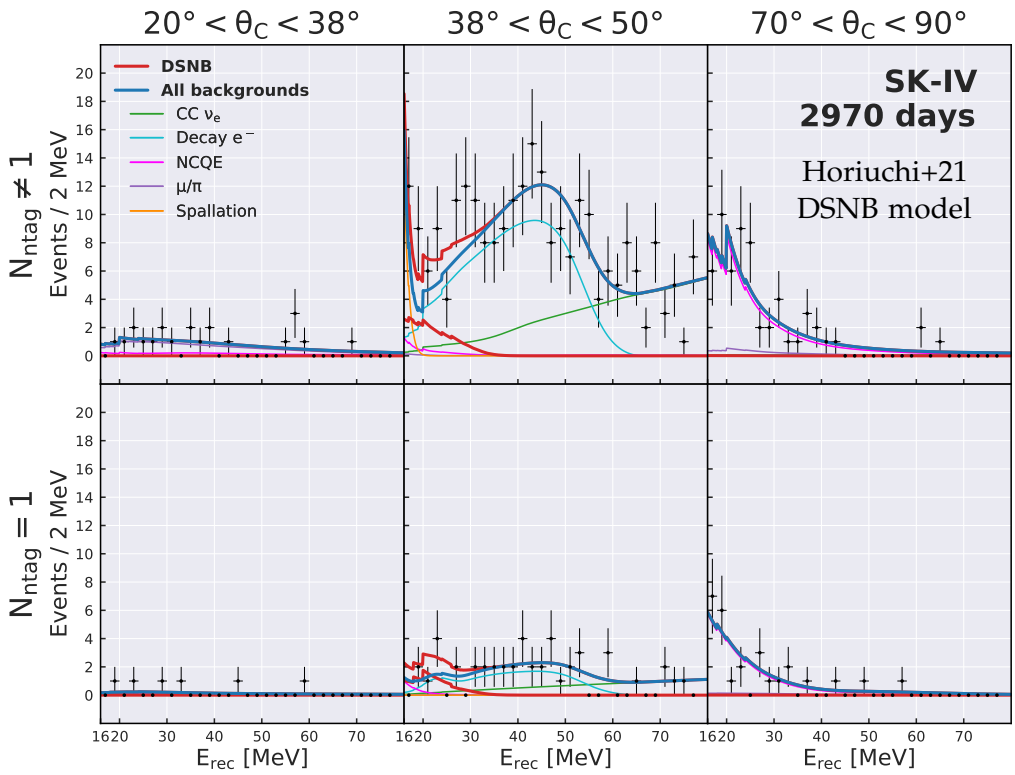

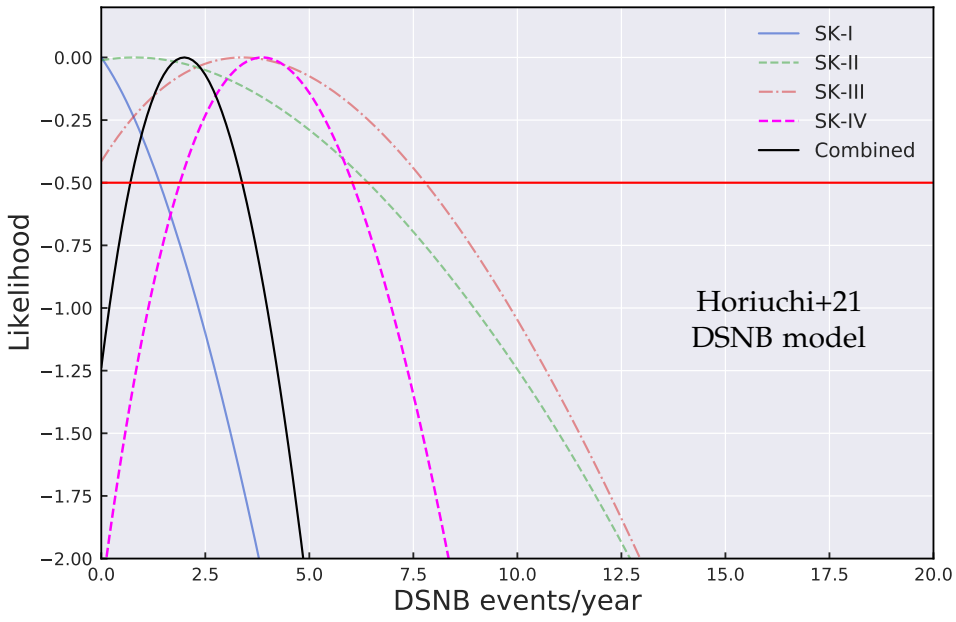

FIGURE C.3: Top:Best-fit signal and background spectra (solid lines) for the SK-IV period, assuming the DSNB flux predicted by the Horiuchi+21 model [19], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I to IV of SK, as well as the combined likelihood.

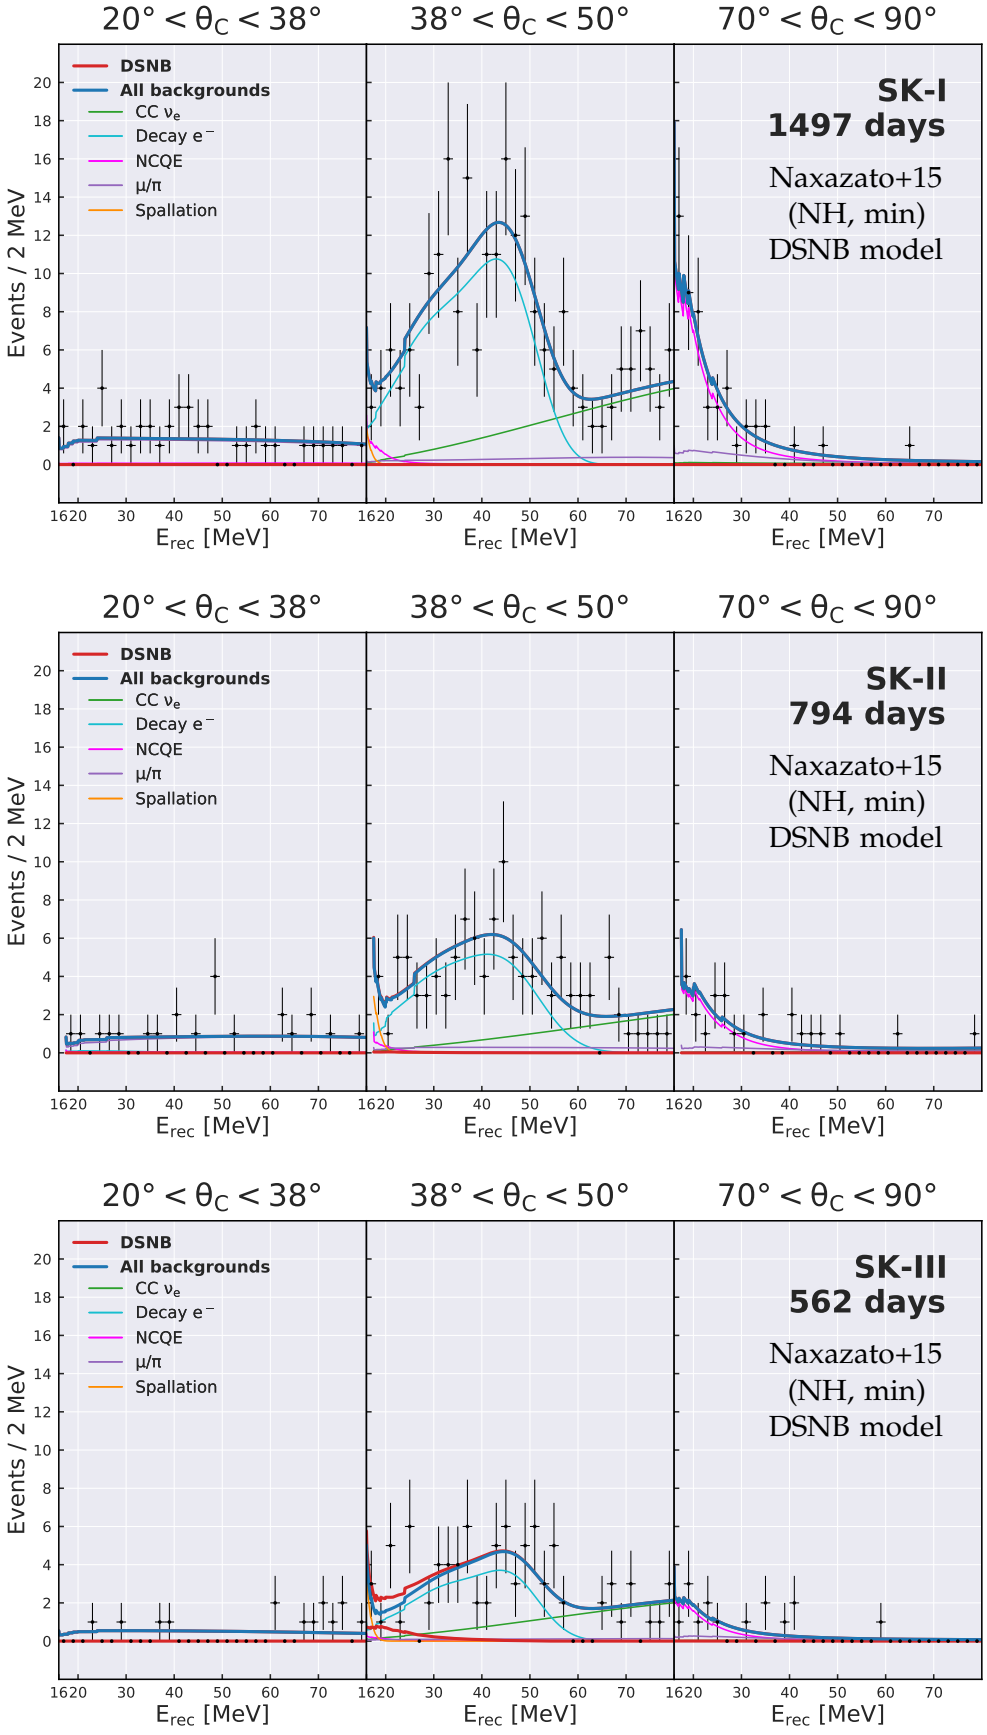

FIGURE C.4: Best-fit signal and background spectra (solid lines) for the SK-I (top), SK-II (middle), and SK-III (bottom) periods, assuming the DSNB flux predicted by the Horiuchi+21 model [26], along with the energy spectrum of the remaining events in data after reduction. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

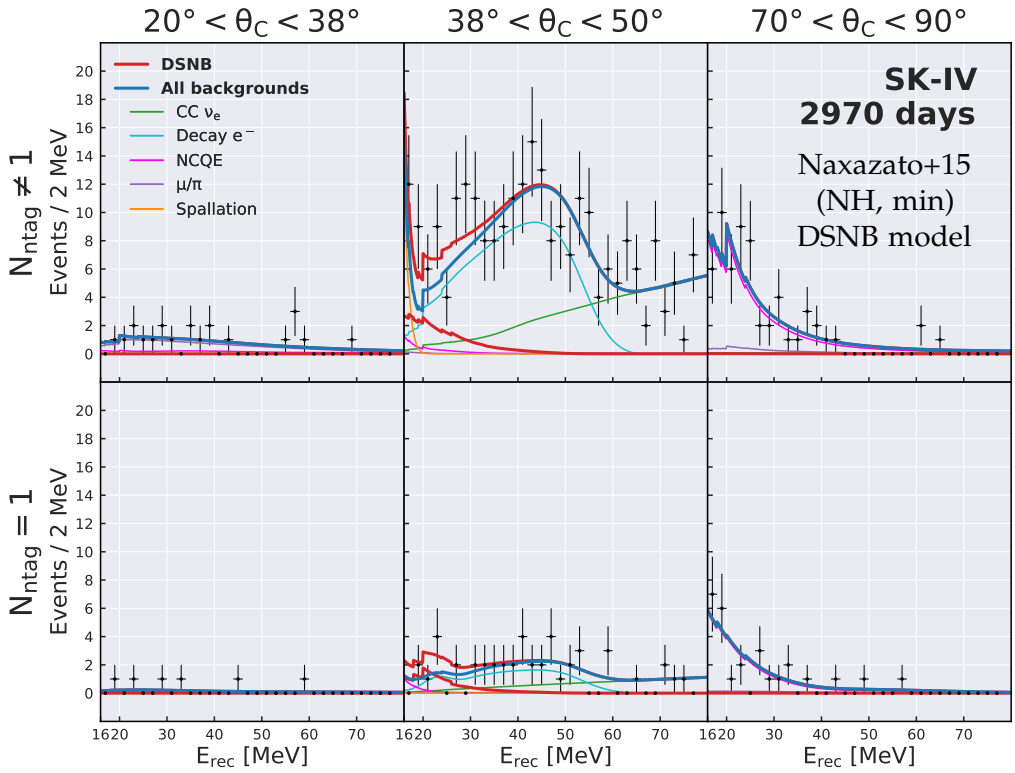

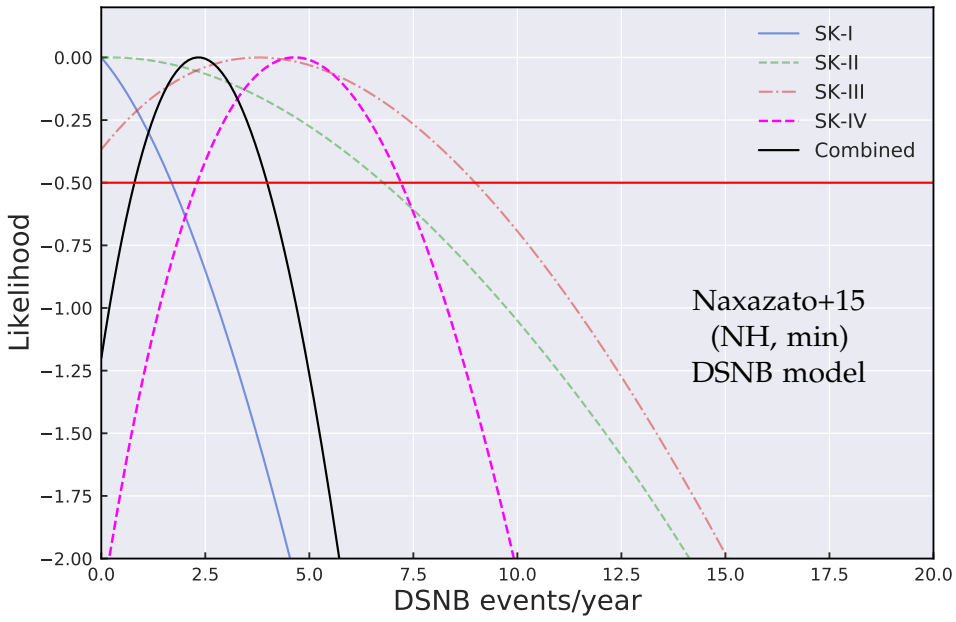

FIGURE C.5: Top:Best-fit signal and background spectra (solid lines) for the SK-IV period, assuming the DSNB flux predicted by the Horiuchi+21 model [26], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I to IV of SK, as well as the combined likelihood.

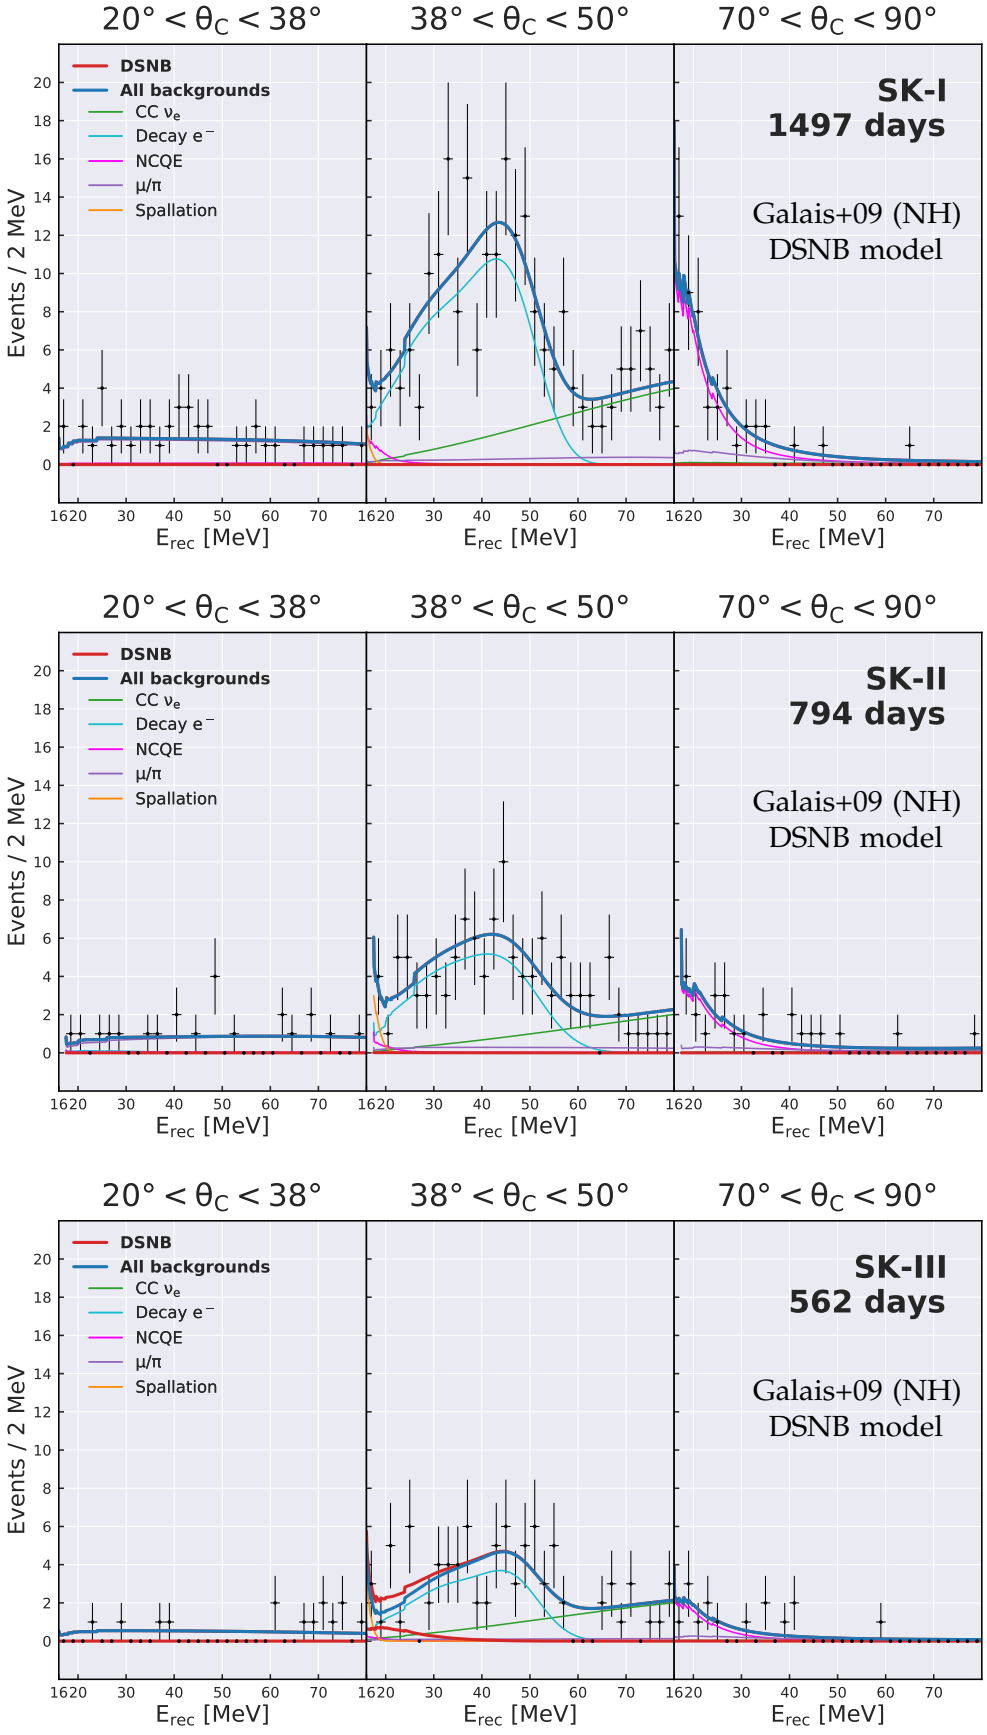

FIGURE C.6: Best-fit signal and background spectra (solid lines) for the SK-I (top), SK-II (middle), and SK-III (bottom) periods, assuming the DSNB flux predicted by the Horiuchi+21 model [22], along with the energy spectrum of the remaining events in data after reduction. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

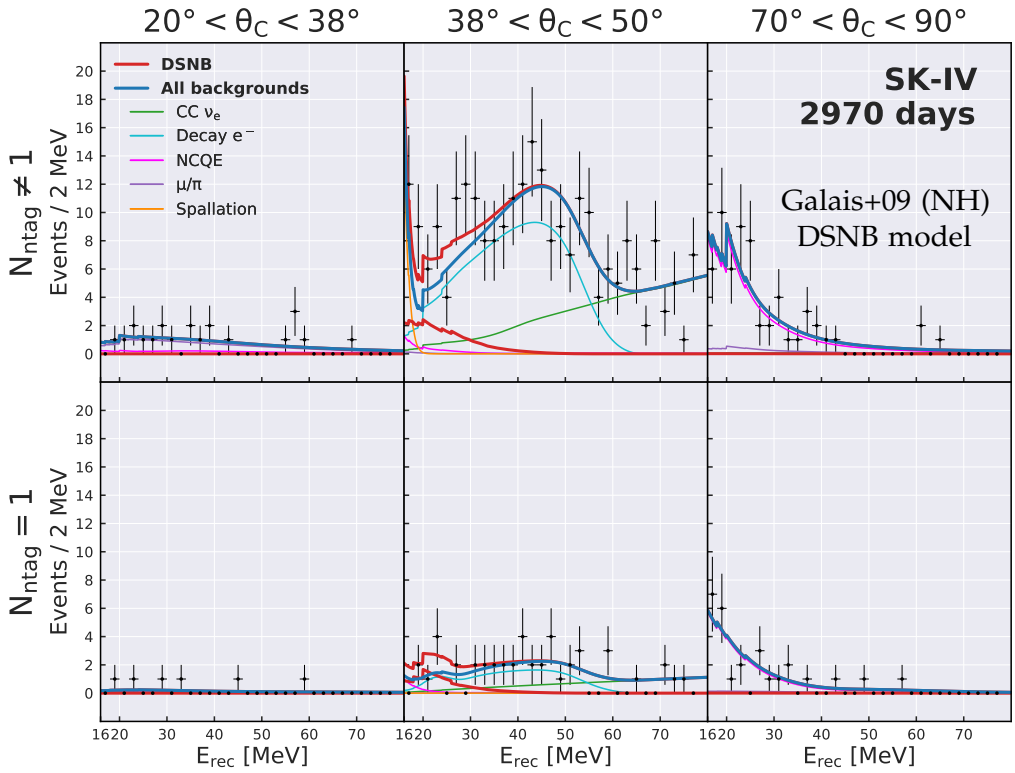

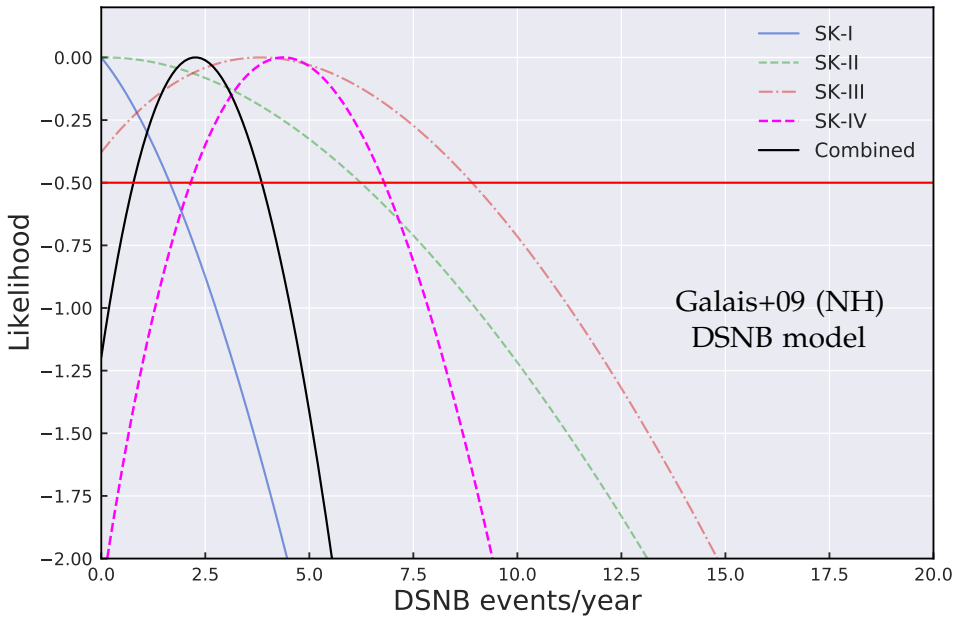

FIGURE C.7: Top:Best-fit signal and background spectra (solid lines) for the SK-IV period, assuming the DSNB flux predicted by the Horiuchi+21 model [22], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I to IV of SK, as well as the combined likelihood.

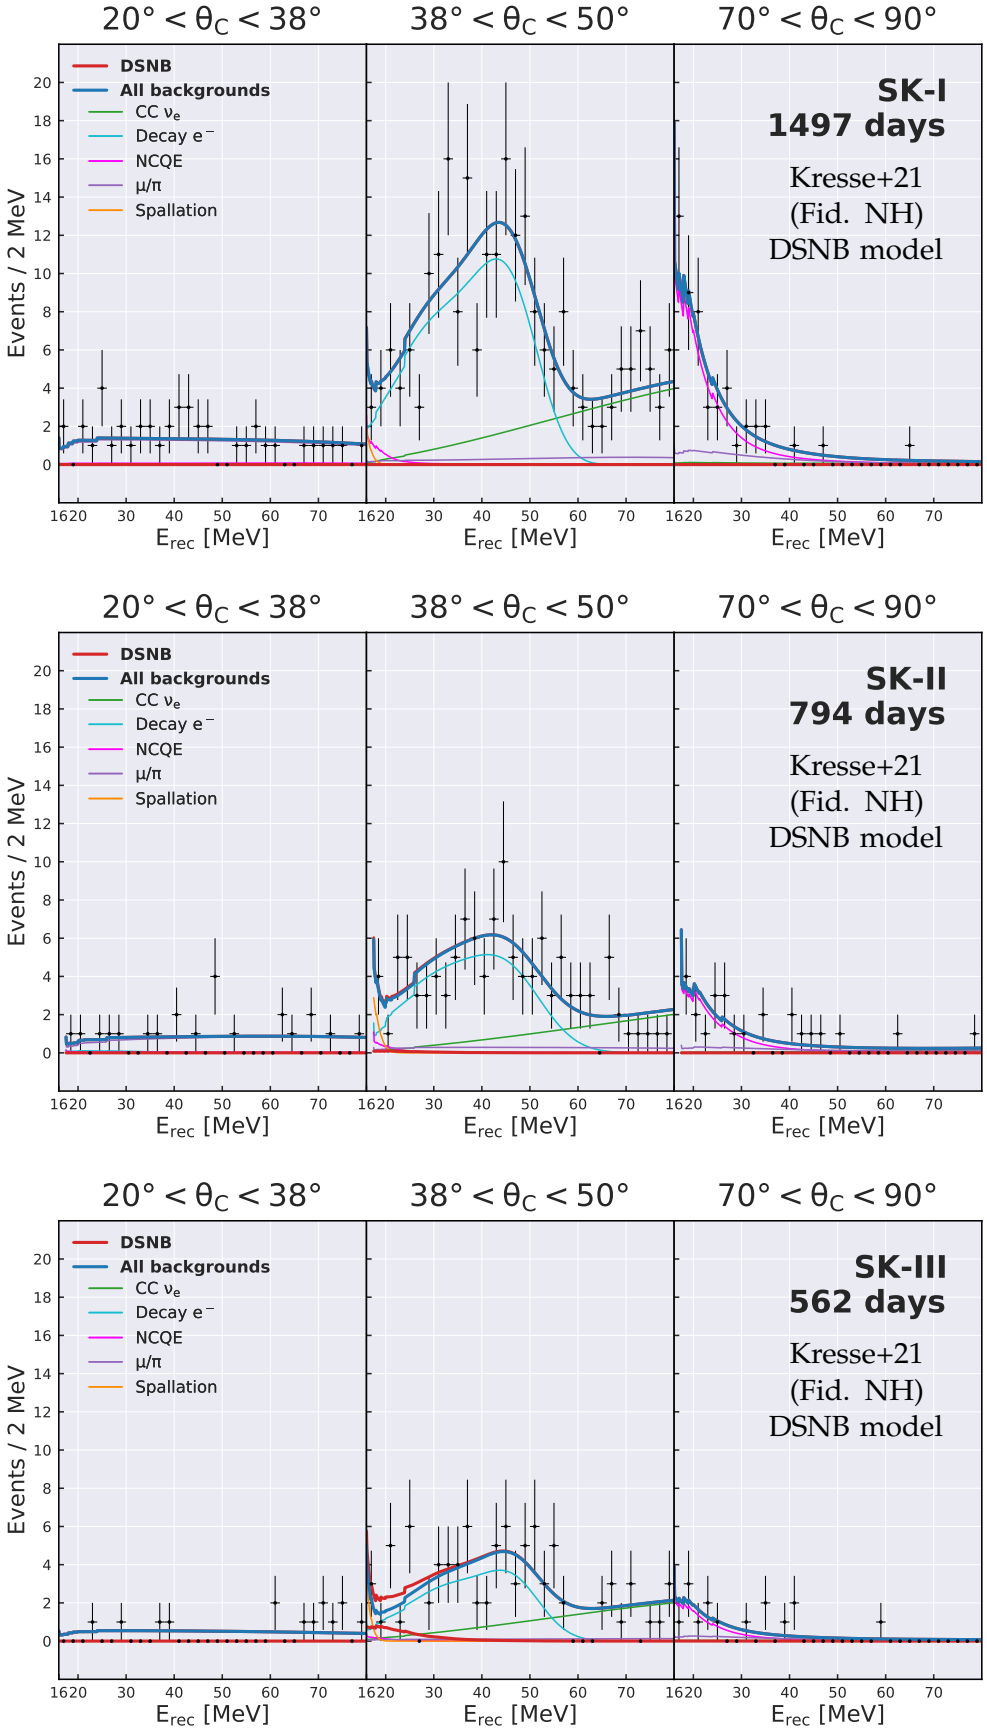

FIGURE C.8: Best-fit signal and background spectra (solid lines) for the SK-I (top), SK-II (middle), and SK-III (bottom) periods, assuming the DSNB flux predicted by the Horiuchi+21 model [25], along with the energy spectrum of the remaining events in data after reduction. Here the data is binned for illustrative purposes only, since the analysis is unbinned.

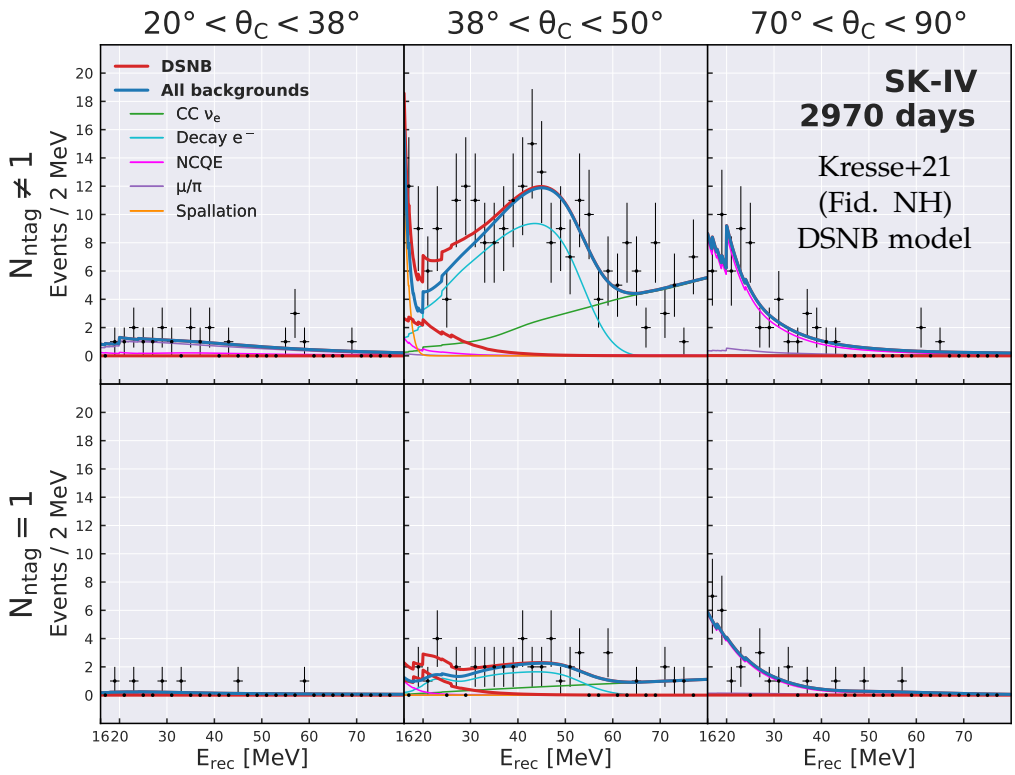

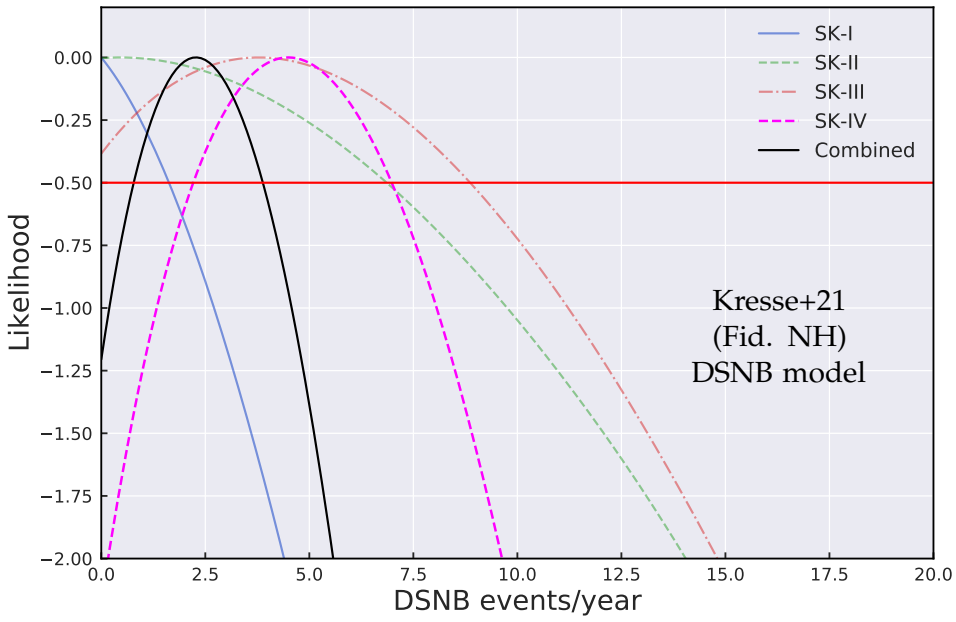

FIGURE C.9: Top:Best-fit signal and background spectra (solid lines) for the SK-IV period, assuming the DSNB flux predicted by the Horiuchi+21 model [25], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I to IV of SK, as well as the combined likelihood.

D
SK-VI neutron tagging BDT

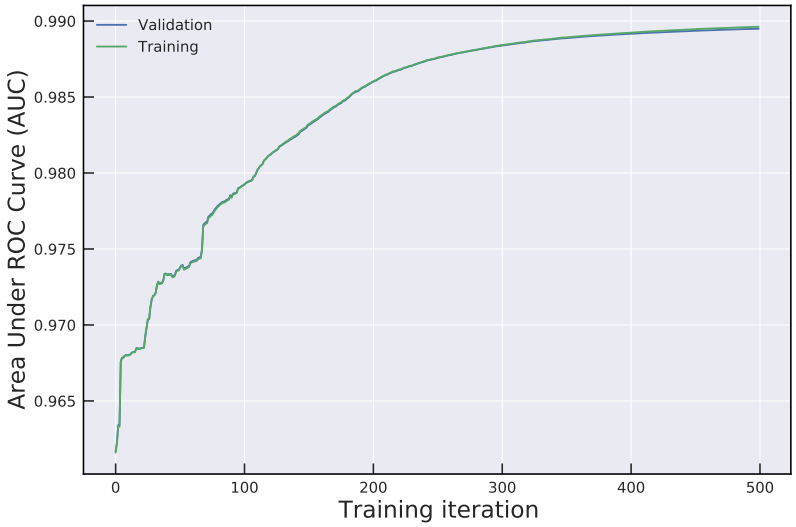

FIGURE D.1: Top: evolution of the SK-VI neutron tagging BDT model performance metric $(1 - \text{AUC})$ as function of the training iteration. The model improves sharply in the first 200 iterations, and begins plateauing after around 500 iterations. The performance on the training set (green line) and validation set (blue line) are both plotted. In practice, the two lines coincide almost completely due to the limited amount of overtraining.

Hyperparameter	Chosen value
Learning rate	0.25
Max tree depth	6
Max number of trees	1500
Early stopping threshold	50
Subsampling fraction	0.97
Signal weight	1.0
Training method	`auto`

TABLE D.1: Hyperparameters chosen for the SK-VI neutron tagging BDT model.

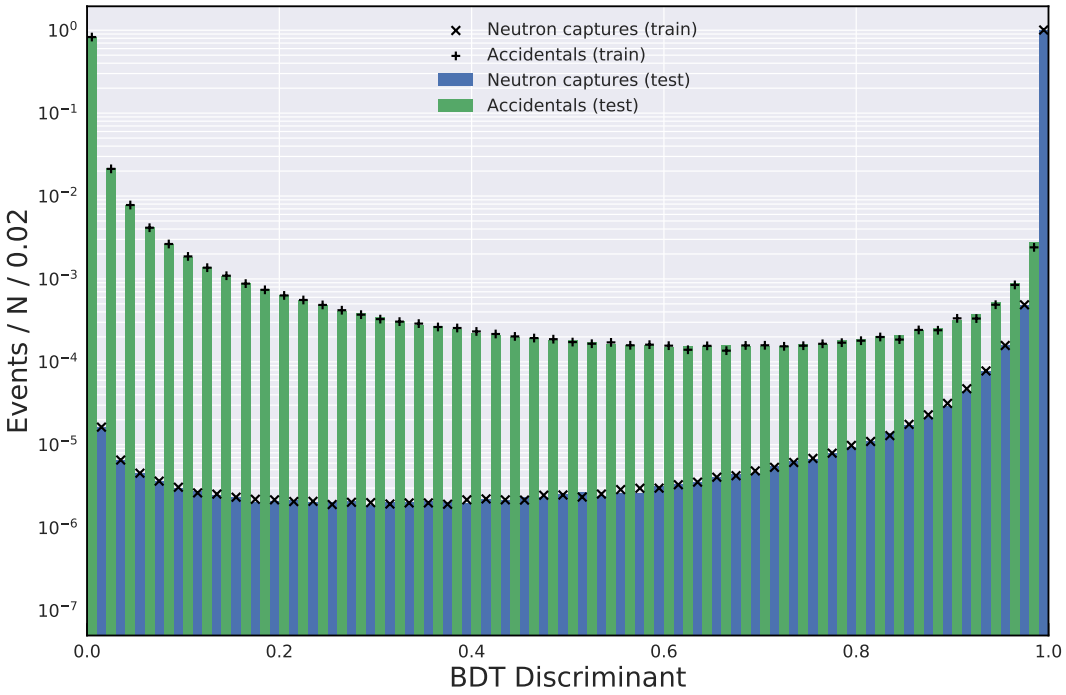

FIGURE D.2: SK-VI neutron tagging BDT output value calculated on the entire SK-IV test dataset, for signal (blue histogram) and background (green histogram). The BDT output is also plotted for the training dataset, for signal ($\times$) and background ($+$). Each distribution is normalized to an area of 1. Notice that the plot is in log scale, so the signal and background distributions are very strongly peaked at 1 and 0, respectively. In the SK-VI dataset, half of the neutron captures occur on Gd.

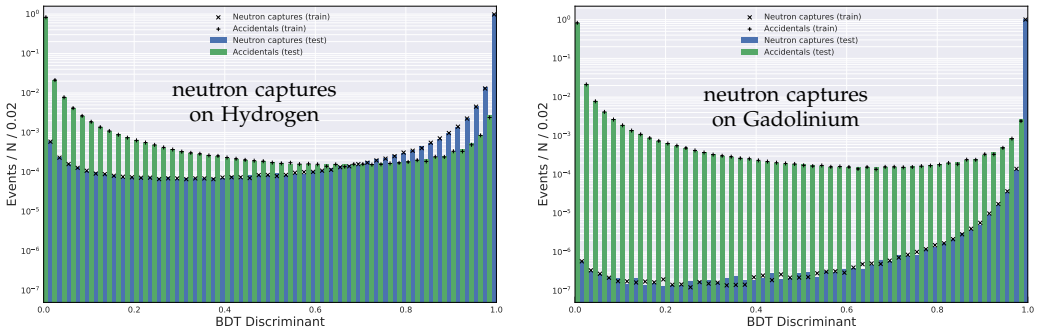

FIGURE D.3: Left: SK-VI neutron tagging BDT output value calculated on the for neutron captures on Hydrogen only (blue histogram) and background (green histogram). Right: SK-VI neutron tagging BDT output value calculated on the for neutron captures on Gadolinium only (blue histogram) and background (green histogram).

E
SK-VI spectral fits

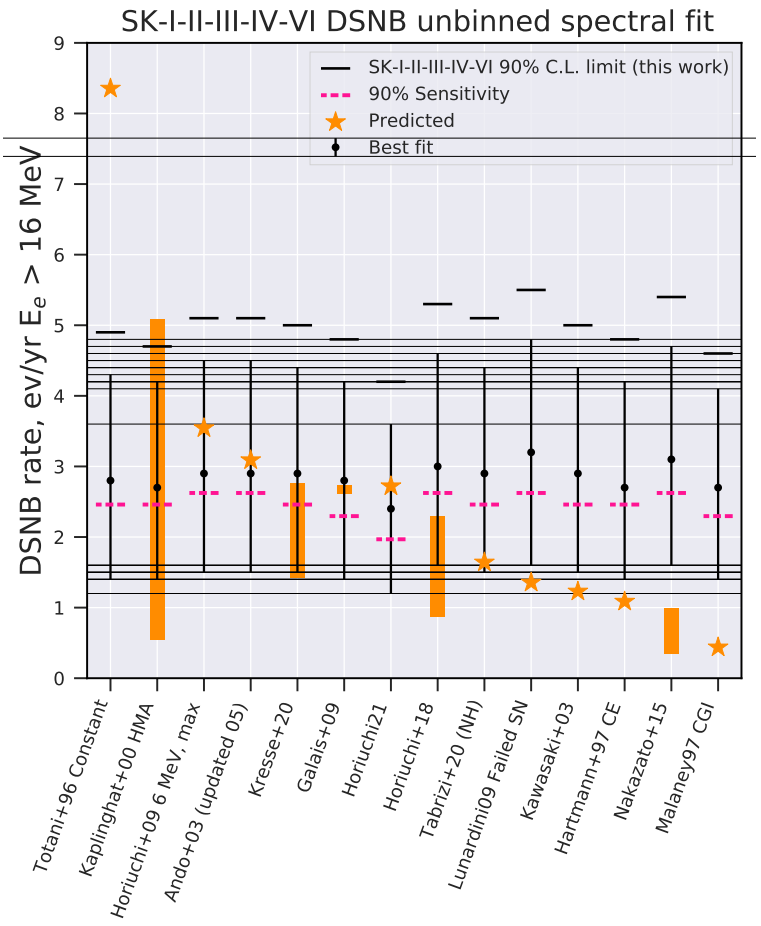

FIGURE E.1: Spectral fit results for the models considered in this work. The black dots and vertical lines are the best-fit DSNB event rates at SK with the associated uncertainties, while the solid black horizontal lines are the 90% C.L. upper limits on the event rate. The dashed magenta lines give the 90% C.L. sensitivity of the search for each model, here defined as 1.64 times the 1σ error on the fit. These values are compared to the theoretical event rate predictions of each model (orange stars and bars). A single star denotes a singular rate prediction, while a bar denotes a range of rate predictions.

Model	Best fit SK6	Best fit All	90% CL limit SK6	90% CL limit All	90% sens. SK6	90% sens. All	Pred.
Totani+96 Constant	$3.8^{+3.3}_{-2.6}$	$1.6^{+0.8}_{-0.8}$	9.1	2.8	5.4	1.3	4.67
Kaplinghat+00 HMA (max)	$4.0^{+3.4}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.4	2.8	5.6	1.5	3.00
Horiuchi+09 6 MeV, max	$3.9^{+3.3}_{-2.6}$	$1.6^{+0.8}_{-0.8}$	9.1	2.8	5.4	1.3	1.94
Ando+03 (updated 05)	$4.0^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.7}$	9.5	2.9	5.7	1.5	1.74
Horiuchi21	$3.5^{+3.3}_{-2.5}$	$1.5^{+0.8}_{-0.7}$	8.9	2.7	5.4	1.3	1.72
Kresse+20 (High, NH)	$4.0^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.7}$	9.5	2.9	5.7	1.5	1.57
Galais+09 (NH)	$3.8^{+3.3}_{-2.6}$	$1.6^{+0.8}_{-0.8}$	8.9	2.8	5.4	1.3	1.56
Galais+09 (IH)	$3.8^{+3.3}_{-2.6}$	$1.6^{+0.8}_{-0.8}$	9.0	2.8	5.4	1.3	1.50
Horiuchi+18 $\xi_{2.5} = 0.1$	$4.0^{+3.3}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.2	2.9	5.4	1.5	1.23
Kresse+20 (High, IH)	$4.0^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.7}$	9.5	2.9	5.7	1.5	1.21
Kresse+20 (Fid, NH)	$4.1^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.7}$	9.6	2.9	5.7	1.5	1.20
Kresse+20 (Fid, IH)	$4.1^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.6	2.9	5.7	1.5	1.02
Kresse+20 (Low, NH)	$4.1^{+3.4}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.6	2.9	5.6	1.5	0.96
Tabrizi+20 (NH)	$4.1^{+3.4}_{-2.7}$	$1.7^{+0.8}_{-0.8}$	9.5	2.9	5.6	1.3	0.92
Kresse+20 (Low, IH)	$4.1^{+3.4}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.6	2.9	5.6	1.5	0.84
Lunardini09 Failed SN	$4.4^{+3.6}_{-2.8}$	$1.7^{+0.9}_{-0.8}$	9.9	3.0	5.9	1.5	0.73
Kawasaki+03	$4.2^{+3.5}_{-2.7}$	$1.7^{+0.9}_{-0.8}$	9.7	2.9	5.7	1.5	0.71
Hartmann+97 CE	$3.9^{+3.4}_{-2.6}$	$1.6^{+0.8}_{-0.8}$	9.2	2.8	5.6	1.3	0.63
Nakazato+15 (max, IH)	$3.9^{+3.3}_{-2.6}$	$1.7^{+0.8}_{-0.8}$	9.1	2.9	5.5	1.3	0.53
Horiuchi+18 $\xi_{2.5} = 0.5$	$4.1^{+3.6}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.9	2.8	5.9	1.5	0.55
Malaney97 CGI	$4.0^{+3.5}_{-2.7}$	$1.6^{+0.9}_{-0.8}$	9.5	2.8	5.8	1.5	0.26
Nakazato+15 (min, NH)	$4.0^{+3.5}_{-2.7}$	$1.7^{+0.8}_{-0.8}$	9.5	2.9	5.7	1.3	0.19

Model	Best fit SK6	Best fit All	90% CL limit SK6	90% CL limit All	1σ sens. SK6	1σ sens. All	Pred.
Totani+96 Constant	$6.7^{+5.7}_{-4.5}$	$2.8^{+1.5}_{-1.4}$	15.8	4.9	9.4	2.5	8.35
Kaplinghat+00 HMA (max)	$6.6^{+5.7}_{-4.4}$	$2.7^{+1.5}_{-1.3}$	15.5	4.7	9.3	2.5	5.09
Horiuchi+09 6 MeV, max	$7.0^{+5.9}_{-4.7}$	$2.9^{+1.6}_{-1.4}$	16.3	5.1	9.7	2.6	3.54
Ando+03 (updated 05)	$7.0^{+6.0}_{-4.7}$	$2.9^{+1.6}_{-1.4}$	16.5	5.1	9.9	2.6	3.09
Kresse+20 (High, NH)	$6.9^{+5.9}_{-4.6}$	$2.9^{+1.5}_{-1.4}$	16.2	5.0	9.7	2.5	2.76
Galais+09 (NH)	$6.4^{+5.7}_{-4.4}$	$2.8^{+1.4}_{-1.4}$	15.3	4.8	9.3	2.3	2.74
Horiuchi21	$5.4^{+5.0}_{-3.9}$	$2.4^{+1.2}_{-1.2}$	13.6	4.2	8.2	2.0	2.72
Galais+09 (IH)	$6.5^{+5.6}_{-4.5}$	$2.7^{+1.5}_{-1.3}$	15.2	4.8	9.1	2.5	2.62
Horiuchi+18 $\xi_{2.5} = 0.1$	$7.3^{+6.0}_{-4.9}$	$3.0^{+1.6}_{-1.4}$	16.8	5.3	9.8	2.6	2.29
Kresse+20 (High, IH)	$6.9^{+5.9}_{-4.6}$	$2.9^{+1.5}_{-1.4}$	16.3	5.0	9.7	2.5	2.14
Kresse+20 (Fid, NH)	$6.8^{+5.8}_{-4.5}$	$2.8^{+1.5}_{-1.4}$	16.0	4.9	9.5	2.5	2.06
Kresse+20 (Fid, IH)	$6.8^{+5.8}_{-4.5}$	$2.8^{+1.5}_{-1.4}$	16.0	4.9	9.5	2.5	1.75
Kresse+20 (Low, NH)	$6.8^{+5.7}_{-4.6}$	$2.8^{+1.4}_{-1.4}$	15.9	4.8	9.3	2.3	1.65
Tabrizi+20 (NH)	$7.1^{+5.9}_{-4.7}$	$2.9^{+1.5}_{-1.4}$	16.5	5.1	9.7	2.5	1.64
Kresse+20 (Low, IH)	$6.8^{+5.7}_{-4.6}$	$2.8^{+1.4}_{-1.4}$	15.9	4.8	9.3	2.3	1.43
Lunardini09 Failed SN	$8.0^{+6.5}_{-5.1}$	$3.2^{+1.6}_{-1.6}$	18.1	5.5	10.7	2.6	1.36
Kawasaki+03	$7.0^{+5.9}_{-4.6}$	$2.9^{+1.5}_{-1.4}$	16.4	5.0	9.7	2.5	1.23
Hartmann+97 CE	$6.5^{+5.7}_{-4.4}$	$2.7^{+1.5}_{-1.3}$	15.5	4.8	9.4	2.5	1.09
Nakazato+15 (max, IH)	$7.0^{+6.0}_{-4.7}$	$3.1^{+1.6}_{-1.5}$	16.4	5.4	9.9	2.6	0.99
Horiuchi+18 $\xi_{2.5} = 0.5$	$6.2^{+5.5}_{-4.2}$	$2.5^{+1.4}_{-1.2}$	15.1	4.4	9.0	2.3	0.87
Malaney97 CGI	$6.3^{+5.7}_{-4.3}$	$2.7^{+1.4}_{-1.3}$	15.2	4.6	9.3	2.3	0.44
Nakazato+15 (min, NH)	$6.8^{+5.9}_{-4.5}$	$2.9^{+1.5}_{-1.4}$	16.0	5.0	9.6	2.5	0.34

TABLE E.1: Top: best-fit values and 90% C.L. upper limits on the DSNB fluxes (in $\mathrm{cm}^{-2}\cdot\mathrm{sec}^{-1}$). Bottom: best-fit values and 90% C.L. upper limits on the DSNB rates (in events$\cdot\mathrm{year}^{-1}$). The fit results are shown for a variety of theoretical models, for the SK-VI and for the combined analysis. Here the upper limits are given for $E_\nu > 17.3$ MeV.

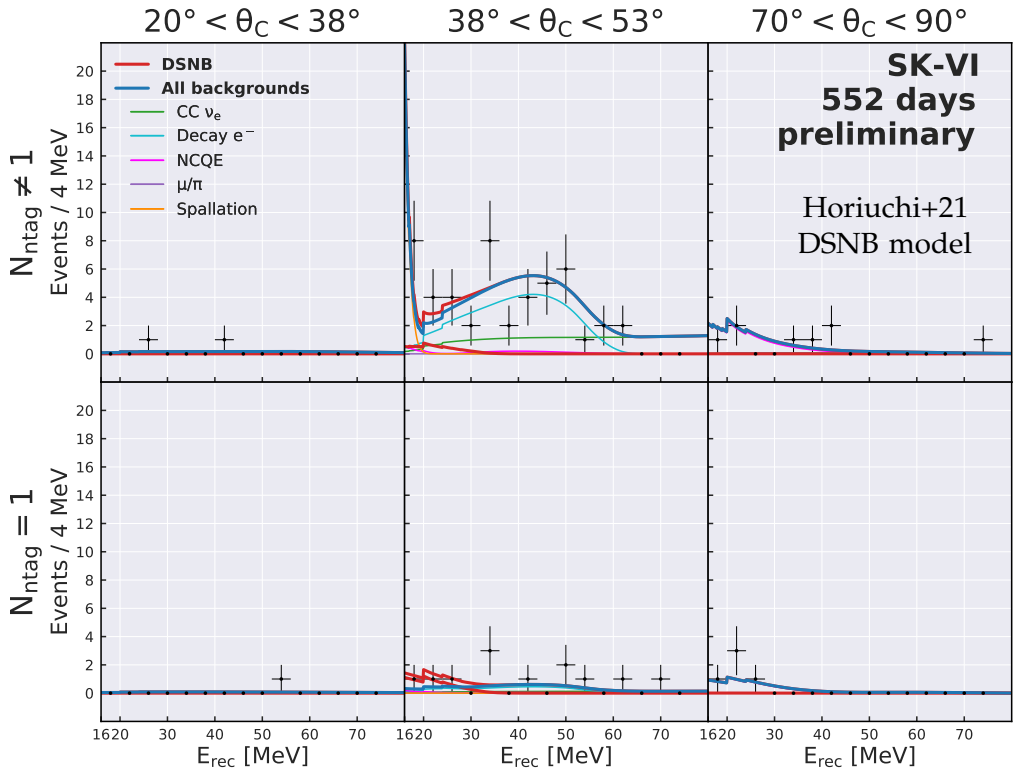

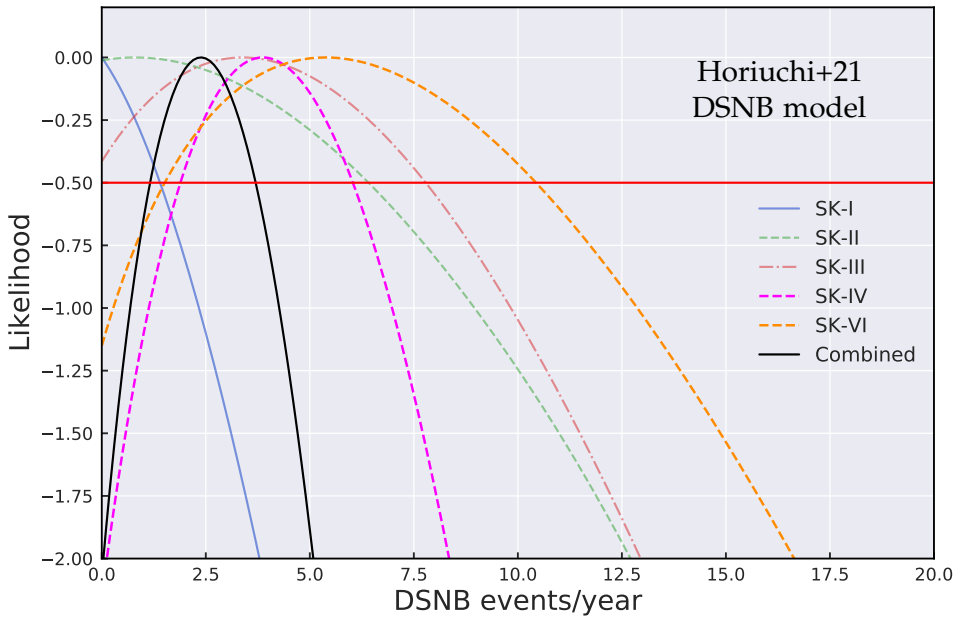

FIGURE E.2: Top:Best-fit signal and background spectra (solid lines) for the SK-VI period, assuming the DSNB flux predicted by the Horiuchi+21 model [19], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I, II, III, IV, and VI of SK, as well as the combined likelihood.

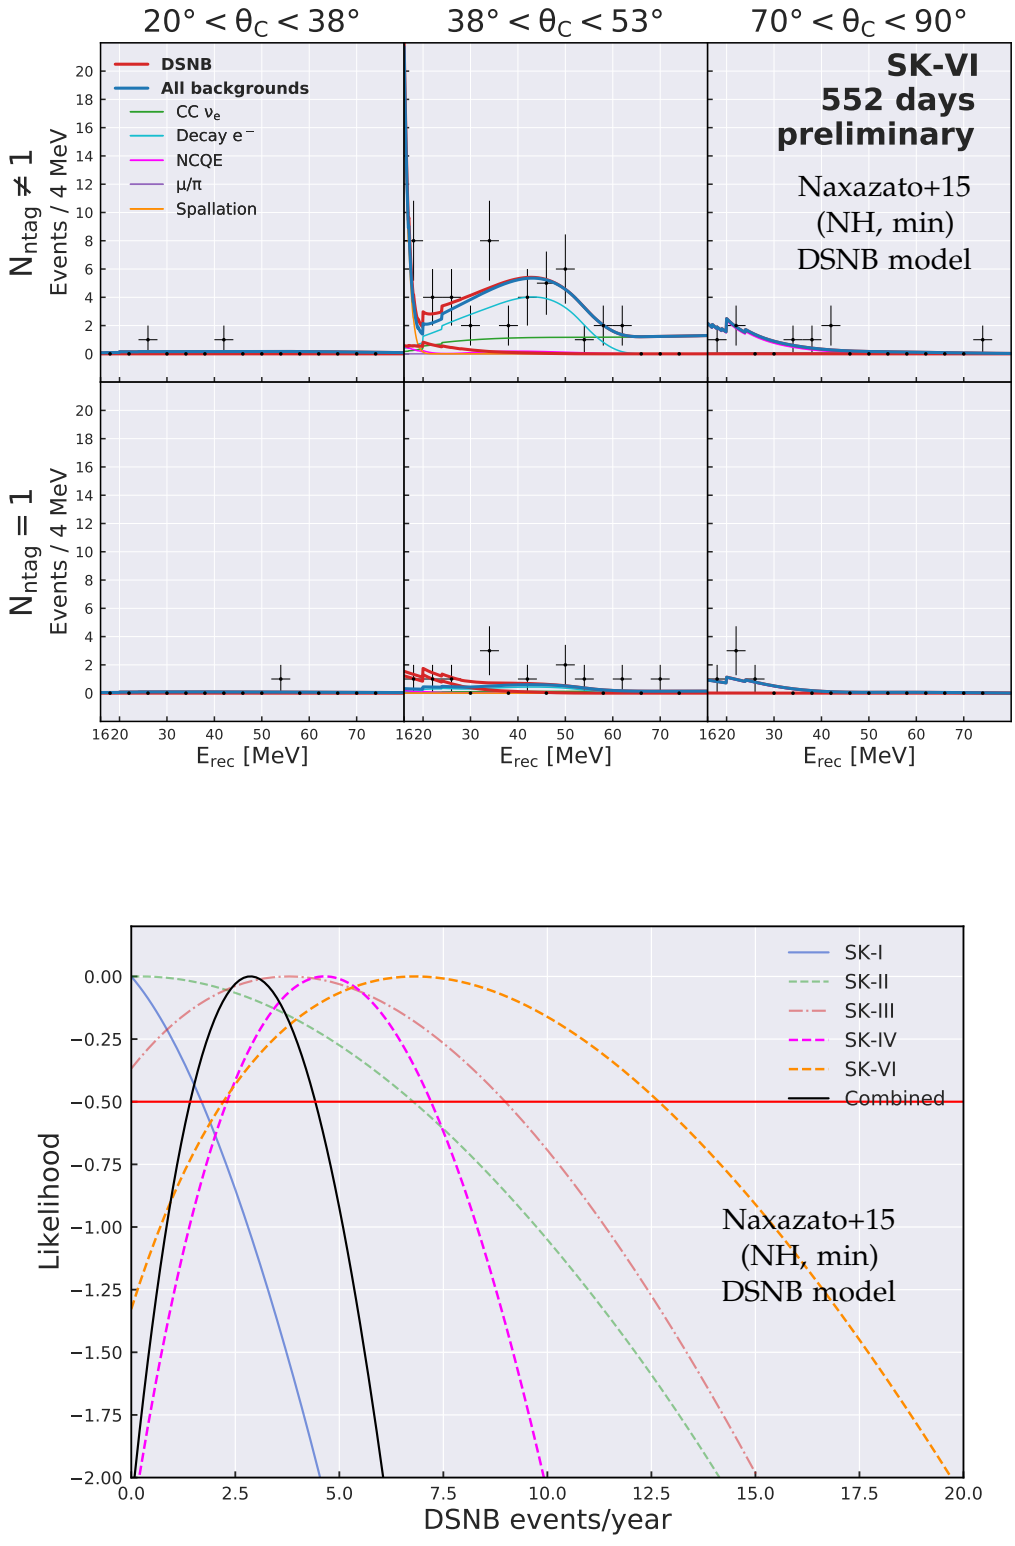

FIGURE E.3: Top:Best-fit signal and background spectra (solid lines) for the SK-VI period, assuming the DSNB flux predicted by the Horiuchi+21 model [26], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I, II, III, IV, and VI of SK, as well as the combined likelihood.

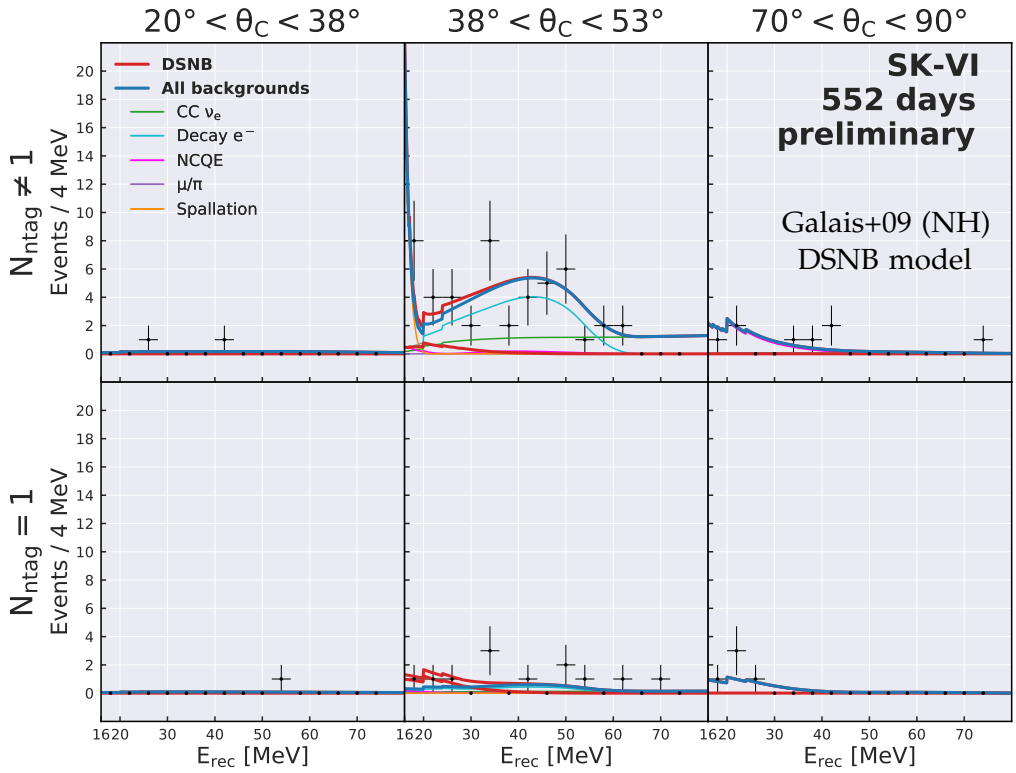

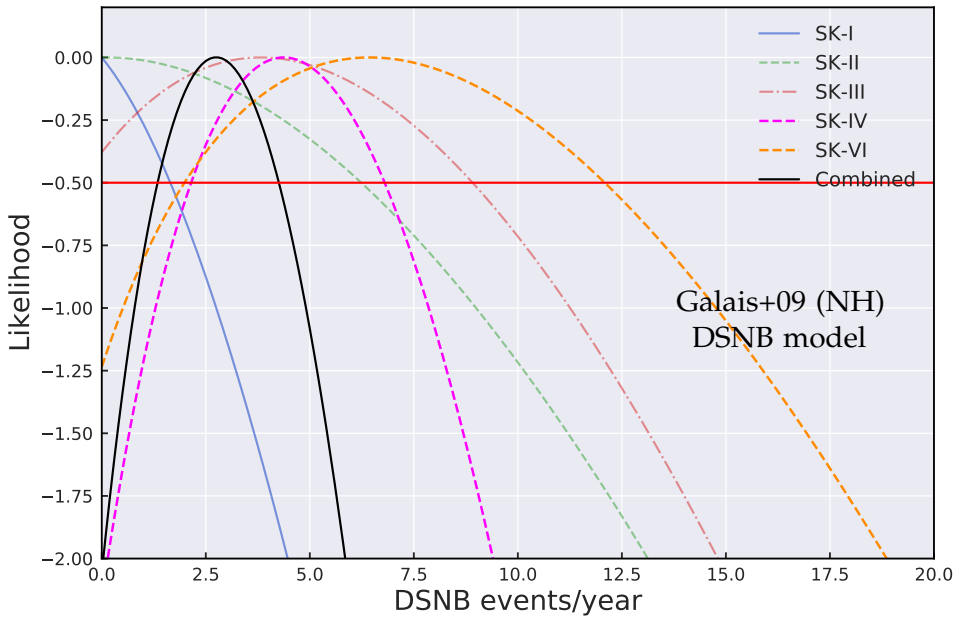

FIGURE E.4: Top:Best-fit signal and background spectra (solid lines) for the SK-VI period, assuming the DSNB flux predicted by the Horiuchi+21 model [22], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I, II, III, IV, and VI of SK, as well as the combined likelihood.

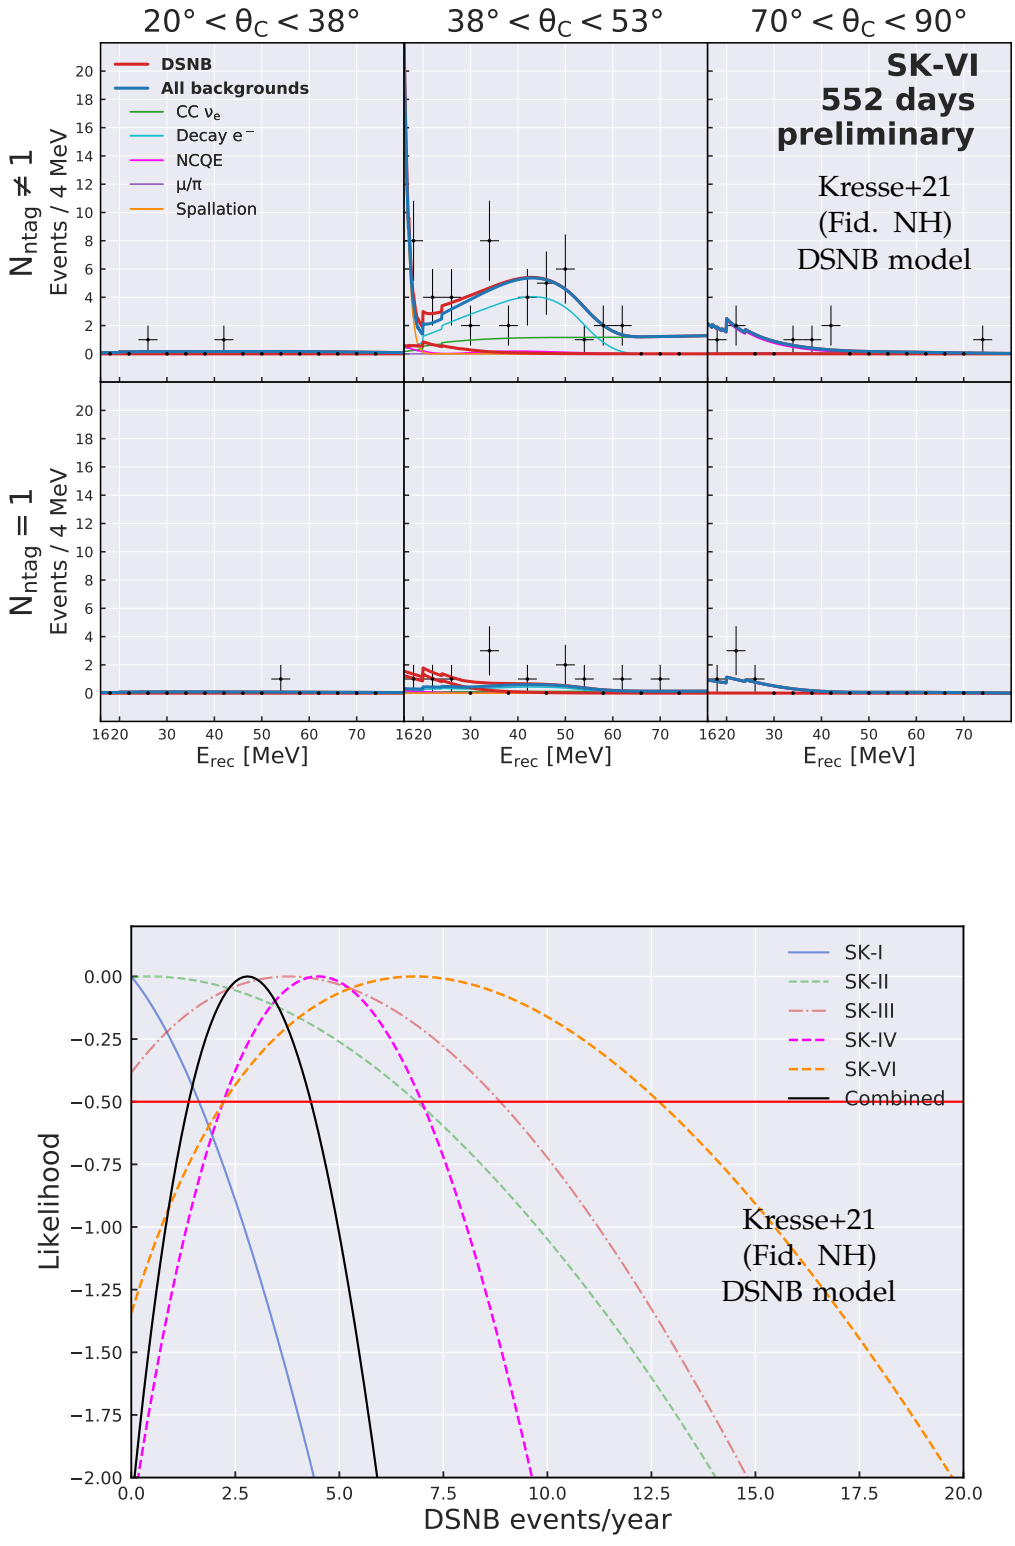

FIGURE E.5: Top:Best-fit signal and background spectra (solid lines) for the SK-VI period, assuming the DSNB flux predicted by the Horiuchi+21 model [25], along with the energy spectrum of the remaining events in data after reduction. The six parameters space regions are the two signal regions and four sidebands, as described in table 4.1. Here the data is binned for illustrative purposes only, since the analysis is unbinned. Bottom: Likelihoods associated with phases I, II, III, IV, and VI of SK, as well as the combined likelihood.

Bibliography

[1] R. P. Feynman. “Space-time approach to nonrelativistic quantum mechanics”. In: *Rev. Mod. Phys.* 20 (1948), pp. 367–387.

[2] J. Schwinger. “On Quantum-Electrodynamics and the Magnetic Moment of the Electron”. In: *Phys. Rev.* 73 (1948), pp. 416–417.

[3] S. Tomonaga. “On a relativistically invariant formulation of the quantum theory of wave fields”. In: *Prog. Theor. Phys.* 1 (1946), pp. 27–42.

[4] F. J. Dyson. “The Radiation Theories of Tomonaga, Schwinger, and Feynman”. In: *Physical Review* 75.3 (Feb. 1949), pp. 486–502.

[5] C. N. Yang and R. L. Mills. “Conservation of Isotopic Spin and Isotopic Gauge Invariance”. In: *Physical Review* 96.1 (Oct. 1954), pp. 191–195.

[6] S. L. Glashow. “Partial Symmetries of Weak Interactions”. In: *Nucl. Phys.* 22 (1961), pp. 579–588.

[7] S. Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266.

[8] D. Galbraith. *UX: Standard Model of the Standard Model*. Online. 2012.

[9] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. en. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29.

[10] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. en. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 30–61.

[11] D. DeCamp et al. “Determination of the number of light neutrino species”. en. In: *Physics Letters B* 231.4 (Nov. 1989), pp. 519–529.

[12] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 978-0-19-850871-7.

[13] Super-Kamiokande Collaboration et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Physical Review Letters* 81.8 (Aug. 1998), pp. 1562–1567.

[14] P. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01.

[15] KATRIN Collaboration. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. en. In: *Nature Physics* 18.2 (Feb. 2022), pp. 160–166.

[16] S. P. Mikheyev and A. Y. Smirnov. “Resonance enhancement of oscillations in matter and solar neutrino spectroscopy”. In: *Yadernaya Fizika* 42 (Jan. 1985), pp. 1441–1448.

[17] L. Wolfenstein. “Neutrino oscillations in matter”. In: *Physical Review D* 17.9 (May 1978), pp. 2369–2374.

[18] K. Zuber. *Neutrino physics*. 2. ed. Series in high energy physics, cosmology and gravitation. Boca Raton, Fla. [u.a.]: CRC Press / Taylor & Francis, 2012. XVIII, 448. ISBN: 9781420064711.

[19] S. Horiuchi et al. “Impact of binary interactions on the diffuse supernova neutrino background”. In: *Phys. Rev. D* 103.4 (2021), p. 043003.

[20] S. Horiuchi, J. F. Beacom, and E. Dwek. “The Diffuse Supernova Neutrino Background is detectable in Super-Kamiokande”. In: *Phys.Rev.D79:083013,2009* (Dec. 2008). arXiv: 0812.3157 [astro-ph].

[21] S. Ando and K. Sato. “Relic neutrino background from cosmological supernovae”. In: *New J. Phys.* 6 (2004), p. 170. arXiv: astro-ph/0410061.

[22] S. Galais et al. “Shock waves in supernovae: New implications on the diffuse supernova neutrino background”. In: *Physical Review D* 81.5 (Mar. 2010), p. 053002.

[23] Z. Tabrizi and S. Horiuchi. “Flavor Triangle of the Diffuse Supernova Neutrino Background”. In: *JCAP* 05 (2021), p. 011.

[24] C. Lunardini. “Diffuse Neutrino Flux from Failed Supernovae”. In: *Physical Review Letters* 102.23 (June 2009), p. 231101.

[25] D. Kresse, T. Ertl, and H.-T. Janka. “Stellar Collapse Diversity and the Diffuse Supernova Neutrino Background”. en. In: *The Astrophysical Journal* 909.2 (Mar. 2021), p. 169.

[26] K. Nakazato et al. “SPECTRUM OF THE SUPERNOVA RELIC NEUTRINO BACKGROUND AND METALLICITY EVOLUTION OF GALAXIES”. en. In: *The Astrophysical Journal* 804.1 (May 2015), p. 75.

[27] S. Horiuchi et al. “Diffuse supernova neutrino background from extensive core-collapse simulations of 8100ăM progenitors”. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (Mar. 2018), pp. 1363–1374.

[28] K. Hirata et al. “Observation of a neutrino burst from the supernova SN1987A”. In: *Physical Review Letters* 58.14 (Apr. 1987), pp. 1490–1493.

[29] R. Svoboda et al. “Neutrinos from Supernova 1987A in the IMB Detector”. In: 26 (Jan. 1987), p. 229.

[30] E. N. Alekseev et al. “Detection of the Neutrino Signal from Supernova 1987A Using the INR Baksan Underground Scintillation Telescope”. In: 26 (Jan. 1987), p. 237.

[31] F. R. Stephenson and D. A. Green. *Historical Supernovae and their Remnants*. Oxford, Clarendon Press, 2002.

[32] T. J. Loredo and D. Q. Lamb. “Bayesian analysis of neutrinos observed from supernova SN 1987A”. In: *Physical Review D* 65.6 (Feb. 2002), p. 063002.

[33] D. N. Schramm and J. W. Truran. “New physics from supernova 1987A”. en. In: *Physics Reports* 189.2 (May 1990), pp. 89–126.

[34] S. A. Kharusi et al. “SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy”. en. In: *New Journal of Physics* 23.3 (Mar. 2021), p. 031201.

[35] S. Abe et al. “Limits on Astrophysical Antineutrinos with the KamLAND Experiment”. en. In: *The Astrophysical Journal* 925.1 (Jan. 2022), p. 14.

[36] Super-Kamiokande Collaboration et al. “Supernova relic neutrino search at super-Kamiokande”. In: *Physical Review D* 85.5 (Mar. 2012), p. 052007.

[37] Y. Suzuki. “The Super-Kamiokande experiment”. In: *The European Physical Journal C* 79 (Apr. 2019).

[38] H. Georgi and S. L. Glashow. "Unity of All Elementary-Particle Forces". In: *Physical Review Letters* 32.8 (Feb. 1974), pp. 438–441.

[39] H. Georgi, H. Quinn, and S. Weinberg. "Hierarchy of Interactions in Unified Gauge Theories". In: *Physical Review Letters* 33 (Aug. 1974), pp. 451–454.

[40] K. S. Hirata et al. "Observation of ^{8}B solar neutrinos in the Kamiokande-II detector". In: *Physical Review Letters* 63.1 (July 1989), pp. 16–19.

[41] Super-Kamiokande collaboration. *Super-Kamiokande collaboration insitutes*. Online. 2022.

[42] S. Fukuda et al. "The Super-Kamiokande detector". en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 501.2 (Apr. 2003), pp. 418–462.

[43] H. Nishino et al. "The new front-end electronics for the Super-Kamiokande experiment". In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 1. Oct. 2007, pp. 127–132.

[44] S. Yamada et al. "Commissioning of the New Electronics and Online System for the Super-Kamiokande Experiment". In: *IEEE Transactions on Nuclear Science* 57.2 (Apr. 2010), pp. 428–432.

[45] L. N. Machado et al. *Pre-Supernova Alert System for Super-Kamiokande*. Tech. rep. arXiv, May 2022.

[46] J. F. Beacom and M. R. Vagins. "GADZOOKS! Antineutrino Spectroscopy with Large Water Cerenkov Detectors". In: *Phys.Rev.Lett. 93 (2004) 171101* (Sept. 2003). arXiv: hep-ph/0309300 [hep-ph].

[47] M. Ikeda. "Status of Super-Kamiokande Gadolinium project". en. In: *Journal of Physics: Conference Series* 2156.1 (Dec. 2021), p. 012150.

[48] K. Abe et al. "First gadolinium loading to Super-Kamiokande". en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1027 (Mar. 2022), p. 166248.

[49] Super-Kamiokande Collaboration et al. "Diffuse supernova neutrino background search at Super-Kamiokande". In: *Physical Review D* 104.12 (Dec. 2021), p. 122002.

[50] A. Coffani. "New studies on cosmogenic induced spallation background for Supernova relic neutrino search in the Super-Kamiokande experiment". PhD thesis. Dec. 2021.

[51] K. Scholberg. "Cross Section Measurements for Supernova Neutrinos". In: vol. 12. JPS Conference Proceedings 12. Journal of the Physical Society of Japan, Dec. 2016.

[52] K. Scholberg. *SNOwGLoBES*.

[53] A. Drukier and L. Stodolsky. "Principles and applications of a neutral-current detector for neutrino physics and astronomy". In: *Physical Review D* 30.11 (Dec. 1984), pp. 2295–2309.

[54] C. J. Horowitz, K. J. Coakley, and D. N. McKinsey. "Supernova observation via neutrino-nucleus elastic scattering in the CLEAN detector". In: *Physical Review D* 68.2 (July 2003), p. 023005.

[55] A. Strumia and F. Vissani. "Precise quasielastic neutrino/nucleon cross-section". en. In: *Physics Letters B* 564.1 (July 2003), pp. 42–54.

[56] P. Vogel and J. F. Beacom. "Angular distribution of neutron inverse beta decay, $\overline{\nu}_e + p \rightarrow e^+ + n$". In: *Physical Review D* 60.5 (July 1999), p. 053003.

[57] J. N. Bahcall, A. M. Serenelli, and S. Basu. "New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes". In: 621 (2005), pp. L85–L88.

[58] A. Goldsack. *SKReact v1.3*. Mar. 2022.

[59] M. Baldoncini et al. "Reference worldwide model for antineutrinos from reactors". In: *Physical Review D* 91.6 (Mar. 2015), p. 065002.

[60] International Atomic Energy Agency. *The Power Reactor Information System (PRIS) and Its Extension to Non-electrical Applications, Decommissioning and Delayed Projects Information*. Tech. rep. 2005.

[61] F. P. An et al. "Improved measurement of the reactor antineutrino flux and spectrum at Daya Bay". en. In: *Chinese Physics C* 41.1 (Jan. 2017), p. 013002.

[62] M. Honda et al. "Atmospheric Neutrino Flux Tables". In: (2011).

[63] Y. Hayato. "A Neutrino Interaction Simulation Program Library NEUT". In: *Acta Phys. Pol. B* 40.2477 (2009).

[64] M. Smy. "Low Energy Event Reconstruction and Selection in Super-Kamiokande-III". In: *30th International Cosmic Ray Conference*. Vol. 5. July 2007, pp. 1279–1282.

[65] Z. Conner. "A study of solar neutrinos using the Super-Kamiokande detector". PhD thesis. University of Maryland, 1997.

[66] The T2K Collaboration et al. "Measurement of neutrino and antineutrino oscillations by the T2K experiment including a new additional sample of ${\ensuremath{\nu}}_{e}$ interactions at the far detector". In: *Physical Review D* 96.9 (Nov. 2017), p. 092006.

[67] Super-Kamiokande Collaboration et al. "Solar neutrino measurements in Super-Kamiokande-IV". In: *Physical Review D* 94.5 (Sept. 2016), p. 052010.

[68] Super-Kamiokande Collaboration et al. "Search for Supernova Relic Neutrinos at Super-Kamiokande". In: *Physical Review Letters* 90.6 (Feb. 2003), p. 061101.

[69] The Super-Kamiokande Collaboration et al. "Solar neutrino measurements in Super-Kamiokande-II". In: *Physical Review D* 78.3 (Aug. 2008), p. 032002.

[70] R. Akutsu. "A Study of Neutrons Associated with Neutrino and Antineutrino Interactions on the Water Target at the T2K Far Detector". PhD Thesis. Tokyo U., 2019.

[71] Super-Kamiokande Collaboration et al. "Solar neutrino measurements in Super-Kamiokande-I". In: *Physical Review D* 73.11 (June 2006), p. 112001.

[72] Super-Kamiokande Collaboration et al. "Solar neutrino results in Super-Kamiokande-III". In: *Physical Review D* 83.5 (Mar. 2011), p. 052010.

[73] T. Totani and K. Sato. "Spectrum of the relic neutrino background from past supernovae and cosmological models". en. In: *Astroparticle Physics* 3.4 (Aug. 1995), pp. 367–376.

[74] M. Kaplinghat, G. Steigman, and T. P. Walker. "Supernova relic neutrino background". In: *Physical Review D* 62.4 (July 2000), p. 043001.

[75] B. Jegerlehner, F. Neubig, and G. Raffelt. "Neutrino oscillations and the supernova 1987A signal". In: *Physical Review D* 54.1 (July 1996), pp. 1194–1203.

[76] J. Meija et al. "Isotopic compositions of the elements 2013 (IUPAC Technical Report)". en. In: *Pure and Applied Chemistry* 88.3 (Mar. 2016), pp. 293–306.

[77] S. F. Mughabghab. "Thermal neutron capture cross sections resonance integrals and g-factors". English. In: (Feb. 2003).

[78] International Atomic Energy Agency. *Database of Prompt Gamma Rays from Slow Neutron Capture for Elemental Analysis*. Tech. rep. IAEA, 2006.

[79] M. Harada. "Geant4 based Simulation Study for Super-Kamiokande". en. In: *Journal of Physics: Conference Series* 1468.1 (Feb. 2020), p. 012255.

[80] K. Hagiwara et al. "Gamma-ray spectrum from thermal neutron capture on gadolinium-157". In: *Progress of Theoretical and Experimental Physics* 2019.2 (Feb. 2019), p. 023D01.

[81] T. Tanaka et al. "Gamma-ray spectra from thermal neutron capture on gadolinium-155 and natural gadolinium". In: *Progress of Theoretical and Experimental Physics* 2020.4 (Apr. 2020), p. 043D02.

[82] L. Marti et al. "Evaluation of gadoliniums action on water Cherenkov detector systems with EGADS". en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 959 (Apr. 2020), p. 163549.

[83] D. Maksimovi, M. Nieslony, and M. Wurm. "CNNs for enhanced background discrimination in DSNB searches in large-scale water-Gd detectors". en. In: *Journal of Cosmology and Astroparticle Physics* 2021.11 (Nov. 2021), p. 051.

[84] A. Y. "Measurement of Neutrino and Antineutrino Neutral-Current Quasielastic-like Interactions and Applications to Supernova Relic Neutrino Searches". PhD thesis. Kyoto University, 2020.

[85] W. T. Winter et al. "The $^{8}\mathrm{B}$ neutrino spectrum". In: *Physical Review C* 73.2 (Feb. 2006), p. 025503.

[86] Super-Kamiokande Collaboration et al. *New Methods and Simulations for Cosmogenic Induced Spallation Removal in Super-Kamiokande-IV*. Tech. rep. arXiv, Nov. 2021.

[87] G. Carminati. "The new Wide-band Solar Neutrino Trigger for Super-Kamiokande". en. In: *Physics Procedia*. 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013 61 (Jan. 2015), pp. 666–672.

ECOLE
DOCTORALE

Titre : À la recherche du fond diffus de neutrinos de supernovas avec Super-Kamiokande: dernières analyses et perspectives pour l'ère du Gd

Mots clés : Neutrinos, supernova, DSNB, Super-Kamiokande, neutrons, astrophysique

Résumé : L'expérience Super-Kamiokande, située au Japon, est une expérience permettant l'étude des neutrinos produits par diverses sources (accélérateur, solaire, atmosphérique) et notamment provenant de sources astrophysiques, en particulier ceux émis par l'effondrement d'une étoile massive lors d'une supernova. En 1987, le prédécesseur de Super-Kamiokande a détecté, pour la seule fois dans l'histoire de l'humanité, l'émission d'une dizaine de neutrinos issus d'une supernova proche située dans le grand nuage de Magellan, un phénomène ne survenant que deux à trois fois par siècle. L'un des enjeux majeurs actuels est de détecter pour la première fois le flux de neutrinos baignant le cosmos et émis par l'ensemble des supernovas produites depuis le début de notre univers. La détection de ce flux, appelé fond diffus de neutrinos de supernovas (DSNB pour Diffuse Supernova Neutrino Background), permettrait en effet d'affiner les modèles des supernovae, ainsi que les modèles cosmologiques en fournissant des indications sur l'histoire de la formation des étoiles, la nucléosynthèse et l'évolution de l'univers.
Cette thèse constitue une analyse de recherche du DSNB dans l'expérience Super-Kamiokande, un détecteur Cherenkov à eau où le canal de détection privilégié est l'interaction beta-inverse, un anti-neutrino électronique interagissant avec un proton pour produire un positron et un neutron. Une attention particulière a été portée sur la réduction des bruits de fonds en développant des stratégies originales basées sur la détection conjointe du positron et du neutron. L'analyse des données enregistrées entre septembre 2008 et juin 2018 (phase SK-IV), a permis de déterminer une limite supérieure du flux de neutrinos du DSNB qui constitue à ce jour la meilleure limite mondiale, excluant ainsi certains modèles théoriques de flux DSNB les plus optimistes. Un deuxième volet original de cette thèse consiste à étendre l'analyse aux données les plus récentes obtenues en présence de sels de gadolinium dissous dans l'eau depuis 2020 (phase SK-Gd), renforçant le signal d'un neutron dans le détecteur. Ce travail montre que la sensibilité obtenue avec les données SK-Gd après environ trois ans de fonctionnement, soit jusqu'à la mi-2023, devrait être statistiquement équivalente à celle obtenue avec la phase SK-IV, ouvrant ainsi la perspective de la découverte du DSNB à moyen terme. Cette thèse offre également un premier aperçu des nouvelles données SK avec Gd, montrant une meilleure sensibilité au DSNB, et une augmentation de la signification de l'excès observé.

Title : Hunting for the Diffuse Supernova Neutrino Background with Super-Kamiokande: latest searches and prospects for the Gd era

Keywords : Neutrinos, supernova, DSNB, Super-Kamiokande, neutrons, astrophysics

Abstract : The Super-Kamiokande experiment, located in Japan, allows for the study of neutrinos from a variety of sources. One of the interests of the experiment is the study of astrophysical neutrinos, and in particular of neutrinos produced during core-collapse supernovae, marking the death of massive stars. In 1987, Super-Kamiokande's predecessor experiment detected a dozen neutrinos originating from a close-by supernova in the Large Magellanic Cloud. The 1987 supernova remains the only instance of supernova neutrino detection to date. Indeed, supernovas that are near enough for neutrino detection are a rare occurrence: roughly three supernovae per century are expected in the Milky Way. In addition to direct supernova detection, one of the main challenges for the experiment is the first detection of the integrated flux from all past supernovae in the observable universe. Detection of this flux, also known as Diffuse Supernova Neutrino Background (DSNB), would allow for the improvement of our understanding of supernovae, as well as the refinement of current cosmological models, providing hints about the history of star formation, about nucleosynthesis, and about the evolution of the universe.
In this thesis, an analysis searching for the DSNB using Super-Kamiokande is presented. In this water Cherenkov detector, the main detection channel for supernova neutrinos is the inverse beta decay, in which an electron antineutrino interacts with a proton, producing a positron and a neutron. Particular attention is given to the background reduction obtained with the detection of the neutron in coincidence with the positron. The analysis of the data from the SK-IV period (between September 2008 and June 2018) allowed for the determination of upper limits on the DSNB flux that are currently the world's tightest, excluding several optimistic models. The second component of this thesis consists in extending the analysis to the most recent data, after doping the detector water with Gadolinium sulfate in 2020 (also known as the SK-Gd phase), significantly strengthening the signal of a neutron in the detector. This work shows that the sensitivity of an SK-Gd search, after just three years of runtime (i.e. in mid-2023), should be statistically equivalent to that obtained with SK-IV, paving the way for discovery of the DSNB in the medium term. This thesis also offers an early look at new SK data with Gd, showing an improved sensitivity to the DSNB, and an increased significance of the observed excess.

Institut Polytechnique de Paris
91120 Palaiseau, France