Reactor Neutrinos in a Gadolinium-Loaded Super-Kamiokande



Alexander Goldsack

Oriel College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary Term, 2022

"*Be curious on your journey.*"
- Bells, Outer Wilds

Acknowledgements

I am first of all thankful to Oxford University and Kavli IPMU for providing a DPhil program that is perfectly matched to what I wanted out of a graduate degree. I'm very glad I have been given the opportunity to work in such a welcoming collaboration so I would also like to thank everyone in Super-K who has helped me in the last few years, especially within the LowE group.

I found nothing but support from my wonderfully-patient partner Sam in my decision to choose a degree that took me to the other side of the world for two years. After returning right at the height of the COVID-19 pandemic, she went above and beyond to ease my non-academic burdens as best as she could while I slowly became an anxious, thesis-writing lump and for that I am eternally grateful.

For all his support getting me to this point, I cannot thank my dad enough. I hold somewhat of a grudge, however, given that the countless hours opening my eyes to the wonders of the universe, be it through documentaries, fostering my love for science fiction or simply lying in the back garden looking at the stars have resulted in an obsession with science that means that I was crazy enough to attempt a doctorate. He let me live in his house when coming back from Japan, however, paying rent with KFC, so maybe I shouldn't be so harsh.

I am thankful to all the Kamioka staff for welcoming me to Japan, especially Masa Ishikura for helping me get settled and deal with everything where my broken Japanese failed me. I am also thankful to Lluis Marti for giving me frequent excuses to go to Shogun Burger and for being an all round good friend to me while in Toyama. I'd also like to thank him, Michael Smy and Scott Locke for their patience with my frequent LowE/WIT questions.

For the many hours-long discussions about physics, living in Japan, and life in general, as well as being an overall great friend, I would very much like to thank Laurence Cook. I think Starbucks would also like to thank him too, given that most of those discussions occurred over coffee and frappuccinos.

Many thanks to Hannah, Matthieu, Raf and especially Ryan Randolph Riches for the weekly insanity that was our D&D sessions. Special thanks to Ryan for being such a good friend, as well as introducing me to most of the people in Japan who I came to call my friends, including Jean Shin and Josh Mueda.

My sanity in early 2020 would not have fared well without the many hours playing Civilisation with Ellie Jones and Aidan Reynolds. Thanks to them both, as well as Amy Hughes and Rory Evans for welcoming me into the D&D group.

For all things Python, I thank my long term coding guru Ryan Cross, who never liked me in the first place, a fact he reminds me of frequently. I also thank him, Cameron, Chris and Dec for helping me unwind in the evenings by making me stress about my poor performance in WoW.

Finally, I would very much like to thank my supervisors, Professors Dave Wark and Mark Vagins, for the endless advice and anecdotes and for generally being in my corner in many ways during my time as a student.

Abstract

In July of 2020, Super-Kamiokande (SK), a 50 kt water Cherenkov detector, entered its new phase: SK-Gd [1]. In this phase, the detector's once-ultrapure water was doped with 0.02% gadolinium sulphate octahydrate by mass, increased to 0.06% in July 2022. This provides far greater sensitivity to neutron captures, thereby increasing the detection efficiency of electron antineutrinos interacting via inverse beta decay. One such source of these neutrinos is the numerous fission products in nuclear reactor cores. To date, only one reactor neutrino interaction has been identified in a large scale water Cherenkov detector [2], in part due to their energy spectrum extending less than 10 MeV above the Cherenkov threshold of positrons in water. A search has been performed for reactor neutrinos in data obtained by SK's Wideband Intelligent Trigger (WIT) system from December 2020 to October 2021, using a boosted decision tree-based selection method. The search intended to identify an increase in event rate consistent with that expected due to the restart of 9 Japanese reactor cores between December of 2020 and October of 2021. Apparent issues with the WIT data between March and August of 2021 motivated the comparison of the rates measured in the data before and after this period. Though the measured rate up to March 2021 of $(1.6 \pm 0.2)\,\mathrm{d}^{-1}$ was consistent with predictions, the rate measured between August and October 2021 was found to be $(8.8 \pm 0.4)\,\mathrm{d}^{-1}$, far in excess of the predicted rate of $3.7\,\mathrm{d}^{-1}$. While not inconsistent with the expected total interaction rate of $9.75\,\mathrm{d}^{-1}$, an unreasonably large signal efficiency would be required for this increase to be solely explained by reactor restarts. One suspected cause of this increase is the persistence of issues with the WIT data from March 2021, continuing after August, when the data was previously thought to be once again stable. As such, this work cannot conclusively say that reactor neutrinos have been identified in SK; however, the Monte Carlo studies presented in this thesis indicate that the experiment has this capability.

Contents

List of Figures viii

List of Tables xiv

1 Introduction 1

2 Neutrinos 4
- 2.1 History and Oscillation . . . 5
- 2.2 Mass Ordering and The MSW Effect . . . 9
- 2.3 Neutrino Sources . . . 11
  - 2.3.1 Cosmic Neutrino Background . . . 11
  - 2.3.2 Solar . . . 12
  - 2.3.3 Atmospheric . . . 12
  - 2.3.4 Accelerators . . . 14
  - 2.3.5 Supernovae and Pre-Supernova Stars . . . 16
  - 2.3.6 Geo-neutrinos . . . 20
  - 2.3.7 Cosmogenic . . . 21
  - 2.3.8 Reactor . . . 22
- 2.4 Inverse Beta Decay . . . 22

3 Super-Kamiokande 26
- 3.1 Detector Structure . . . 27
- 3.2 Detection Principle . . . 32
- 3.3 Electronics and DAQ . . . 34
- 3.4 Data Periods . . . 36
- 3.5 BONSAI Low Energy Event Reconstruction . . . 37
- 3.6 The Wide-Band Intelligent Trigger (WIT) . . . 41
- 3.7 Monte-Carlo . . . 53

4 A Gadolinium Loaded Super-K 56
- 4.1 Motivation . . . 56
- 4.2 Water Systems . . . 57
- 4.3 EGADS . . . 59

4.4 Deployment and Validation 61
4.5 WIT Coincidence Trigger 63

5 Reactor Neutrinos 72
5.1 Nuclear Reactors 72
5.2 Nuclear Power in Japan 79
5.3 Neutrino Production 82
5.4 Reactor Neutrino Experiments 85
5.5 Monitoring Nuclear Reactors 89
5.6 SK-Gd and Reactor Neutrinos 91
5.7 Simulating Reactor Neutrinos: SKReact 93

6 Restart Detection 100
6.1 Introduction and Objective 100
6.2 Sources of Background 105
6.3 Cut Optimisation with SPLASH 108
6.4 BDT Selection 117
6.5 Background Prediction and Data Availability 127
6.6 Results 136

7 Conclusion 144

Appendices

A Japanese and South Korean Nuclear Reactors 148

B Coefficients for Antineutrino Emission Spectra of Reactor Fuels 151

C Fuel Contribution Fractions 153

References 155

List of Figures

2.1 The proton-proton chain and CNO cycle, modified from [17]. 13
2.2 A diagram of the accelerator neutrino beam production process, including magnetic horns for focusing the desired pions and rejecting the oppositely charged pions to produce a neutrino or antineutrino beam. Adapted from [23]. 14
2.3 The luminosity (top left) and average energy (bottom left) of ν_e, $\bar{\nu}_e$ and ν_x (all other flavours) with time from core "bounce" due to neutron degeneracy pressure, and the energy spectrum of $\bar{\nu}_e$ at different times from this bounce (right) [36]. 18
2.4 The spectra of geo-neutrinos produced in the decays 2.26 (^{238}U chain, solid black), 2.27 (^{232}Th chain, dash-dotted red) and 2.28 (^{40}K chain, dashed blue) [45]. The solid vertical line indicates the inverse beta decay interaction threshold of 1.806 MeV. 21
2.5 The cross section for inverse beta decay on free protons, calculated using Strumia and Vissani's parameterisation given in [58]. 24

3.1 A cross section of the SK detector (partially filled with water). . . . 28
3.2 Photograph of the structure which supports the PMTs and separates the ID and OD. The fibreglass rear of the ID PMT cases can be seen within the structure. Taken during the 2018 tank open work. 29
3.3 Photographs of the SK inner detector taken during the 2018 tank open work. 29
3.4 A photograph from the bottom of the SK OD looking up, showing how Tyvek is used to optically isolate the OD from the ID and encourage reflection within the OD. The Hamamatsu R1408 PMTs and wavelength-shifting light collection plates can also be seen. Taken during the 2018 tank open work. 30
3.5 The formation of a Cherenkov ring in a medium of refractive index n by a charged particle of velocity $v > c/n$. The particle is either created or given energy by a neutrino interaction at A and stops producing Cherenkov radiation at B. 33
3.6 A sharp muon-like Cherenkov ring event (a) compared to a fuzzier electron-like ring (b) for two T2K events in SK [24]. 34

3.7 The Super-Kamiokande data acquisition (DAQ) system, in place since SK-IV except for the WIT PCs which were added in 2015. . . 35
3.8 STORE Goodness for simulated events and raw WIT data (BG). Events below the black line (in the hatched region) are rejected in WIT. 45
3.9 Radial position for simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by Clus-Fit. Events above the black line (in the hatched region) are rejected in WIT. . . 46
3.10 Height of simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by Clus-Fit. Events outside of the two black lines (in the hatched region) are rejected in WIT. . 47
3.11 Radial position for simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by BONSAI. The disconnect in the BG distribution at the trigger threshold is a symptom of the WIT process applying FV cuts to some events, and is not physical. Events above the black line (in the hatched region) are rejected in WIT. 48
3.12 Height of simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by BONSAI. The disconnect in the BG distribution at the trigger threshold is a symptom of the WIT process applying FV cuts to some events, and is not physical. Events outside of the two black lines (in the hatched region) are rejected in WIT. 49
3.13 BONSAI goodness of simulated events and raw WIT data (BG) that pass the preceding stages of WIT. Events below the black line (in the hatched region) are rejected by WIT. 50
3.14 N18 of simulated events and raw WIT data (BG) that pass the preceding stages of WIT. Events below the black line (in the hatched region) are rejected by WIT. 51

4.1 Schematic of the SK-Gd water systems [1]. 58
4.2 The EGADS hall. 60
4.3 The conductivity of water samples taken at a range of heights throughout the SK ID over the course of first Gd loading, in 2020. The sampling system encountered issues between the 23^{rd} and 29^{th} of July. 62
4.4 An example of neutron capture times and the resultant fit for 30 mins of AmBe calibration data with the source placed in the centre of the detector. Measurements were carried out for a range of source placements. 64

4.5 The logic of the WIT coincidence trigger, including the trigger identification numbering scheme. 66
4.6 The distance between prompt and delayed event vertices for simulated gadolinium and hydrogen capture, reconstructed by online BONSAI, compared to the distance between sequential events in a sample of WIT data considered to be background.Gadolinium concentration is taken to be that after first loading (0.02% by mass). 67
4.7 The time difference between prompt and delayed events for simulated gadolinium and hydrogen capture, reconstructed by online BONSAI, compared to the time between sequential events in a sample of WIT data considered to be background. Gadolinium concentration is taken to be that after first loading (0.02% by mass). 68
4.8 WIT detection efficiencies when relaxing the hit threshold (and therefore STORE goodness threshold) for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs. 69
4.9 WIT detection efficiencies when relaxing the N18 threshold for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs. 70
4.10 WIT detection efficiencies when relaxing the BONSAI goodness threshold for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs. 71

5.1 The simulated electron-antineutrino spectra with associated errors produced by individual fuels in a 3423 MW core of the Ohi 3 PWR reactor. 85
5.2 Ratio $R^{\text{exp}}_{a,HM}$ of measured and expected IBD yields for a selection of reactor experiments as a function of reactor-detector distance L for the Huber-Mueller model of reactor antineutrino spectra, taken from [94]. The horizontal green band shows the average ratio $\bar{R}_{HM}$ and its uncertainty, indicating a 2.5σ deviation from the expected. . . . 89
5.3 Left: Allowed parameter regions (at 1σ, 90%, 2σ, 99%, 3σ confidence limit for 2 degrees of freedom) for fixed $\theta_{13} = 8.6°$ from the analysis of KamLAND data (solid green contours with the best fit marked by a green star) and from the combined analysis of solar data for the GS98 model (full regions with the best fit marked by black star) and the AGSS09 model (dashed void contours with the best fit marked by a white dot). The orange contours are the results of a global analysis for the GS98 model using the previous SK4 neutrino data. Right: $\Delta\chi^2$ dependence on Δm^2_{21} for the same four analyses after marginalizing over θ_{12}. Taken from NuFit 5.1 (2021) [96]. 90

5.4 The simulated neutrino emission spectrum on production s^{-1} from the Ohi-3 pressurised water reactor (PWR) core. 94
5.5 The simulated 2018 reactor spectrum after oscillation s^{-1} at SK, with the Ohi-3 pressurised water reactor (PWR) core highlighted. . 96
5.6 The Gaussian fit (solid lines) of BONSAI reconstructed energy for simulated positrons of energies from 1 to 10 MeV, area normalised to the efficiency of detection. Dotted lines at true energy given for reference. 97
5.7 The process for converting interacted spectrum to reconstructed spectrum. 98
5.8 The interacted and smeared spectrum over the course of 2018, with the Ohi-3 core highlighted, simulated by SKReact. 99
5.9 The SKReact GUI. 99

6.1 The length of the data files in the WIT pair data, calculated using the hardware trigger counter differences in each file. 102
6.2 The length of the data files in the WIT pair data, calculated using the hardware trigger counter differences in each file, removing files of length greater than 87.38 s. 102
6.3 The size of the data files in the WIT pair data, removing files of length greater than 87.38 s. 103
6.4 The simulated reactor neutrino interaction rate from late 2020 to late 2021, with the ramp up simulated using sigmoid functions. A ramp up of 62 days is used in this analysis. 105
6.5 Vertex z as reconstructed by BONSAI for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data. 112
6.6 Vertex r as reconstructed by BONSAI for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data. 113
6.7 BONSAI reconstructed energy for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data. 113
6.8 OVAQ goodness variable for simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data. 114

6.9 EffWall for simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data. 115
6.10 BONSAI reconstructed distance between simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data. 115
6.11 Simulated neutron capture times corresponding to 0.02% Gd concentration (MC). BG is time between prompt and delayed events in raw pre-reactor restart WIT data, identified by the WIT coincidence trigger. 116
6.12 The input variable importance for a BDT trained on the 2019 reactor spectrum, including the β isotropy parameters. 123
6.13 The input variable importance for a BDT trained on the simulated 2021 reactor spectrum, the BDT which formed the final reactor sample selection. 124
6.14 The final tree in the ensemble BDT used in this analysis. 125
6.15 The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG). 126
6.16 The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG), with a >0.2 BDT score cut. 127
6.17 The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG), with a >0.95 BDT score cut. 128
6.18 A comparison of all BDT variables for a 24h run in the pre (85766) and post (86102) datasets. 131
6.19 A comparison of the BDT score between events in a 24h run in the pre (run 85766) and post (run 86102) datasets. 132
6.20 The predicted BG and reactor rate over the entire dataset, binned weekly, with no file quality requirements. Errors are statistical only. 133
6.21 The predicted BG and reactor rate over the entire dataset, weekly binned, with file quality requirements. Errors are statistical only. . . 135
6.22 The pair rate in the pre and post periods, compared to the expected BG and reactor rates. The red line represents the reactor rate with 100% efficiency. Errors are statistical only. 136
6.23 The pair rate in the pre and post periods, weekly binned, compared to the expected BG and reactor rates, with the livetime calculated using assumed 100% run lengths. The red line represents the reactor rate with 100% efficiency. Errors are statistical only. 137
6.24 The required signal (MC) efficiency in the pre period needed to explain the observed pair rate increase for a range of BG rate increases and relative signal efficiency increases. 138

6.25 The required signal (MC) efficiency in the pre period needed to explain the observed pair rate increase for a range of BG rate increases where the MC efficiency increases as a factor of the BG increase. . . 139
6.26 The difference between the known run length and run length derived from the sum of all calculated file lengths. 140
6.27 The pair rate in the pre and post periods, weekly binned, compared to the expected BG and reactor rates, with the livetime calculated using the file lengths as opposed to assumed 100% livetime run lengths. Errors are statistical only. 141
6.28 The pair rate across the whole WIT pair dataset, without quality cuts which would remove the pre→post intermediary data. Livetimes are calculated using the summed file lengths. Errors are statistical only. 142
6.29 The pair rate across the whole WIT pair dataset, without quality cuts which would remove the pre→post intermediary data. Livetimes are calculated using the run lengths. Errors are statistical only. . . 143

List of Tables

3.1 The data-taking phases of SK. 38

3.2 Hit-counter trigger thresholds and their approximate trigger rates in SK as of September 2021. The range of hits for the triggers do not overlap; for example, the SLE trigger only activates for events with 34-46 hits in the trigger window. The OD trigger is not independently monitored and is only considered when an ID trigger is activated. . 42

5.1 Operable commercial nuclear power plant designs throughout the world [73]. The column "GWe" refers to the total electrical power capacity, measured in GW. 77

5.2 Nuclear power reactors currently operating in Japan and South Korea, ordered by $(P_{th})/(R_{SK})^2$ where P_{th} is thermal power output capacity of the reactor and R_{SK} is the distance from the core to the SK detector [75]. An asterisk by a plant's name denotes a South Korean plant. 81

5.3 Japanese reactor core restarts during the SK-Gd phase, ordered by date of restart. The expected increase in IBD events in the SK ID due to the restart is included, as well as the expected pre-restart rate, calculated from SKReact as detailed in section 5.7. 92

6.1 Criteria of radioactive impurities in $Gd_2(SO_4)_3 \cdot 8\,H_2O$ powder and the physics goals driving the requirements [1]. 106

6.2 The predicted pre/post results based on pre offline data and post online. An asterisk denotes a predicted value. 135

A.1 All nuclear reactors in Japan and South Korea, ordered by $(P_{th})/(R_{SK})^2$ where P_{th} is thermal power output capacity of the reactor and R_{SK} is the distance from the core to the SK detector [75]. An asterisk by a plant's name denotes a South Korean plant. A status of pending refers to plants that have applied for a restart and are pending approval from either the local government or the NRA. 148

B.1 Coefficients α_{pk} of the Mueller *et al.* fifth-order polynomial parameterisation of antineutrino flux from elements $k = {}^{235}U$, ${}^{238}U$, ${}^{239}Pu$ and ${}^{241}Pu$. In the column $\delta\alpha_{pk}$ the 1σ error on α_{pk} are given [78]. . 151

C.1 ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu fission (f_i) or power (p_i) fractions for a range of reactor designs, compiled by [101]. 154

List of Abbreviations

AAS Atomic absorption spectrometer

ABWR Advanced boiling water reactor

AFT After (trigger)

AGR Advanced gas-cooled reactor

ANN Artificial neural network

ATM Analogue timing module

BDT Boosted Decision Tree

BG Background

BNO Baksan Neutrino Observatory

BONSAI Branch Optimisation Navigating Successive Annealing Iterations

BooNE Booster Neutrino Experiment

BOREX Boron Solar Neutrino Experiment

BSM Beyond the Standard Model (of particle physics)

BWR Boiling water reactor

CνB Cosmic Neutrino Background

CC Charged current (interaction)

CCSN Core collapse supernova

CL Confidence limit

CMB Cosmic microwave background

CNN Convolutional neural network

CP Charge-parity (operator)

CP-1 Chicago Pile-1 (reactor)

DAQ Data acquisition

DC Double Chooz

DirKS Directional KS (test)

DOF Degrees of freedom

DSNB Diffuse supernova neutrino background

DT Deuterium-tritium (neutron generator)

DUNE Deep Underground Neutrino Experiment

EGADS Evaluating Gadolinium's Action on Detector Systems

FBR Fast breeder reactor

FIFO First-in-first-out

FNR Fast neutron reactor

FOM Figure of merit

FPGA Field programmable gate array

FV Fiducial volume

GADZOOKS! Gadolinium Antineutrino Detector Zealously Outperforming Old Kamiokande, Super!

GCR Gas cooled reactor

GdLS Gadolinium doped liquid scintillator

GEANT-3 Geometry and Tracking

GRB Gamma ray burst

GUI Graphical user interface

GUT Grand Unified Theory

GZK Greisen-Zatsepin-Kuzmin (neutrinos)

HE High energy (trigger)

HEU Highly enriched uranium

HK Hyper-Kamiokande

HM Huber-Mueller (model)

HPGe High-purity germanium

HV High voltage

IAEA International Atomic Energy Agency

IBD Inverse beta decay

ICRR Institute for Cosmic Ray Research

ID Inner detector

IH Inverted hierarchy (neutrino mass)

ILL Institute Laue-Langevin

IMB Irvine–Michigan–Brookhaven (experiment)

INGRID Interactive Neutrino GRID

IO Inverted ordering (neutrino mass)

IPMU Kavli Institute for the Physics and Mathematics of the Universe

JUNO Jiangmen Underground Neutrino Observatory

K2K KEK-to-Kamioka

KamiokaNDE Kamioka Nucleon Decay Experiment

KamLAND Kamioka Liquid Scintillator Antineutrino Detector

KamLAND-Zen Kamioka

KEK The High Energy Accelerator Research Organization

KS Kolmogorov-Smirnov (test)

LArTPC Liquid argon time projection chamber

LE Low energy (trigger)

LEU Low-enriched uranium

LF Load factor

LINAC Linear accelerator

LMC Large Magellanic Cloud

LS Liquid scintillator

LSND Liquid Scintillator Neutrino Detector

LWGR Light water graphite-moderated reactor

$M_\odot$ Mass of the Sun

MC Monte Carlo

ML Machine learning

MSG Multiple Scattering Goodness

MSW Mikheyev-Smirnov-Wolfenstein (effect).

MW Mega watt (unit)

NC Neutral current (interaction)

ND280 Near Detector at 280 m

NH Normal hierarchy (neutrino mass)

NO Normal ordering (neutrino mass)

NOvA NuMI Off-axis ν_e Appearance

NPT Treaty on the Non-Proliferation of Nuclear Weapons

NRA Nuclear Regulation Authority

NuMI Neutrinos at the Main Injector

OD Outer detector

OVAQ One-dimensional Variable of Vertex and Angular Reconstruction Quality

pe Photo-electrons (unit)

PHWR Pressurised heavy water reactor

PMNS Pontecorvo-Maki-Nakagawa-Sakata (matrix)

PMT Photo-multiplier tube

PRIS Power Reactor Information System

PWR Pressurised water reactor

QAC Charge to analogue converter

QBEE QTC-based electronics with ethernet

QTC Quantum tunelling composite

RAA Reactor antineutrino anomaly

RENO Reactor Experiment for Neutrino Oscillation

SBN Short-Baseline Neutrino (program)

SHE Super high energy (trigger)

SK Super-Kamiokande

SKG4 SK Geant4 (MC software)

SLE Super low energy (trigger)

SM Standard Model (of particle physics)

SN Supernova

SNO Sudbury Neutrino Observatory

SPLASH Super-K Python Library of Analysis Software and something beginning with H

SRN Supernova relic neutrinos

SSM Standard solar model

STEREO Sterile Reactor Neutrino Oscillations

STORE Simultaneous Trigger and Online Reconstruction

T2K Tokai-to-Kamioka

TAC Time to analogue converter

TOC Total organic carbon reduction

TOF Time of flight

TPC Time projection chamber

UCI University of California, Irvine

UDEAL Underground Device Evaluating Attenuation Length

UF Ultra filtration

UHECR Ultra high energy cosmic rays

UV Ultraviolet

WatChMaL Water Cherenkov Machine Learning

WATCHMAN Water Cherenkov Monitor for Antineutrinos

WC Water Cherenkov

WIT Wide-band Intelligent Trigger

1 Introduction

Now is an exciting time for the Super-Kamiokande (SK) experiment and the world of neutrino physics research in general. While the field is gearing up for an era of extraordinary-scale, high-precision experiments with the approaching Hyper-Kamiokande and DUNE experiments, SK, the 50 kt water Cherenkov detector now in its 26th year of operation, has expanded the neutrino horizon in an altogether different direction.

In the summer of 2020, SK entered its SK-Gd phase where its once ultrapure water was doped with 0.02% gadolinium sulphate octahydrate by mass, providing a neutron capture target with $\sim 100,000\times$ the cross section of pure water, and one which releases a signal that SK can detect with almost 100% accuracy. Neutrons are produced in a range of processes in SK, but of particular interest is the inverse beta decay reaction, where an electron antineutrino interacts with a proton to produce a positron and a neutron.

Though the primary motivation for the SK-Gd project is to detect the spectrum of neutrinos produced by distant supernovae, this thesis explores its capabilities to detect electron antineutrinos produced with great intensity in the cores of terrestrial nuclear reactors nearly 200 km from the detector. Such a search has never been performed in SK due to the difficulty in detecting these low energy neutrinos of a few MeV, and to date only one reactor neutrino event has ever been identified in

a water Cherenkov detector, during the SNO+ experiment's pure water phase [2]. The vastly increased sensitivity to neutron capture and therefore inverse beta decay makes the detection of reactor neutrinos possible for the first time in SK.

An ideal stage is also set for this analysis in terms of nuclear reactor activity. Following the 2011 Tohoku earthquake and subsequent meltdown of three cores of the Fukushima Daiichi plant, Japan has been operating at a fraction of its nuclear power capacity as reactors have undergone the long process of conforming to increased safety standards and seeking approval from local governments to approve their restarts. 2021 saw a number of these restarts, most notably some select reactors of the Takahama, Ohi and Mihama plants, all less than 200 km from SK.

There are two main motivations for this search, the first is that of a purely academic nature. In the world of reactor neutrino experiments, most detectors are situated up to ∼1 km from their reactor complex of interest and only one sought to measure the oscillation probability from a range of plants at a longer distance: the Kamioka Liquid Scintillator Antineutrino Detector (KamLAND) experiment, situated in the same mine as Super-Kamiokande. KamLAND provides the only non-solar probe of Δm^2_{21} and therefore the only measurement using antineutrinos as opposed to neutrinos, a measurement that somewhat disagrees with the global average. Prior to an update to the day/night asymmetry results from SK in 2020, there was a 2σ difference between the combined SK and SNO results and KamLAND's measurement, however this difference has now been reduced to around 1.4σ [3]. While this holds little statistical significance, it remains that KamLAND is the only experiment to probe the ∼200 km baseline reactor neutrino realm. A repetition of this measurement with SK would therefore affirm that this difference is indeed put to rest.

The second motivation for detecting reactor neutrinos at SK is the coveted opportunity to provide a real-world application of particle physics. This is due to its potential to support nuclear non-proliferation. Proving that a neutrino detector can monitor the activity of a nuclear reactor at a significant standoff

distance provides a new tool in the world's arsenal in the effort to police the illegal production of nuclear weapons.

Given the unprecedented nature of this search, this thesis describes the various new tools and analysis approaches produced in an attempt to detect reactor neutrinos for the first time in SK. Chapter 2 provides a historical and theoretical introduction to neutrinos. Chapter 3 then details the SK detector and its approach to detecting these elusive particles. This chapter also introduces a crucial utility of the reactor neutrino search: the Wideband Intelligent Trigger (WIT) system. The major evolution of SK into the SK-Gd phase is described in chapter 4, including the introduction of a coincidence trigger into WIT in support of this search which seeks to improve the efficiency of detecting IBD pairs with little impact on the WIT data stream. Nuclear reactors and their neutrino production is described in chapter 5 and here the SKReact software is detailed. This is a reactor neutrino simulator developed to support all analyses concerning reactor neutrinos in SK. The multivariate approach taken to attempt to isolate a sample of reactor neutrinos from the vast swathes of WIT data is presented in chapter 6, utilising the newly-developed SPLASH framework, firstly through traditional methods and then through the use of boosted decision trees. The results of the search are also presented here, as is a discussion of the consistency of the detected rates with those expected at the start and end of 2021.

2
Neutrinos

Contents

2.1 History and Oscillation 5
2.2 Mass Ordering and The MSW Effect 9
2.3 Neutrino Sources 11
2.3.1 Cosmic Neutrino Background 11
2.3.2 Solar 12
2.3.3 Atmospheric 12
2.3.4 Accelerators 14
2.3.5 Supernovae and Pre-Supernova Stars 16
2.3.6 Geo-neutrinos 20
2.3.7 Cosmogenic 21
2.3.8 Reactor 22
2.4 Inverse Beta Decay 22

Neutrinos (ν) are spin-$\frac{1}{2}$, electromagnetically neutral fermions of non-zero but small mass. They interact via the weak force as one of three leptonic "flavours" (ν_e, ν_μ and ν_τ) corresponding to the three charged leptons. The Standard Model (SM) of particle physics predicts that in a given weak interaction vertex, a neutrino can only interact with a neutrino or charged lepton of the same flavour; however, they exhibit a property not predicted by the Standard Model, where their flavour and mass eigenstates are mixed. This allows them to oscillate between different flavours as they propagate. A flavour eigenstate different to that at the time of

production can therefore be observed at some distance from the source. They have a significantly lower mass than any other massive particle. Model-dependent cosmological measurements constrain the neutrino mass to $<0.17\,\text{eV}$ [4], while measurements of beta decay in tritium set a limit of $1.1\,\text{MeV}$ [5]. Their total interaction cross section is far smaller than any other particle in the Standard Model. Neutrino experiments therefore often focus on maximising the integrated flux of neutrinos through their detector volumes, be it through proximity to their source or in constructing the largest-volume particle detectors in the world.

2.1 History and Oscillation

In 1930 Woflgang Pauli postulated that a weakly-interacting, neutrally charged particle could account for the missing energy observed in β^- decay [6]. This missing energy was evident in the continuous energy spectrum of the outgoing electron and had no clear explanation at the time. It wasn't until 1956 that Clyde L. Cowan and Frederick Reines confirmed its existence by detecting electron antineutrinos produced in a nuclear reactor [7]. The neutrinos interacted via the interaction

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{2.1}$$

known as inverse beta decay (IBD), detailed in section 2.4, with protons in tanks of cadmium chloride-doped water, producing a positron and a neutron. The positrons would annihilate, producing a pair of gammas in tanks of liquid scintillator sandwiching the water tanks. Photo-multiplier tubes (PMTs) would detect this light, identifying the positron. To confirm the PMT signal was from a positron produced in IBD, the neutron must also be detected in coincidence, motivating the cadmium-chloride doping of the water; cadmium has a large neutron capture cross section. The neutron would capture on the cadmium, producing another gamma up to 17 microseconds after the pair annihilation which would also be detected by the PMTs in the scintillator. The identification of this IBD signal at a rate proportional to the power output of the reactor confirmed the existence of neutrinos.

Headed by Raymond Davis Jr. in the 1960s, the Homestake experiment sought to measure the flux of neutrinos emitted by fusion reactions in the Sun by counting argon atoms produced in the reaction

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e^-. \tag{2.2}$$

John N. Bahcall provided a theoretical value from the standard solar model (SSM), but the experiment only measured a flux of around a third of Bahcall's calculation [8]. This gave rise to the so-called solar neutrino problem, which saw Bahcall's calculations and Davis' experimental methods come under scrutiny over the following years to try and explain the discrepancy.

At the time, neutrinos were thought to be massless, but one potential explanation for the flux deficit was that if neutrinos had mass, they could "oscillate" between flavours as they propagate. The idea of neutrino oscillation analogous to kaon oscillation was proposed by Bruno Pontecorvo in 1957 [9], with the formal theoretical neutrino mixing framework being developed by Ziro Maki, Masami Nakagawa and Shoichi Sakata in 1962 [10].

In this framework, the neutrino flavour eigenstates $\nu_{e,\mu,\tau}$ are mixed with the mass eigenstates $\nu_{1,2,3}$ via the "Pontecorvo-Maki-Nakagawa-Sakata" (PMNS) matrix U_{PMNS} like

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{PMNS} \times \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \tag{2.3}$$

The PMNS matrix itself is defined as

$$U_{PMNS} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}, \tag{2.4}$$

though it is often parameterised as a combination of three rotation matrices, one for each rotation in phase between eigenstates. In this form it is given by

$$U_{PMNS} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix} \tag{2.5}$$

where $s_{ij} = \sin\theta_{ij}$, $c_{ij} = \cos\theta_{ij}$ and θ_{ij} is the mixing angle between states i and j. δ_{CP} is the CP violating phase, determining the difference in oscillation between neutrinos and antineutrinos. δ_{CP} has potentially grand implications on the formation of the universe and the matter-antimatter asymmetry problem in some models, and is the focus of investigation for many large scale neutrino experiments.

Equation 2.5 assumes Dirac neutrinos, where the neutrino and antineutrino have different lepton numbers; however, it is possible that they are Majorana particles which would mean that neutrinos are their own antiparticle. In this case, two extra phases $\alpha_{1,2}$ must be included to fully describe this parameterisation of the PMNS matrix; however, this property of the neutrino has not been confirmed and so this thesis will exclusively consider Dirac neutrinos.

The evolution of a neutrino mass eigenstate from $x, t = 0$ to $x = x$, $t = t$ is given by

$$|\nu(x,t)\rangle = e^{-i\phi} |\nu(0,0)\rangle \tag{2.6}$$

where $\phi = Et - px$ is the phase of the state with energy and momentum E and p. It can be shown that for $\phi = Et + Px$, the difference $\Delta\phi = \phi_1 - \phi_2$ for two neutrino flavour states after propagating some distance $x = L$ is given by

$$\Delta\phi \approx \frac{\Delta m^2 L}{2E} \tag{2.7}$$

where $\Delta m^2 = m_1^2 - m_2^2$.

The probability of a neutrino with energy E initially of flavour α being detected as flavour β at some distance L from the source is given by

$$P_{\alpha\to\beta}(L,E) = |\langle\nu_\beta|\nu_\alpha\rangle|^2. \tag{2.8}$$

If a two neutrino universe is considered and δ_{CP} is assumed to be 0 for simplicity (though recent results hint towards a maximal value of δ_{CP} [11]), using equations 2.6 and 2.7 and a two-flavour version of the PMNS matrix, this can be rewritten as

$$P_{\alpha\to\beta}(L,E) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{2E}\right) \qquad \alpha \neq \beta \tag{2.9}$$

where θ is the mixing angle between the two mass states. The trigonometric dependence of probability with distance is the reason this phenomenon is dubbed neutrino "oscillation".

The Super-Kamiokande (SK) experiment and the Sudbury Neutrino Observatory (SNO) provided the crucial evidence required to confirm neutrino oscillation.

The "smoking gun" evidence for neutrino oscillation was SK's discovery that the flux of atmospheric neutrinos demonstrated a zenith-angular dependence. These neutrinos are produced when primary cosmic rays interact with the upper atmosphere, detailed in 2.3.3, producing a mix of muon and electron neutrinos and antineutrinos at a ratio of $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e) \approx 2$. In the absence of oscillation, this flux ratio should be isotropic at SK as these interactions occur approximately evenly around the Earth, and the neutrinos will have negligible interaction with the Earth itself. Instead, there was a significant difference in the flux of up-going muon neutrinos (created on the far side of the Earth travelling up through the Earth to SK) compared to down-going [12]. This measurement is consistent with the idea that the up-going neutrinos oscillated away to other flavours, while the down-going neutrinos created in the atmosphere above the detector primarily remain as muon-type due to their shorter propagation length.

To definitively confirm that the neutrinos had oscillated away, the flux of all neutrino flavours must be shown to be consistent with Bahcall's predictions. SK is only sensitive to ν_e and ν_μ, but SNO used heavy water (^2H_2O) as its target material, permitting the neutral current (NC) interaction with deuterium ^{2}H

$$\nu_\alpha + {}^2\mathrm{H} \to p + n + \nu_\alpha, \tag{2.10}$$

which is sensitive to all flavours of neutrino. If solar neutrinos were indeed oscillating to different flavours by the time they reached Earth, SNO would still be able to make a complete measurement of the solar neutrino flux in this interaction

channel. SNO's measurement was consistent with Bahcall's prediction for the total flux produced by the solar ^{8}B fusion chain [13].

SNO's verification of the total solar neutrino flux and SK's measurement of the zenith angle dependency of atmospheric neutrino flux confirmed the existence of neutrino oscillation. The 2015 Nobel Prize in Physics was jointly awarded to Arthur B. McDonald of SNO and Takaaki Kajita of SK for this discovery.

2.2 Mass Ordering and The MSW Effect

The oscillation probability $P_{\alpha\to\beta}$ (equation 2.9) introduces a problem when trying to obtain a complete measurement of all of the oscillation parameters. As the Δm^2 term is contained within a $\sin^2$ function, measurement of $P_{\alpha\to\beta}$ alone is insufficient to determine its sign. Expanding to three flavours and therefore three mass eigenstates results in further degeneracy in the form of three Δm^2 terms, introducing the neutrino mass hierarchy problem. In vacuum oscillation as derived above, only the differences in magnitude of the mass states can be measured and not the sign of these differences, nor their absolute masses. It is therefore unclear which of the three mass states is the heaviest and which is the lightest.

This problem has been partially solved thanks to the modification of the oscillation probability when traversing matter. While all flavour eigenstates undergo coherent forward scattering off of electrons in matter, the electron flavour eigenstate can do this via the charged current interaction channel, unavailable to the muon and tau flavour eigenstates. This modifies the potential experienced by the electron eigenstate exclusively, resulting in an effective potential difference between V_e and V_a for the ν_e and ν_a (other flavour) eigenstates given by

$$V = V_e - V_a = \sqrt{2} G_f n_e \tag{2.11}$$

where G_F is the Fermi coupling constant and n_e is the number density of electrons in the medium [14]. This modifies the Hamiltonian of the system as $H_0 \to H_m = H_0 + V$ leading to a change in the eigenstates ν_1 and ν_2 and their eigenvalues:

$$\nu_1, \nu_2 \rightarrow \nu_{1m}, \nu_{2m} \tag{2.12}$$

$$\frac{m_1^2}{2E}, \frac{m_2^2}{2E} \rightarrow H_{1m}, H_{2m} \tag{2.13}$$

where ν_{1m} and ν_{2m} are the neutrino eigenstates in the presence of matter, with respective Hamiltonians H_{1m} and H_{2m}.

The equation of motion for a two-neutrino state of ν_e and ν_a in matter is therefore given by

$$i\frac{d}{dx}\begin{pmatrix}\nu_e\\ \nu_a\end{pmatrix} = \left[\frac{\Delta m^2}{4E}\begin{pmatrix}-\cos 2\theta & \sin 2\theta\\ \sin 2\theta & \cos 2\theta\end{pmatrix} + \begin{pmatrix}\frac{V}{2} & 0\\ 0 & -\frac{V}{2}\end{pmatrix}\right]\begin{pmatrix}\nu_e\\ \nu_a\end{pmatrix} \tag{2.14}$$

which has clear dependence on the sign of Δm^2. Given the high electron density in the Sun, solar neutrinos experiencing this effect have been utilised to determine the sign of Δm^2_{12}, concluding that $m_2 > m_1$. The sign of Δm^2_{31} is yet to be measured and determining whether neutrinos exhibit a normal hierarchy (NH, $m_3 > m_2 > m_1$) or inverted (IH, $m_2 > m_1 > m_3$) is a key physics goal of some modern experiments. These experiments primarily study electron neutrino appearance in beam or atmospheric neutrino samples that experience matter effects as they propagate through the Earth; however, this measurement is relatively difficult given the vastly lower electron density and propagation length when compared to the solar analysis.

In matter, the flavour states of a two-flavour ν_e and ν_a state can be expressed in terms of the mass eigenstates and a single modified mixing angle θ_m like so:

$$|\nu_e\rangle = \cos\theta_m |\nu_{1m}\rangle + \sin\theta_m |\nu_{2m}\rangle \tag{2.15}$$

$$|\nu_a\rangle = \cos\theta_m |\nu_{2m}\rangle - \sin\theta_m |\nu_{1m}\rangle . \tag{2.16}$$

θ_m can be expressed as

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta}{R} \tag{2.17}$$

where R is the resonance factor, given by

$$R \equiv \left(\cos 2\theta - \frac{2EV}{\Delta m^2}\right)^2 + \sin^2 2\theta. \tag{2.18}$$

For $n_e \to 0$ as in vacuum, $V \to 0$ as there is no longer a modified potential due to the presence of electrons. In this case, $R \to 1$ and $\theta_m \to \theta$, recovering vacuum oscillation. In the limit of very high electron density, $V \to \infty$ and $\sin^2 2\theta_m \to 0$, suppressing any oscillation effects whatsoever.

As $\theta_m = \theta_m(n_e)$, a resonance can occur. The condition for this resonance is given by

$$2EV = \Delta m^2 \cos 2\theta. \tag{2.19}$$

If the electron density changes slowly enough, such as for a solar neutrino exiting the Sun, the resonance condition can be satisfied, which maximises equation 2.17 and maximises oscillation. This resonance with adiabatically changing electron density is known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect and is clearly observed in solar neutrinos [15].

2.3 Neutrino Sources

There exists an abundance of neutrino sources due to the relative ease of their production; neutrinos are exceptionally light and are therefore produced in a large variety of weak interactions involving charged leptons. The nature of the largest contributing sources of neutrinos at Earth are detailed here.

2.3.1 Cosmic Neutrino Background

Neutrinos were created in abundance in the high energies of the very early universe. These neutrinos decoupled from matter around one second after the Big Bang and, thanks to their weak interaction, are expected to permeate the universe today, forming the Cosmic Neutrino Background (CνB). The density of the CνB is predicted through Big Bang cosmology to be extremely high, at about $336\,\mathrm{cm}^{-3}$ [16];

however, they are of a particularly low energy given their age, with a temperature of $\sim$1.95 K, and so CνB have yet to be detected. Their flux is dependent on their energy which is in turn dependent on the neutrino's absolute mass. This is not yet known to a degree where it can be determined whether CνB neutrinos are relativistic, but if they are, they would be the largest neutrino flux contributor at Earth.

2.3.2 Solar

The greatest contribution to the measured flux of neutrinos at Earth is due to nuclear interactions in the Sun. There are a number of fusion reaction chains which produce electron neutrinos, the primary contributor being the proton-proton chain. This begins with the fusion of two protons to form deuterium in the reaction

$$p + p \rightarrow {}^{2}\mathrm{H} + e^{+} + \nu_{e}. \tag{2.20}$$

The resulting deuterium then typically fuses with another proton in the interaction

$${}^{2}\mathrm{H} + p \rightarrow {}^{3}\mathrm{He} + \gamma. \tag{2.21}$$

The helium-3 from this reaction can then interact via one of four branches, detailed in figure 2.1, which also demonstrates an alternative production of helium-3 via the so-called "pep" reaction.

Solar neutrinos are of great interest due to their well understood production and well known and large distance from their source. They have been studied extensively in SK, detailed in [18][19][20] and [21].

2.3.3 Atmospheric

The Earth's upper atmosphere is constantly being bombarded by high energy particles known as (primary) cosmic rays produced throughout the universe from a wide range of sources. The cosmic ray flux is dominated by protons. When these protons interact with the Earth's upper atmosphere, they produce pions. The charged pions then decay, primarily via the interactions

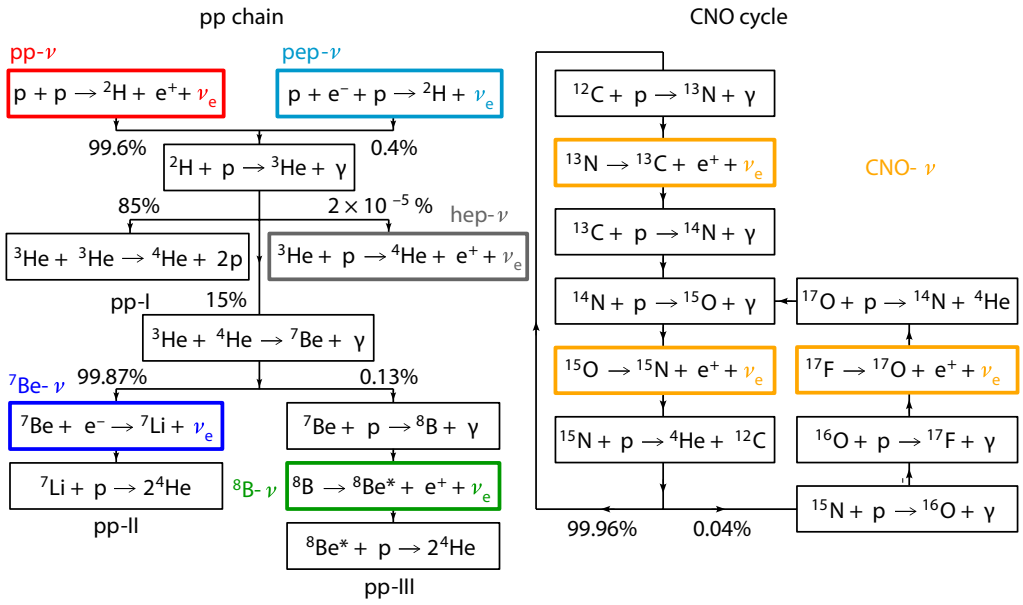


Figure 2.1: The proton-proton chain and CNO cycle, modified from [17].

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \tag{2.22}$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu \tag{2.23}$$

while the neutral pions undergo electromagnetic showering. The muons from the charged pions then also decay, predominantly via the channels

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \tag{2.24}$$

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu \tag{2.25}$$

resulting in a ratio of $(\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e) \approx 2$, with some minor contribution from higher order processes. Atmospheric neutrinos are of particular interest as their flux and energy spectrum can be well constrained through the measurement of their accompanying muon flux. They also provide a variable baseline from production to detection, dependent on zenith angle. This was instrumental in the discovery of neutrino oscillation; SK observed a strong zenith angle dependency on the flux of muon neutrinos [12].

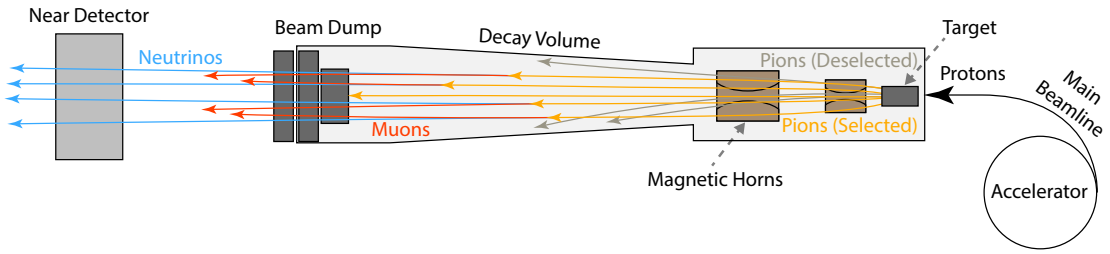


Figure 2.2: A diagram of the accelerator neutrino beam production process, including magnetic horns for focusing the desired pions and rejecting the oppositely charged pions to produce a neutrino or antineutrino beam. Adapted from [23].

2.3.4 Accelerators

While neutrinos themselves cannot be electromagnetically accelerated, the same process which produces atmospheric neutrinos can also be used to create a beam of neutrinos in an accelerator. Protons can be accelerated into a proton-rich target in the direction of a neutrino detector, which produces pions and, subsequently, a beam of neutrinos. There are many benefits to this approach, and SK detects accelerator neutrinos as part of the T2K (Tokai-to-Kamioka) experiment. T2K is in many ways an evolution of the K2K (KEK-to-Kamioka) experiment which directed a neutrino beam from the KEK (The High Energy Accelerator Research Organization) laboratory in Tsukuba, Japan, towards SK. It was the first long-baseline accelerator neutrino experiment, measuring muon neutrino disappearance and paving the way for future accelerator neutrino experiments [22].

Firstly, the beam can be chosen to be primarily composed of either neutrinos or antineutrinos, with some contamination. Once the pions have been created, they travel some distance before decaying. During this time, a magnetic field can be applied to focus either π^+ or π^-, defocusing the other. This allows the beam composition to be switched between neutrino and antineutrino modes by reversing the polarisation of the magnetic selection field; π^+ for ν_μ, π^- for $\bar{\nu}_\mu$. This is crucial for measuring the δ_{CP} parameter, as the oscillation probability of neutrino and antineutrino beams can be directly compared with with relatively well-understood initial fluxes. A diagram of the neutrino beam production process is shown in figure 2.2.

Another advantage of creating a beam of neutrinos is the ability to fine-tune their energy spectrum. As well as modifying the energy of the incident protons themselves, the kinematics of the decay of the selected pions can be utilised. Along the beam axis, the kinematics of the pions' decay results in a broad neutrino energy spectrum. Moving off-axis to the beamline, however, results in a tighter energy spectrum at the expense of reduced flux and average energy. The advantage of a tighter energy spectrum is that the oscillation maxima for neutrinos in the beam in that direction will be similarly tightly spread, given the L/E dependency. Beam experiments can therefore use the off axis angle and baseline from beam source to detector to tune the oscillation maxima to be tightly peaked around the position of their detector, reducing systematic uncertainties. This technique is employed by the T2K (2.5°, $\bar{E} = 0.6$ GeV, $L = 295$ km) [24] and NOvA (0.8°, $\bar{E} = 2.0$ GeV, $L = 810$ km) [25] experiments. Some experiments, such as K2K and the upcoming DUNE experiment [26], opt instead for an on-axis approach, favouring the higher flux and average energy over a tighter energy spectrum.

To constrain the systematic uncertainties in the beam's properties as well as to measure interaction cross sections, supplementary detectors are often deployed close to the beam source. These are typically small scale detectors situated a fraction of the distance from the beamline compared to the primary detector. They are therefore often referred to as near detectors. Primary detectors of experiments are sometimes referred to as far detectors in contrast and are typically significantly larger than the near detectors to account for the loss of flux with distance. Near detectors can either utilise a similar detection technique as the far detector such as in the NOvA experiment, or standalone detector designs, such as T2K's ND280 and INGRID detectors.

Accelerator-based beam experiments are at the forefront of the measurement of the δ_{CP} phase. T2K currently provide the tightest constraint on this important oscillation parameter, favouring maximum CP violation at a confidence level of approximately 3σ [11]. Two major next-generation experiments seek to definitively measure δ_{CP}: Hyper-Kamiokande, a 258 kt evolution of the Super-Kamiokande

detector utilising an upgraded version of the beam used in T2K [27]; and DUNE, a large-scale liquid argon time projection chamber (LArTPC) based in the United States. DUNE will utilise an upgraded version of the high-power NuMI beam produced at Fermilab, Chicago [26], used in the MINERvA [28], MINOS [29] and NOvA [25] experiments.

2.3.5 Supernovae and Pre-Supernova Stars

By far the most instantaneously intense but fleeting source of neutrinos are supernovae (SNe), specifically core-collapse SNe (CCSNe). As a massive ($\gtrsim 8\ M_{\odot}$) star reaches the end of its life, its core of higher-mass materials can collapse under its own immense gravitational pressure, triggering a supernova.

As hydrogen is exhausted in a massive star, its core condenses, increasing the temperature to a point where more massive elements can be fused, balancing the gravitational potential energy with the outward pressure due to fusion. The core itself remains stable due to electron degeneracy pressure. This is the outward pressure which emerges due to Pauli's exclusion principle, causing electrons to resist occupying non-degenerate states in extremely high density materials.

As the supply of lighter elements is exhausted, equilibrium is maintained through the fusion of the heaviest fusible nuclei such as carbon, neon, oxygen and eventually silicon for sufficiently massive stars. Silicon fusion is the final stage of a massive star's evolution, as silicon fuses to iron. Iron does not undergo fusion due to it being the inflection point of nucleon binding energy per nucleon: the point at which fusion becomes an energy sink rather than a source. Once a star begins its carbon burning phase, the high temperature is offset primarily by neutrino cooling, more so with neon and oxygen, peaking with silicon. Neutrino cooling is the process where neutrinos produced through pair production, plasmon decay and neutrino Bremsstrahlung take energy away from the high-temperature stellar core. This is an exceptionally short phase in the star's evolution when considering the timescale of its $\sim 10^7$ years-long life up to this point, from ~ 1000 y for the carbon burning phase down to a few days for silicon burning. As the star begins to fuse silicon, neutrino

production increases to a level that may be detectable by neutrino experiments on Earth, which can be used as an early warning signal of a CCSN.

A pre-SN alarm was recently introduced into SK which estimates that a pre-SN silicon-burning star could be detected by a gadolinium-dope SK, detailed in chapter 4, up to 600 pc away with a false alarm rate of 1 per century [30][31]. KamLAND, a liquid-scintillator neutrino detector based in the same mine as SK, has an implementation of a pre-SN alarm [32].

As the star continues to fuse silicon, the iron core accumulates beyond the Chandrasekhar limit of $1.44\,M_{\odot}$. This mass limit is the point at which the electron degeneracy pressure can no longer resist gravitational collapse. At this point, electron capture on protons in the iron core produces electron neutrinos and neutron rich nuclei. The neutrons in these nuclei can then undergo β^- decay, producing electron antineutrinos. The reduction in electrons in the core results in reduced electron degeneracy pressure, triggering the collapse.

During this period, the core's density is such that neutrinos cannot diffuse outwards faster than the collapse, trapping them. If the mass of the core is more than a few Solar masses, then neutron degeneracy is not sufficient to decelerate the collapse and the core will directly form a black hole with no supernova [33]. Below this mass, the density of the core reaches the point where neutron degeneracy pressure becomes significant and sufficient to decelerate the collapse, releasing the trapped neutrinos over the course of tens of milliseconds and generating a shockwave that propagates outwards into the mantle of the star. This shockwave violently ejects the outer layers of the star, leaving only the core. The outer boundary of this shockwave is the source of the intense light observed from a supernova; however, the exact mechanism required to produce such luminosity is a matter of debate and likely involves some form of absorption of the initial neutrino burst [34][35]. This initial neutrino burst is the highest average energy and flux of neutrinos produced over the course of a CCSN, shown in figure 2.3.

The super-heated core is sufficiently hot where neutrino-antineutrino pairs are formed through thermal emission at a rate greater than through electron capture

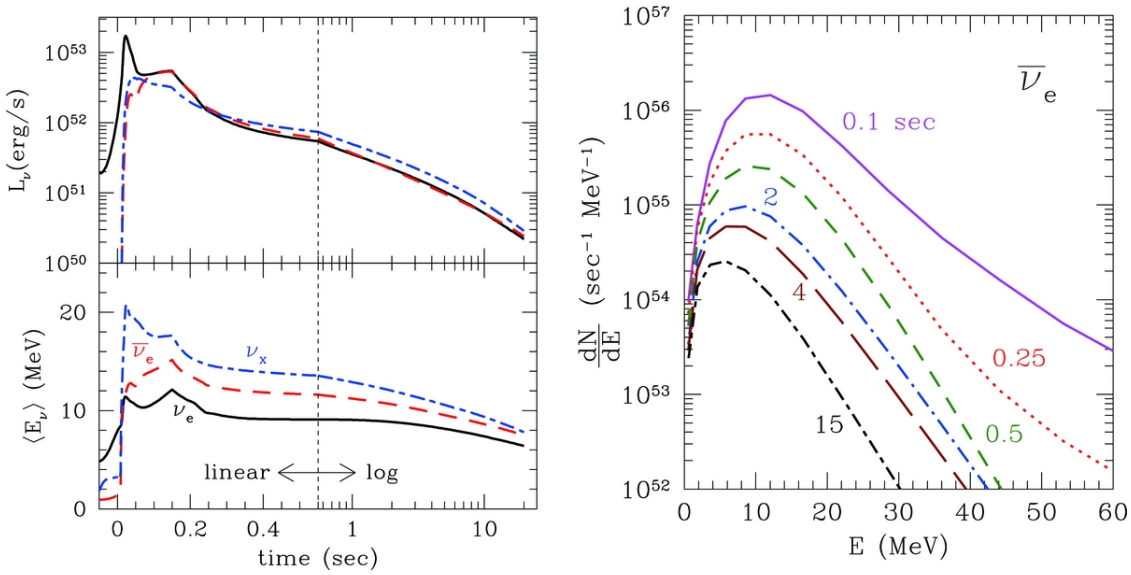


Figure 2.3: The luminosity (top left) and average energy (bottom left) of ν_e, $\bar{\nu}_e$ and ν_x (all other flavours) with time from core "bounce" due to neutron degeneracy pressure, and the energy spectrum of $\bar{\nu}_e$ at different times from this bounce (right) [36].

or β^- decay, rapidly cooling the core over a period of ~ 10 s, after which point neutrino luminosity drops significantly. While the electromagnetic radiation from a CCSN is luminous enough to outshine an entire galaxy, 99% of the total 3×10^{53} erg released by a CCSN is carried away by the neutrinos. As the core cools, the neutrino production rate falls, as does their average energy.

SNe from the Milky Way or nearby galaxies would be a plentiful source of neutrinos. Unfortunately, given the sub-surface, electromagnetically invisible nature of a progenitor star nearing core collapse, their happening is impossible to predict. The detection of silicon burning neutrinos is the only outward warning sign, and these leave little time to prepare for the burst, around a day at most and just a few hours for more distant stars in the Milky Way [30][31]. They are also infrequent. The galactic rate of CCSN is predicted to be 1.63 ± 0.46 $(100\,\mathrm{yr})^{-1}$ based on massive star counts, extragalactic SN rates and ^{26}Al abundance, a CCSN by-product [37]. An alternative approach models the observability of galactic CCSN, considering visible luminosity and absorption due to interstellar dust. This observability model is then combined with historical records of visible SNe to predict a rate of $4.6^{+7.4}_{-2.7}$ $(100\,\mathrm{yr})^{-1}$ [38].

A total of 25 neutrinos produced by CCSNe have been identified to date. All were from SN1987A on February 23, 1987, the only nearby CCSN since the discovery of the neutrino. Kamiokande-II detected 12 [39], the Irvine-Michigan-Brookhaven experiment (IMB) 8 [40] and the Baksan Neutrino Observatory (BNO) 5 [41]. SN1987A occurred approximately 51.4 kpc from Earth, outside of the Milky Way in the Large Magellanic Cloud (LMC), a nearby satellite galaxy. Given the significant progress made in neutrino detection technologies since SN1987A and prevalence of detectors of various technologies around the world, a nearby CCSN in or near the Milky Way would result in a significant increase in SN neutrinos detected. SK alone expects between 2,000 and 3,000 events in its 22.5 kt fiducial volume for an SN burst at 10 kpc using the minimum flux predicted by the Nakazato SN burst model [42]. For many modern neutrino detectors, an important aspect of their operation is maximising livetime, i.e. the time the detector is online and capable of detecting a CCSN.

While detecting the burst from a nearby CCSN with modern detectors would greatly advance our understanding of the core collapse mechanism, an alternative approach is to probe neutrinos produced in supernovae far beyond the local cluster. These are known as supernova relic neutrinos (SRN) given the far older point of emission compared to local CCSN. Galactic supernovae are rare; however, it follows that the wider a net you cast the greater the rate of CCSN, with the payoff of a severely reduced flux for more distant supernovae. This integrated flux forms the diffuse supernova neutrino background (DSNB) with the current upper limit set by SK at 2.7 $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ for neutrino energies greater than 17.3 MeV [43]. This is many orders of magnitude lower than the flux expected from a galactic CCSN, but has the advantage of being constant and acting as a probe into all supernovae within the observable universe. Measuring the DSNB is the one of the main motivations for the SK-Gd project [44].

2.3.6 Geo-neutrinos

Geo-neutrinos are produced when nuclides in the Earth's interior undergo β^- decay. These decays are also responsible for a significant portion of the heat radiated by the Earth's interior, at a rate proportional to the rate of neutrino production. Measuring the production of geo-neutrinos would therefore indicate the contribution of radioactive decay to the $\sim 46\,\mathrm{TW}$ terrestrial surface heat flux, providing valuable information for various studies into the Earth's interior [45].

The majority of the heat produced by radioactive decay in the Earth's interior are from the elements ^{238}U, ^{323}Th and ^{40}K in the decay schemes:

$$^{238}\mathrm{U} \rightarrow {}^{206}\mathrm{Pb} + 8\alpha + 8e^- + 6\bar{\nu}_e + 51.7\,\mathrm{MeV} \tag{2.26}$$

$$^{232}\mathrm{Th} \rightarrow {}^{208}\mathrm{Pb} + 6\alpha + 4e^- + 4\bar{\nu}_e + 42.7\,\mathrm{MeV} \tag{2.27}$$

$$^{40}\mathrm{K} \rightarrow {}^{40}\mathrm{Ca} + e^- + \bar{\nu}_e + 1.31\,\mathrm{MeV} \tag{2.28}$$

with energy spectra shown in figure 2.4.

At low energies, inverse beta decay is the dominant interaction mode for electron antineutrinos, with a threshold of 1.806 MeV. Figure 2.4 shows that this excludes the majority of the ^{238}U and ^{232}Th spectra and the entirety of the ^{40}K spectrum from detection via inverse beta decay. To add to the difficulty of geo-neutrino detection, a nuclear power reactor will dominate the flux of electron antineutrinos close to the IBD threshold unless a detector is sufficiently far away from such a plant. This significantly increases the difficulty of identifying geo-neutrinos in detectors situated in Western Europe, the eastern United States, India and East Asia (particularly Japan, South Korea and the east coast of China) due to the high density of nuclear power plants in these regions. Despite this, the leading measurements of the geo-neutrino flux are jointly provided by the KamLAND (Kamioka Observatory, Japan) and Borexino (Gran Sasso, Italy) experiments [46]. The large scale scintillator experiment JUNO aims to improve upon this during periods where its nearby reactors, its principle interest, are offline [47].

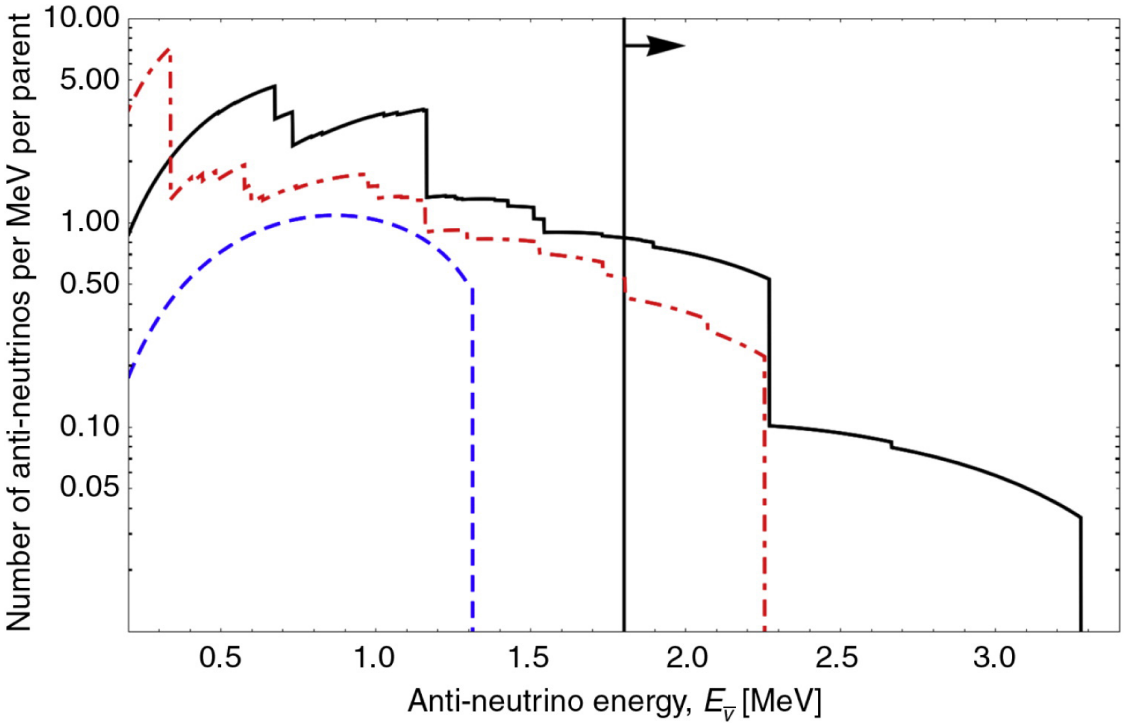


Figure 2.4: The spectra of geo-neutrinos produced in the decays 2.26 (^{238}U chain, solid black), 2.27 (^{232}Th chain, dash-dotted red) and 2.28 (^{40}K chain, dashed blue) [45]. The solid vertical line indicates the inverse beta decay interaction threshold of 1.806 MeV.

2.3.7 Cosmogenic

By far the highest energy neutrinos observed on Earth are those produced by particles accelerated by undetermined astrophysical objects. The first evidence for these most likely extragalactic neutrinos was found at the IceCube neutrino observatory, which also holds the best constraint on their flux [48]. IceCube is a PMT-string based detector that uses a cubic kilometer of Antarctic ice as its interaction medium, aiming to identify Cherenkov radiation produced by TeV-PeV-scale neutrinos, though a subarray of the detector, DeepCore, is sensitive to neutrinos $\gtrsim$10 GeV. IceCube detected neutrinos of energy $1.041^{+0.132}_{-0.144}$ and 1.141^{+143}_{-133} PeV in 2011 and 2012 respectively [49], as well as the highest energy neutrino ever recorded (later in 2012), with an energy of 2.004^{+236}_{-262} [50].

PeV scale neutrinos are predicted to be produced when ultra-high energy cosmic rays (UHECR), typically protons, collide with CMB photons, dubbed Greisen-Zatsepin-Kuzmin (GZK) neutrinos. There are a number of suggested sources of

UHECRs. Gamma ray bursts (GRBs) are one potential source. GRBs are short-lived high-luminosity electromagnetic events, predominantly due to supernovae, but also due to the merger of binary neutron stars. Low-luminosity GRBs may provide mild enough radiation fields with a sufficient energy budget to produce UHECRs; however, there is debate as to if their energies truly are sufficient [51]. Relativistic SNe are also postulated to accelerate particles to UHECR territories.

Blazars are another potential source of UHECRs. These are active galactic nuclei with powerful jets that may have sufficiently strong magnetic fields to accelerate UHECRs. Such powerful jets may also be produced when a star passes close to a supermassive black hole, where the star becomes tidally disrupted, tearing it apart [52]. Pulsars are another source of strong magnetic fields [53], as are galaxies with a high star-formation rate inducing strong magnetic superwinds [54].

A multimessenger collaboration identified that the direction of a high-energy neutrino event detected by IceCube was consistent with a blazar [55], confirming that blazars are a source of high-energy astrophysical neutrinos.

2.3.8 Reactor

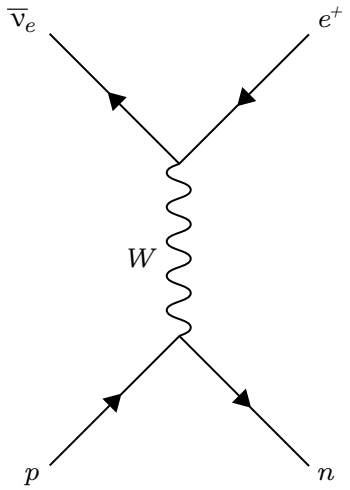
Nuclear reactors are potent sources of terrestrial neutrinos, producing large numbers of electron antineutrinos when fission fragments undergo β^- decay. Reactor neutrinos are discussed in greater detail in chapter 5.

2.4 Inverse Beta Decay

Electron antineutrinos of a few MeV total energy primarily interact with water via inverse beta decay. This reaction has the form

$$\bar{\nu}_e + p \rightarrow e^+ + n \tag{2.29}$$

with the Feynman-like diagram



Due to the difference in mass between the reactants and products, an energy threshold of 1.806 MeV is imposed for this reaction. The positron's low mass relative to the neutron results in it carrying most of the incident neutrino's energy, offset by the nucleon mass difference $\Delta = m_n - m_p = 1.293\,\text{MeV}$. Up to 15 MeV the angular distribution of the positron is nearly isotropic, and only slightly forward-going at higher energies, meaning for the purposes of most low energy analyses, the incident neutrino's direction cannot be determined from the positron. Some directionality is retained in the neutron's recoil as demonstrated by the CHOOZ reactor neutrino experiment; however, the recoil distance was measured to be only $(1.7 \pm 0.4)\,\text{cm}$ [56]. Reconstructing the incident neutrino's direction from the neutron is therefore a difficult task only possible for experiments with fine vertex reconstruction resolution and high statistics, Double Chooz, for example [57]. This method is not commonly utilised.

The cross section for IBD on free protons increases with energy. Strumia and Vissani provide a parameterisation of this cross section given by:

$$\sigma_{IBD} \approx 10^{-43}\text{cm}^2 p_e E_e E_\nu^a \tag{2.30}$$

where

$$a = -0.07056 + 0.02018 \ln E_\nu - 0.001953 \ln^3 E_\nu \tag{2.31}$$

and $E_e = E_\nu - \Delta$ [58]. The cross section below 10 MeV calculated using this parameterisation is shown in figure 2.5.

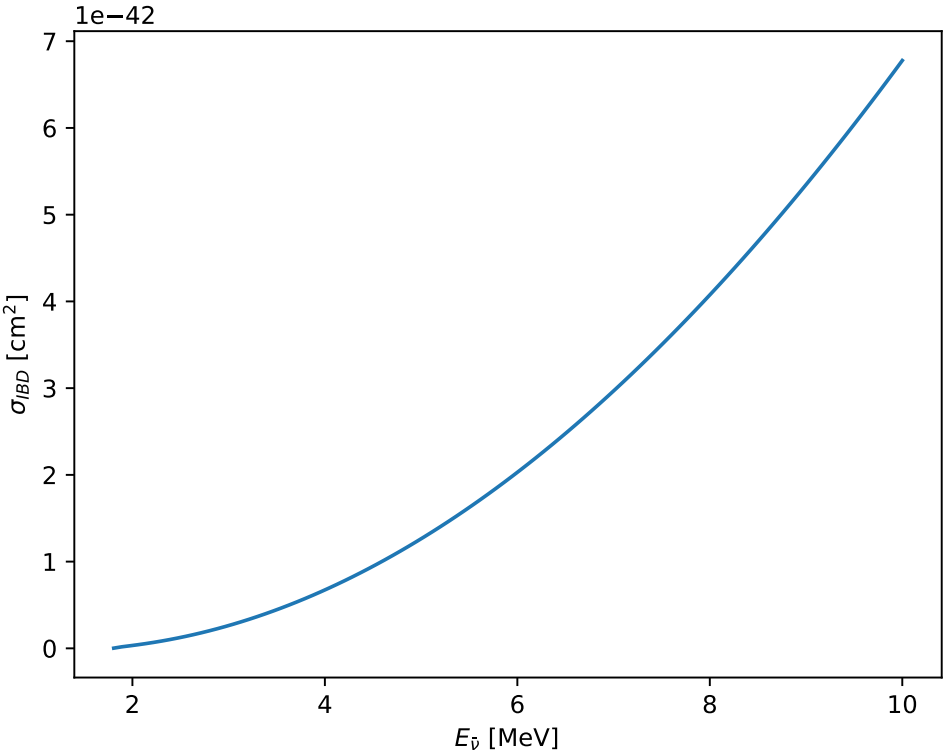


Figure 2.5: The cross section for inverse beta decay on free protons, calculated using Strumia and Vissani's parameterisation given in [58].

While the positron is of primary interest given it indicates the incident neutrino energy, the neutron can be used to tag an IBD event in cases where the positron cannot easily be identified. In the absence of a hadronic calorimeter, a neutron can be identified when it is captured by a nucleus, typically referred to as neutron capture. A short time after production, the neutron thermalises and is quickly captured by atomic nuclei. The time difference between the IBD interaction and the neutron capture varies depending on the material in which the neutron is produced; an abundance of nuclei with a relatively high neutron capture cross section will decrease the capture time. This time difference leads to the neutron capture often being referred to as the "delayed" event to the positron's "prompt".

When a neutron captures on a proton as with capture on hydrogen in water, a 2.22 MeV photon is produced due to the change in binding energy. While this is typically of lower energy than the prompt positron, the identification of two events

in temporal and spatial coincidence can help in removing background that is not typically tightly coincident. This is difficult in practice, however, if the vertex of the neutron capture cannot be well reconstructed, or background events mimic a neutron capture signal. This motivates the use of materials that produce a more distinct signal following neutron capture. One such material is gadolinium, as used in the SK-Gd upgrade detailed in chapter 4 [44].

3
Super-Kamiokande

Contents

3.1 Detector Structure 27
3.2 Detection Principle 32
3.3 Electronics and DAQ 34
3.4 Data Periods 36
3.5 BONSAI Low Energy Event Reconstruction 37
3.6 The Wide-Band Intelligent Trigger (WIT) 41
3.7 Monte-Carlo 53

The Super-Kamiokande (SK) detector is a large-scale water Cherenkov (WC) detector located in Gifu prefecture, Japan, primarily operated by the Institute of Cosmic Ray Research (ICRR) of the University of Tokyo. It is an evolution of the Kamioka Nucleon Decay Experiment (KamiokaNDE, often stylised as Kamiokande) which began data taking in 1983 and was decommissioned in 1997 following its Kamiokande-III phase. The primary aim of Kamiokande was to detect proton decay, a theorised process which could provide evidence for a Grand Unified Theory (GUT) of the weak, strong and electromagnetic forces. As neutrinos became a greater focus of the particle physics community towards the end of the 20th century, it became clear that Kamiokande also served as a particularly effective neutrino detector. Super-Kamiokande was therefore proposed: a much larger detector

of similar design with a primary aim of confirming or denying the existence of neutrino oscillation. A powerful detector in its own right, SK is also used as the far detector in the highly successful T2K long baseline neutrino oscillation experiment, detecting neutrinos produced in a beam some 295 km away in the village of Tokai, Ibaraki prefecture. It was also used similarly as part of the K2K experiment, the first long-baseline accelerator neutrino experiment, detecting neutrinos produced at KEK in Tsukuba. SK began taking data in April 1996 and will continue to operate at least until its successor, Hyper-Kamiokande (HK), is operational, which is expected to be sometime in 2027.

3.1 Detector Structure

The detector is situated in a mine operated by the Kamioka Mining Company under Mt. Ikeno which straddles the Gifu-Toyama prefectural boundary, though SK itself is entirely contained within Gifu prefecture. In any given direction, there is at least a kilometre of rock, helping to reduce the flux of atmospheric muons passing through SK to $\sim 10^{-5}\times$ the flux at the Earth's surface. A schematic of the detector's layout can be seen in figure 3.1.

SK consists of 50 kt of gadolinium-doped ultrapure water contained in a 39 m diameter stainless-steel cylindrical tank of height 42 m. Concrete fills the space between the outside of the tank and the rock, apart from above, where a large, accessible dome houses the high voltage (HV) and data acquisition (DAQ) systems for the detector. Ports for access to the detector are also found on the top of the tank. Some ports are light tight for insertion of calibration materials or sampling systems, and some act as human access ports for the rare occasions where direct access to the detector is needed. Within the tank is a 55 cm thick stainless-steel structure which supports the detector's many PMTs, as seen in figure 3.2. This structure separates the 32 kt of the so-called inner detector (ID) from the outer detector (OD) and its outer side leaves 2.2 m of room between it and the wall of the tank. Modular sub-structures are bolted together to form the support structure, with each module holding a 4×3 set of ID PMTs.

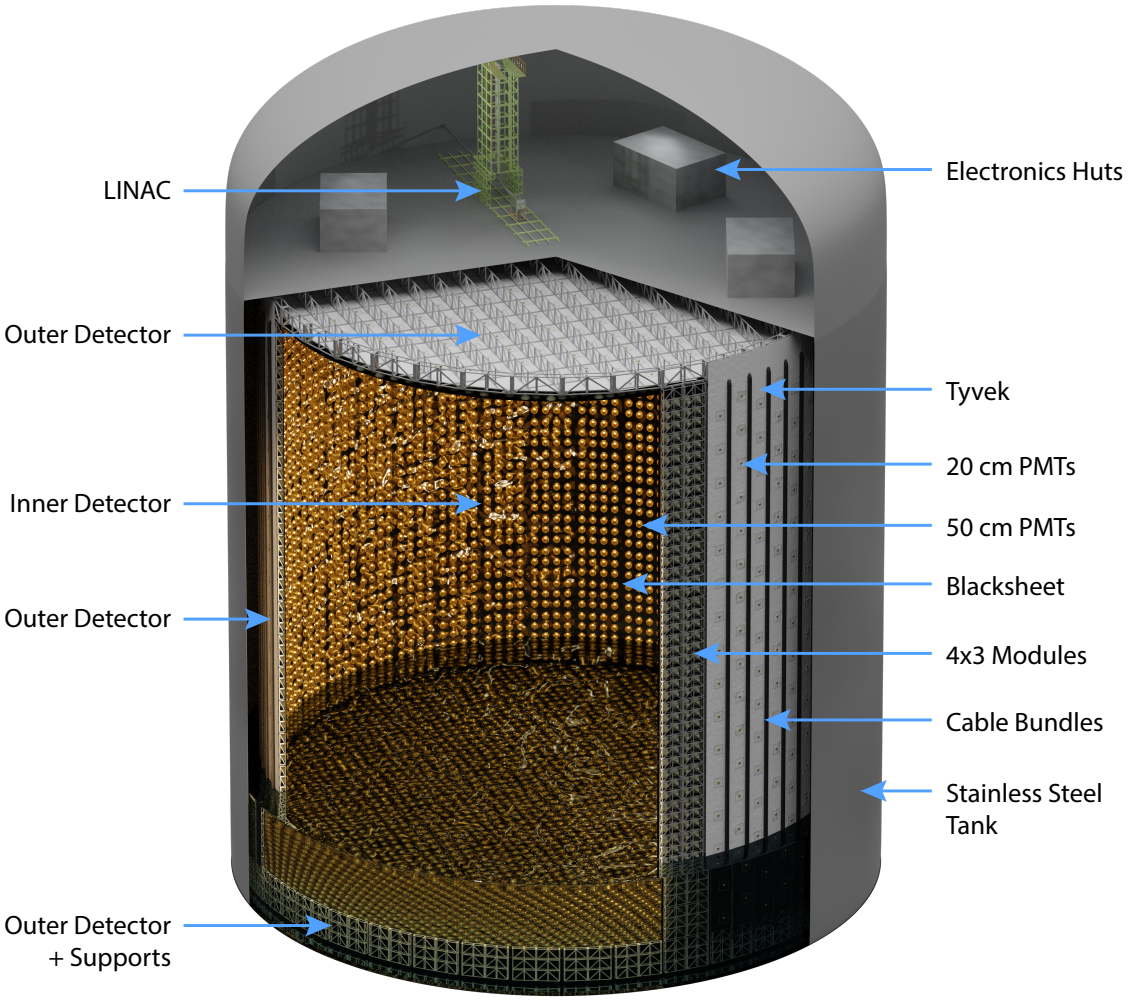


Figure 3.1: A cross section of the SK detector (partially filled with water).

The inner detector contains 11,129 Hamamatsu R3600-05 50 cm diameter PMTs on both the walls and endcaps of the ID, and are used to detect light from neutrino interactions. The quantum efficiency of these PMTs peaks at 21% at 360-400 nm, and the collection efficiency for photoelectrons at the first dynode is $> 70\%$ [59]. The ID has approximately 40% photo cathode coverage. A black plastic sheet is mounted around each of these PMTs to prevent light from crossing the ID/OD barrier. A subset of the ID volume, > 2 m from the PMT plane in most analyses (22.5 kt of water), is known as the fiducial volume (FV). Events whose interaction vertex is reconstructed outside the FV are rejected in an attempt to remove events due to radioactive decay that originate close to the PMTs or from the PMT glass itself.

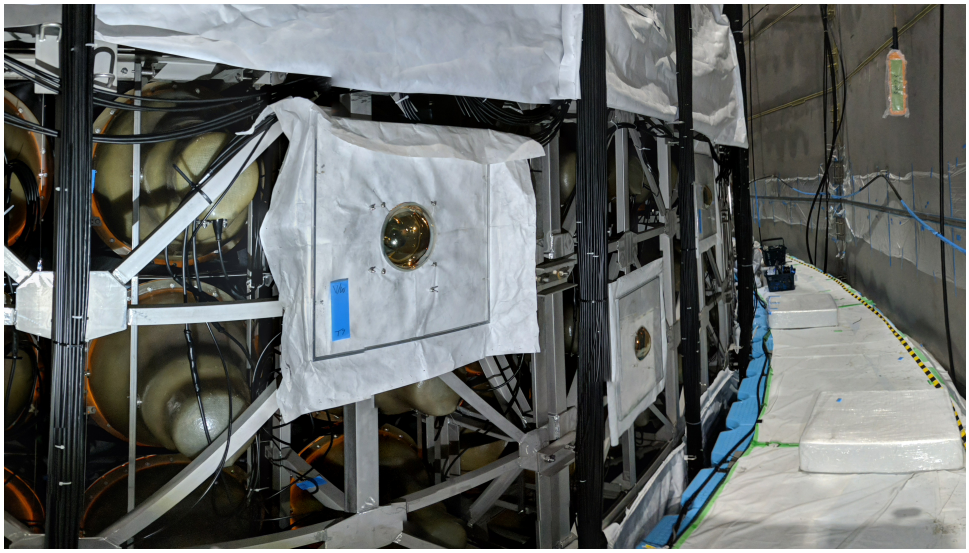
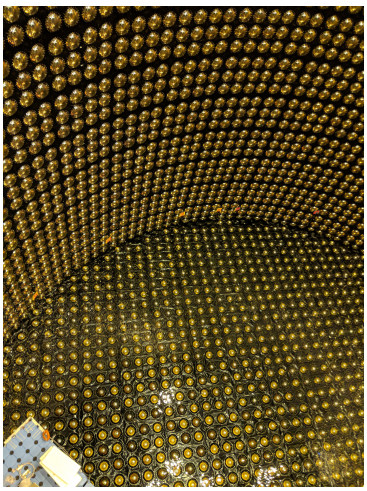
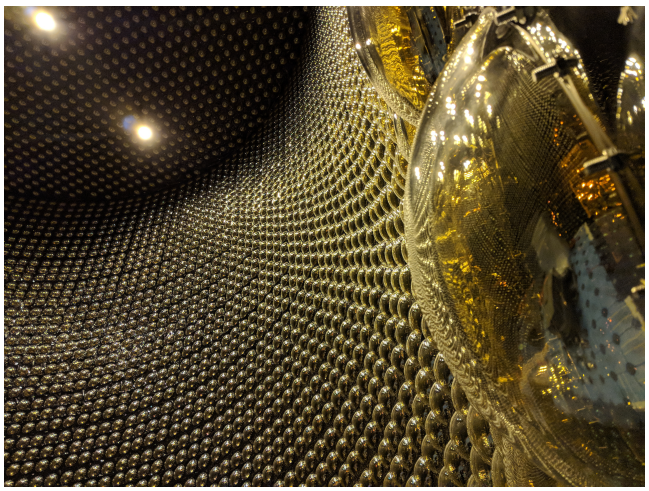


Figure 3.2: Photograph of the structure which supports the PMTs and separates the ID and OD. The fibreglass rear of the ID PMT cases can be seen within the structure. Taken during the 2018 tank open work.



(a) View from the main access hatch on top of the tank.



(b) Just above the floor of the ID, with a Hamamatsu R3600-05 PMT in the foreground.

Figure 3.3: Photographs of the SK inner detector taken during the 2018 tank open work.

Photographs of the inner detector, showing the PMT arrays, can be seen in figure 3.3.

In the OD, 1,885 outward-facing R1408 20 cm diameter PMTs are mounted evenly around the structure. Each PMT is coupled to a 50 × 50 cm acrylic wavelength shifting plate designed to act as a light guide, improving the light collection for these smaller, more sparsely arranged PMTs. Light collection is further improved by sheets of Tyvek (an opaque paper-like material manufactured by DuPont) which



Figure 3.4: A photograph from the bottom of the SK OD looking up, showing how Tyvek is used to optically isolate the OD from the ID and encourage reflection within the OD. The Hamamatsu R1408 PMTs and wavelength-shifting light collection plates can also be seen. Taken during the 2018 tank open work.

covers the outer side of the structure. This Tyvek is white on the outward-facing side, encouraging light reflection and subsequently improving the light detection efficiency of the OD PMTs. On the inward-facing side, the Tyvek is black to make the material opaque and to absorb any light which penetrates the ID black sheet. The Tyvek is hung on the structure in such a way that strives for optical isolation of the ID from the OD as can be seen in figure 3.4.

Tyvek is also hung on the inner wall of the tank to further encourage light reflection, also seen in figure 3.4. This tyvek lacks a black backing as there is no need for it to be fully opaque and must be relatively lightweight. These sheets are hung from the very top of the tank on the wall and are only connected at the very bottom, so being lightweight eases installation and prevents them from ripping under their own weight.

The purpose of the OD is to filter out events that originate outside of the inner detector; a neutrino interaction in the ID should have no signal in the OD,

as all charged particles are produced at the interaction vertex. Despite the rock overburden, cosmic muons can still reach SK, at a flux of about $6 \times 10^{-8} \mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1}$, resulting in an event rate of about 2 Hz [59]. The OD is essential in removing these events, as well as any other events where charged particles may find their way into the ID from outside the tank, such as those produced in radioactive decay in the walls surrounding the detector.

Also found in the OD are 26 sets of large "compensation coils" attached to the inner walls of the tank, both horizontally and vertically. These carry a carefully controlled current to induce a magnetic field opposite to the Earth's geomagnetic field, counteracting it to provide a neutral field in the detector volume. The geomagnetic field, while relatively weak, affects the trajectory of particles in the PMTs, reducing their collection efficiency [60]. These compensation coils reduce this effect.

The cables feeding power to the PMTs and carrying photo-current information from them are gathered in large bundles climbing up the outside of the structure. These cable bundles exit the tank through light-tight ports and are routed to one of four huts placed on top of the tank, which contain the HV systems and DAQ modules. To simplify the tuning of the voltage being sent to the PMTs, all cables are approximately the same length from the PMTs to the HV modules, requiring a large number of cable trays on the top of the tank to contain the excess length.

Just outside the dome above the tank, in one of the hallways leading away from the detector, is the LINAC. This is an ultra-low intensity, low energy ($\sim$ few MeV) electron LINAC which can be piped into one of 3 calibration ports on the top of the tank, all on the LINAC half of the hall. There are a range of different beam pipe lengths that can be inserted into the port, allowing for injection at different heights in the detector. LINAC calibration usually occurs around once a year, with extra calibration campaigns when it is deemed necessary, usually if the detector status has undergone some change while running.

In the SK coordinate system, the origin is at the centre of the detector, with positive z being vertically upwards. x increases positively from the LINAC side of the detector to the entrance of the SK hall on the opposite side of the dome.

3.2 Detection Principle

SK relies on the principle that charged particles travelling faster than the phase velocity of light in a medium emit electromagnetic radiation. This is known as Cherenkov radiation and is primarily produced in the ultraviolet spectrum, but its distribution extends into the visible.

As the light is produced by a moving source travelling faster than its fixed speed of c/n, where n is the refractive index of the medium, it takes the form of a cone. This cone's opening half-angle θ is given by:

$$\cos\theta = \frac{1}{n\beta} \tag{3.1}$$

as shown in figure 3.5.

Given that the opening angle of the Cherenkov cone is dependent on the particle's velocity, this angle can be reconstructed using the projected ring image detected by SK's PMTs and the energy of the particle can be inferred. This then gives an indication of the energy of the incident neutrino, depending on the kinematics of the interaction. Electrons and positrons below 10 MeV only travel a few cm in water before they either lose enough energy through scattering to drop below the Cherenkov threshold or are captured by an atom. Given that SK's reconstruction resolution is $\sim$1 m, such events can therefore be treated as a single interaction vertex.

The amount of Cherenkov light produced by a particle decreases as its energy approaches that of the Cherenkov threshold: the point at which a particle's velocity is equal to the local phase velocity of light. For electrons and positrons in pure water, this translates to a total energy of 0.767 MeV or 0.257 MeV kinetic. For electrons and positrons in SK with total energy less than 10 MeV above the Cherenkov threshold, there are not enough PMTs hit by light to reliably reconstruct the

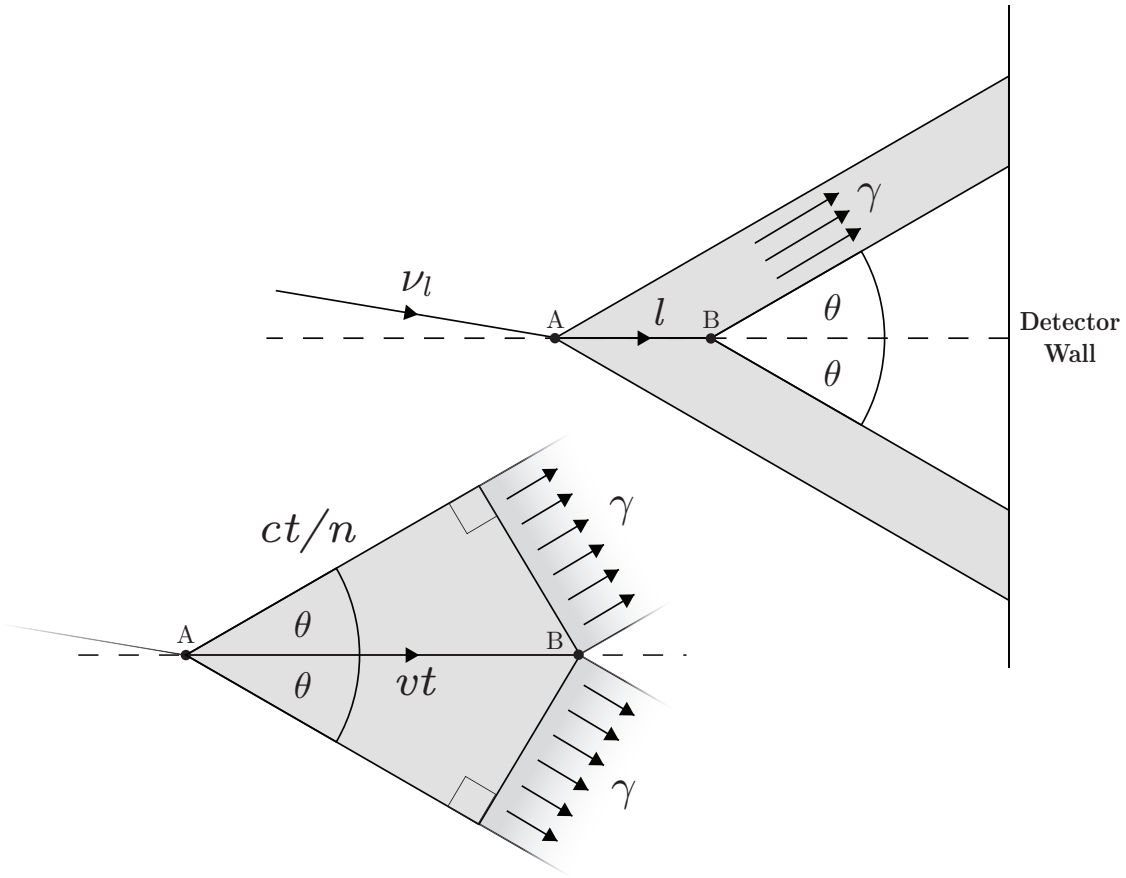


Figure 3.5: The formation of a Cherenkov ring in a medium of refractive index n by a charged particle of velocity $v > c/n$. The particle is either created or given energy by a neutrino interaction at A and stops producing Cherenkov radiation at B.

Cherenkov cone and infer the particle's energy using its angle. Instead, the number of PMTs hit by light is used to reconstruct the energy, as the energy dependence of Cherenkov emission is theoretically well known. The photo-coverage of the detector (40%) and individual PMT quantum and collection efficiencies are factored into this reconstruction. The approach to low energy reconstruction is discussed in further detail in section 3.5.

While SK's low energy analyses are limited to electron or positron detection, muon neutrino interactions are of particular interest in the energy regime above the muon's Cherenkov threshold of around 160 MeV. Indeed, this is crucial in SK's role in the T2K experiment, which produces primarily muon-type neutrinos. It is important to isolate the muon and electron-like events from one another to

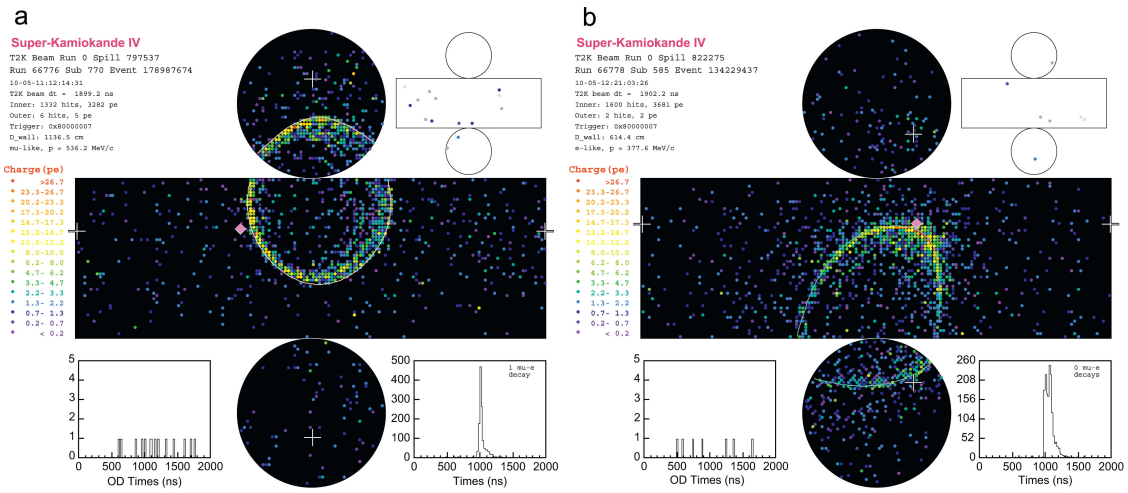


Figure 3.6: A sharp muon-like Cherenkov ring event (a) compared to a fuzzier electron-like ring (b) for two T2K events in SK [24].

determine the flavour of the incident neutrino.

To differentiate between the two lepton flavours, the sharpness of the Cherenkov ring patterns is evaluated. Given its relatively low mass, an electron will scatter multiple times before it stops producing Cherenkov radiation, its direction changing with each interaction. This effectively results in a large number of rings overlaying onto one another, each with slightly different direction, leading to a blurred ring with many PMT hits straying from the ring that corresponds to the average direction of the electron. This is in contrast to the heavier muon whose paths are typically straight with minimal scattering, producing a sharp ring. This difference in sharpness can be seen visually in figure 3.6.

3.3 Electronics and DAQ

SK's ID PMTs require a high voltage across them to work efficiently, ranging from 1.7 kV to 2 kV. This is provided by HV systems located in the electronics huts on top of the SK tank, supplied along the same cable that also carries the signal to the DAQ systems in the same huts. Each PMT must have its HV tuned uniquely, and to simplify this task all PMTs have approximately equal length cables of 70 m.

Up until August 2008, the PMT signals were processed by custom-built modules called analogue timing modules (ATMs) comprised of charge-to-analogue (QAC)

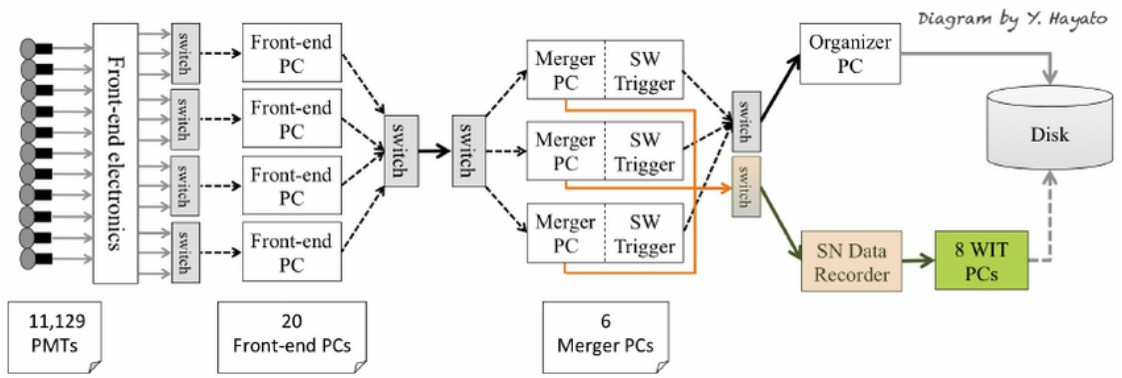


Figure 3.7: The Super-Kamiokande data acquisition (DAQ) system, in place since SK-IV except for the WIT PCs which were added in 2015.

and time-to-analogue converters (TAC). For each PMT there were two pairs of QAC/TAC that would allow for the readout of multiple sequential hits in the case where a PMT would be hit by light in quick succession. These modules were replaced with quantum tunelling composite (QTC)-based electronics with ethernet QBEE. The QBEE modules provide a significant upgrade to the charge and time resolution of the ATMs, improving to beyond the intrinsic resolution of the PMTs themselves [60]. There are 550 QBEEs, each with 24 input channels.

The hit data is stored in first-in-first-out (FIFO) memory by field programmable gate arrays (FPGAs) on the QBEE boards and passed along to the 20 front-end computers. These computers time-order the hit data as 20 ms blocks, with 1 ms overlap to cover events that would otherwise straddle two blocks.

As the vast majority of data generated by the SK detector is not due to neutrino interactions, online triggers are required to curate the data sent to long term storage. In the ATM era, this triggering was handled by hardware, only sending data to the front-end computers if the number of PMT hits in a 200 ns window exceeds a certain threshold. With the QBEEs a trigger is issued at a rate of 60 kHz, mimicking the ATM's true hardware trigger, and the front-end computers deal with triggering through software. A schematic of the SK DAQ system is given in figure 3.7 which also shows the Wide-band Intelligent Trigger (WIT) computers. These were added in 2015 and are further detailed in section 3.6.

3.4 Data Periods

To date, there are 7 distinct data-taking periods of SK, named SK-I through to SK-VII, generally separated by periods where the detector is shut down for a significant amount of time, typically for general maintenance or upgrades. The transitions between SK-V and SK-VI, and SK-VI and SK-VII had no detector downtime between them as they signalled the point where gadolinium had finished being added to the detector water, which occurred during while the detector was operational (see chapter 4). SK-VI represents the first stage of what is often collectively referred to as SK-Gd given the significant change from ultrapure to Gd-loaded water. These 7 periods are split into "runs" of data-taking which allows for easier segmentation and selection of offline data. These runs are capped at 24 h but may be shorter if a run was stopped due to some change in the detector. This most often happens when calibration data is being taken, in which case runs are marked as "test" runs which can be easily ignored during offline analyses not concerning calibration.

First data was taken in April of 1996, beginning the SK-I period, ending in July 2001 with the tank being drained for general maintenance and inspection. In November of 2001, part-way through refilling the tank to start the SK-II period, around 7,000 PMTs were suddenly destroyed. This was likely due to the implosion of a single PMT whose glass was weakened, triggering a cascade of implosions throughout the water-filled volume of the detector. Those outside of the water remained intact. The tank was again drained and the remaining functional PMTs, $\sim$47% of the total in SK-I, were redistributed evenly throughout the detector, resulting in around 19% photo-coverage. The PMTs were also fitted with a casing. This casing is made of two parts: a fibreglass section covering the opaque rear of the PMT, and a transparent acrylic dome in front. These sections are bolted together and are designed to mitigate the pressure waves created were a PMT to implode, preventing a cascade. The acrylic front does not significantly affect the photo-collection of the PMT.

SK-II began in October 2002 with reduced photo-coverage, running until October 2005. SK-III restored the original photo-coverage with the purchase of new PMTs,

running from August 2006 to August 2008. SK-IV brought in upgrades to the DAQ systems, replacing the ATM boards with new QBEEs, running from September 2008 to May of 2018, the longest data taking period for SK so far. Although much of the work on the gadolinium processing and water systems were made during SK-IV, there remained a significant amount of work to prepare the detector for Gd loading between SK-IV and SK-V. As well as the changes to the water systems inaccessible during data taking, the structure of the detector had to be meticulously cleaned in preparation for the first time SK would run without exclusively pure water in its tank.

The upgrades in preparation for SK-Gd were complete, and the tank filled with ultrapure water in January 2019, beginning the SK-V data taking period. A short shutdown in February 2020 included using remotely controlled underwater drones to make minor changes to the water inlet pipes at the base of the inner detector. Drones were employed to avoid the slow process of draining and refilling the tank for relatively minor modifications.

Following delays due to the COVID-19 pandemic, the planned date of Gd loading was pushed back from March 2020 to July 14th. Loading was completed on August 17, 2020, marking the beginning of the SK-Gd (SK-VI) phase of the experiment. Further gadolinium loading began on June 1st, 2022, acting as an interim phase before the final loading up to the ultimate target concentration. This loading was completed on July 5th, 2022 and signalled the start of the SK-VII phase which continues as of the writing of this thesis.

A summary of the phases of SK is given in table 3.1.

3.5 BONSAI Low Energy Event Reconstruction

A low energy event in SK will generally produce around 6 photoelectrons per MeV above Cherenkov threshold [61]. It takes photons up to 220 ns to traverse the detector, which defines a coincidence window in which PMT hits could have come from the same light production vertex. The dark noise rate of individual PMTs in SK (when a PMT registers a hit in the absence of light) is around 5 kHz so in this window around 12 PMTs would register dark noise hits on average. For events

Phase	Start Date	End Date	Notes
SK-I	April, 1996	July, 2001	First data taking, neutrino oscillation discovery.
SK-II	October, 2002	October, 2005	Reduced photo-coverage (PMT implosion cascade).
SK-III	August, 2006	August, 2008	Restored photo-coverage.
SK-IV	September, 2008	May, 2018	Upgraded DAQ systems, longest phase to date.
SK-V	January, 2019	August, 2020	Ready for Gd loading.
SK-VI	August, 2020	July, 2022	First phase of SK-Gd, 0.02% concentration.
SK-VII	July, 2022	In progress as of August, 2022	Increased Gd concentration of 0.06%.

Table 3.1: The data-taking phases of SK.

above 10 MeV, the number of true hits (> 60 pe) would dominate in this time frame; however, for lower energy events the dark noise hits become comparable in number to the true hits from the particle, skewing reconstruction. This lower energy regime is also home to significant radioactive background, primarily from the PMT glass and fibreglass casings. Developed in response to the severely decreased photo-coverage during the SK-II period, BONSAI (Branch Optimisation Navigating Successive Annealing Iterations) offers superior reconstruction quality of low energy events below 10 MeV (and especially below 5 MeV) over previous techniques and is crucial in reducing this radioactive background.

BONSAI defines a likelihood $\mathcal{L}$ for a given vertex based on timing residuals of all PMT hits following time of flight (TOF) subtraction. For a TOF-corrected vertex hypothesis $(\vec{x}, t_0)$ with PMT hit times t_i, each with a TOF t_i^{TOF} from the emission vertex, the likelihood is given by:

$$\mathcal{L}(\vec{x}, t_0) = \sum_{i=1}^{N_{hits}} \log(P(t_i - t_i^{TOF} - t_0)) \tag{3.2}$$

where P is the probability distribution function of the PMT photoelectron signal timing. The Cherenkov photon hits will be coincident to within a few ns given their

3 ns timing resolution. This likelihood is therefore a strong discriminator for non-coincident hits due to dark noise hits that are essentially randomly distributed within the event window. The best fit vertex would therefore have the maximum likelihood.

While the likelihood could be calculated for every possible vertex, with some resolution, this would be an incredibly computationally intensive process. Instead, BONSAI initially considers four-hit combinations of PMTs. For every combination, the hit vertices are used to solve for an emission vertex, considering TOF from emission to PMT hit. These test vertices serve as starting points for the main vertex search algorithm. In this search, BONSAI employs simulated annealing. In nature, annealing is the process of heating and cooling a material, producing larger crystals if cooled sufficiently slowly, reducing defects in the material. Simulated annealing is used in search algorithms to find a global maximum when local maxima are extreme enough for traditional algorithms to misidentify them as global. It is analogous to the natural process in that the system's "temperature" begins high, with more aggressive search patterns, and is slowly decreased to zero, at which point the algorithm completes its search and "freezes".

In BONSAI's case, the temperature of the system is the range of likelihoods that vertices must fall within to remain active, defined by a tuneable percentage of the difference between the best and worst fits. From the starting vertices, the likelihood is calculated for a set of twelve neighbouring vertices. These points are defined by the area vectors of a dodecahedron z-aligned with the reconstructed direction of the origin vertex. An optimal thirteenth point is also determined from 3D quadratic interpolation of the other twelve. All vertices that fall within the range of likelihoods of the origin remain active, and form a new "branch". This process is repeated along each branch, successively reducing the active likelihood range. If the vertex around which new vertices are calculated itself has a likelihood within the current allowed range, a new set of 12+1 vertices are evaluated in the successive iteration, this time within a reduced radius around the origin vertex. This process fine-tunes the location of vertices with a strong likelihood, resulting in a high-precision fitted vertex once the system freezes.

This process means a small fraction of the volume of the detector is explored, significantly reducing the number of computationally expensive likelihood calculations. It also results in a fine-tuned vertex not necessarily linked to the original grid of vertices tested in the case of a brute-force grid search.

In addition to the likelihood, a goodness criterion is evaluated for the final vertex $\vec{v}$, given by

$$g(\vec{v}) = \sum_{i=1}^{N_{hits}} w_i e^{-0.5((t_i - |\vec{x}_i - \vec{v}|/c)/\sigma)^2} \tag{3.3}$$

where w_i are Gaussian hit weights based on the resolution of the hit time residuals (60 ns), t_i are the PMT hit times, $\vec{x}_i$ the PMT hit locations and σ is the effective timing resolution for Cherenkov events (5 ns) [61]. This acts as a powerful test of the timing residual coincidence and is used in many analyses.

The direction is fit using a weighted average of all PMT hit vectors from the fitted vertex. Each hit is weighted by its likelihood multiplied by the Gaussian function of the hit time residual. This method also produces $\cos\theta_C$ for the Cherenkov cone half angle θ_C. While this is used in high energy analyses to determine the energy of the particle, as mentioned in section 3.2, for low energy events it is more accurate to simply count the effective number of PMT hits with some considerations. This is also known as the effective number of hits, denoted N_{eff}. The calculation of N_{eff} considers the TOF, number of active PMTs, dark noise rate and various other factors, described in [20]. A fourth-order polynomial function is used to convert N_{eff} to energy, obtained from MC simulations and tuned and validated using LINAC and deuterium-tritium (DT) calibration.

Though not strictly a part of the BONSAI reconstruction, two variables, known as OVAQ and effWall, are often used as a quality metric for events, calculated using BONSAI information. OVAQ is defined as

$$\text{OVAQ} = g^2 - \text{DirKS}^2 \tag{3.4}$$

where g is the BONSAI goodness (equation 3.3) and DirKS is a directional KS test of the PMT hit pattern. OVAQ was found to be a powerful general-purpose measure of event reconstruction quality and is used in many low energy analyses.

The second variable, effWall, is defined as the distance from the vertex to the wall of the tank, along the *negative* of the reconstructed event direction. Given that a large amount of the low energy backgrounds in SK originate in the walls of the ID, especially from the PMT glass, cuts on effWall can significantly reduce background rates. The distance along the positive event direction (toWall) as well as the closest distance from the vertex to the wall (fromWall) are also used in SK analyses, though less so than effWall.

3.6 The Wide-Band Intelligent Trigger (WIT)

Prior to 2015, the main triggering in SK was handled by hit-counters. The number of PMT hits in a 200 ns window is checked against a number of thresholds, issuing a trigger corresponding to the highest threshold it passes. There is also an OD trigger, flagging an event if the number of OD PMT hits exceeds a threshold, used to remove events entering from outside the ID. There are four levels of trigger thresholds, detailed in table 3.2. Events passing these thresholds trigger the DAQ to save a 40 μs time window of hits centred on the 200 ns trigger window, or 1.3 μs in the case of the super low energy (SLE) trigger, given the high data rate in this regime. In the case of a super high energy (SHE) trigger with no OD trigger, an after (AFT) trigger is also issued. This saves all hits in a 500 μs window after the initial trigger window, with the aim to save neutron captures and other secondary particles such as those produced in spallation.

The abundance of background in the low energy regime is reflected in the high SLE trigger rate. 34 PMT hits corresponds approximately to an energy threshold of 3.7 MeV, assuming 12 hits due to dark noise. The SLE trigger is effectively the lowest a simple hit-counter threshold can be set without overloading SK's DAQ systems and quickly filling up long term storage.

Name	Abbreviation	NHits Threshold	Rate
Super Low Energy	SLE	34	≈15 kHz
Low Energy	LE	47	∼40 Hz
High Energy	HE	50	∼10 Hz
Super High Energy	SHE	52	≈60 Hz
Outer Detector	OD	22	Unknown

Table 3.2: Hit-counter trigger thresholds and their approximate trigger rates in SK as of September 2021. The range of hits for the triggers do not overlap; for example, the SLE trigger only activates for events with 34-46 hits in the trigger window. The OD trigger is not independently monitored and is only considered when an ID trigger is activated.

To lower the energy threshold further, the nature of the background contributing to the trigger rate must be considered. At low energies, the background is overwhelmingly due to dark noise, as well as radioactive decay in the glass and fibreglass casings of the PMTs. The radioactive decay events typically reconstruct close to the PMT plane as a result, which motivates the SK fiducial volume boundary being set 2 m from the PMTs. If the vertex of an event can be quickly reconstructed online, a fiducial volume requirement can be applied to events before being sent to long term storage. The quality of the event and its reconstruction can also factor into this decision.

The Wide-band Intelligent Trigger was introduced to the SK triggering systems in 2015, aiming to take a more nuanced approach to online triggering in support of pushing the energy threshold of SK as low as possible [62]. WIT reconstructs events in stages of increasing granularity and by extension processing time. This gives the ability to reject events at each stage based on their reconstruction vertex and quality, minimising processing power being wasted on events that are of low quality or that were likely produced outside of the fiducial volume.

Despite its intelligent nature, the vast majority of the data that WIT collects is dark noise and radioactive decay events, simply as a symptom of operating so close to the Cherenkov threshold. Regardless, it is an extremely powerful system as it independently triggers on these low energy events without being overwhelmed, in contrast to the hit-counter triggers which must rely on after-triggers to higher energy events, with reduced efficiency. There is also an advantage to the large amount of

background in WIT: unprocessed WIT data can be collectively treated as background for the purposes of optimisation of sample selections and tracking background rates. For a given window of data in WIT, the total event rate is many orders of magnitude higher than the rate of neutrino interactions in WIT's energy range, providing a large collective background sample. This is described in more detail in chapter 6.

There are 6 key stages to WIT:

1. Pre-Trigger

   Akin to the existing hit counters, a pre-trigger rejects events with less than a certain number of hits in a 230 ns window, corresponding to the longest time for light to traverse the SK tank with an extra 10 ns to give some leeway. This number is not fixed, but is a certain threshold of hits above a dynamically calculated dark noise rate. The dark noise rate is typically around 12 in SK-Gd, and the threshold above dark noise is currently set to 12. In SK-IV the threshold was set to 11; however, an increase in trigger rate following the 2018 tank open motivated a higher threshold. Although this trigger rate began to decline, it again increased following Gd loading, and so the threshold will likely remain at 12 for the foreseeable future.

2. Simultaneous Trigger and Online Reconstruction (STORE) Goodness

   The first stage of BONSAI is then used to generate a set of test vertices. These are calculated using combinations of four PMT hits, detailed in section 3.5. For each vertex, the TOF-corrected coincidence of PMT hits in a 5 ns window is calculated, given by

$$\text{STORE goodness} = \sum_{i}^{N_{hits}} \exp 0.5 \left(\frac{\Delta t_i}{5\,\text{ns}} \right)^2 \tag{3.5}$$

   where Δt_i is the TOF-corrected hit time for the ith PMT. Δt_i is calculated simply by

$$\Delta t_i = t_i - \frac{|v_e - v_i|}{c_{group}} - t_e \tag{3.6}$$

for uncorrected hit time t_i, hit vertex v_i and emission time and vertex t_e and v_e. Given the powerful timing resolution[1] of the SK PMTs of $\sim$3 ns [63], this serves as a strong discriminator to identify events with a large number of coincident hits, indicating a true event.

The highest STORE goodness for all test vertices is selected. If this goodness is below 0.6 $\times$ the pre-trigger threshold, the event is discarded. Additionally, if the goodness exceeds 1 $\times$ the pre-trigger threshold the event is deemed to have particularly coincident hits, relaxing some later trigger requirements. Figure 3.8 shows how the distribution of STORE goodness for background events is shifted downwards when compared to reactor positrons and Gd captures, with a sharper fall-off at higher values. Hydrogen capture events also have a higher distribution than the background; however, far less so than Gd capture.

[1]Compared to the vast majority of PMTs in use worldwide, the timing resolution of the PMT design used in SK is relatively poor. This is due to the large size of the photocathode surface necessitating a longer travel time for the photoelectrons from the point of emission to the dynode. This travel time is significant; however, when considering the travel time of light throughout the detector, it is comparatively small.

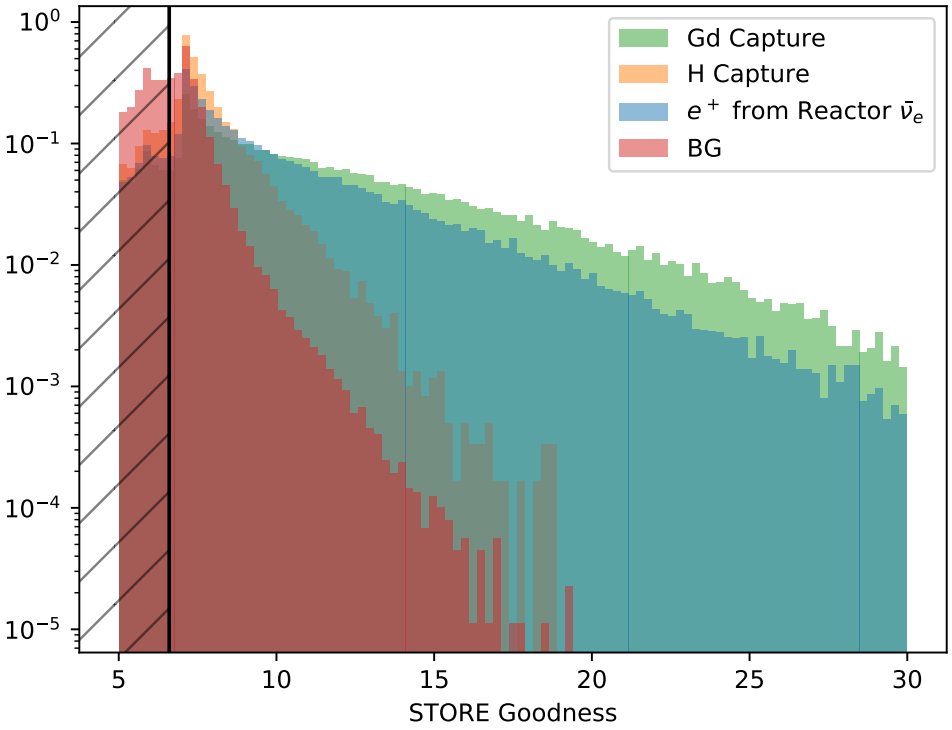


Figure 3.8: STORE Goodness for simulated events and raw WIT data (BG). Events below the black line (in the hatched region) are rejected in WIT.

3. Clus-Fit Fiducial Volume

   A vertex is reconstructed using the Clus-Fit fast fitter algorithm. This fitter defines clusters of PMTs which may contribute to the event, constrained by the 5 ns window of STORE and merges the largest clusters to form a set of PMTs from which the vertex is reconstructed in a process similar to BONSAI. An initial fiducial cut is then applied to this event, unless STORE flagged the event as having highly coincident hits. Given the relative low fidelity of the Clus-Fit vertex, events must be at least 1 m from the PMT wall, in contrast to the typical 2 m cut in most SK analyses. It is possible an event may fail to reconstruct using Clus-Fit due to a lack of good PMT hit clusters. In this case, if STORE flagged the event as having highly coincident hits this step is skipped, otherwise the event is discarded. Figures 3.9 and 3.10 show the large

number of events that reconstruct through Clus-Fit close to the PMT wall, removed with this cut.

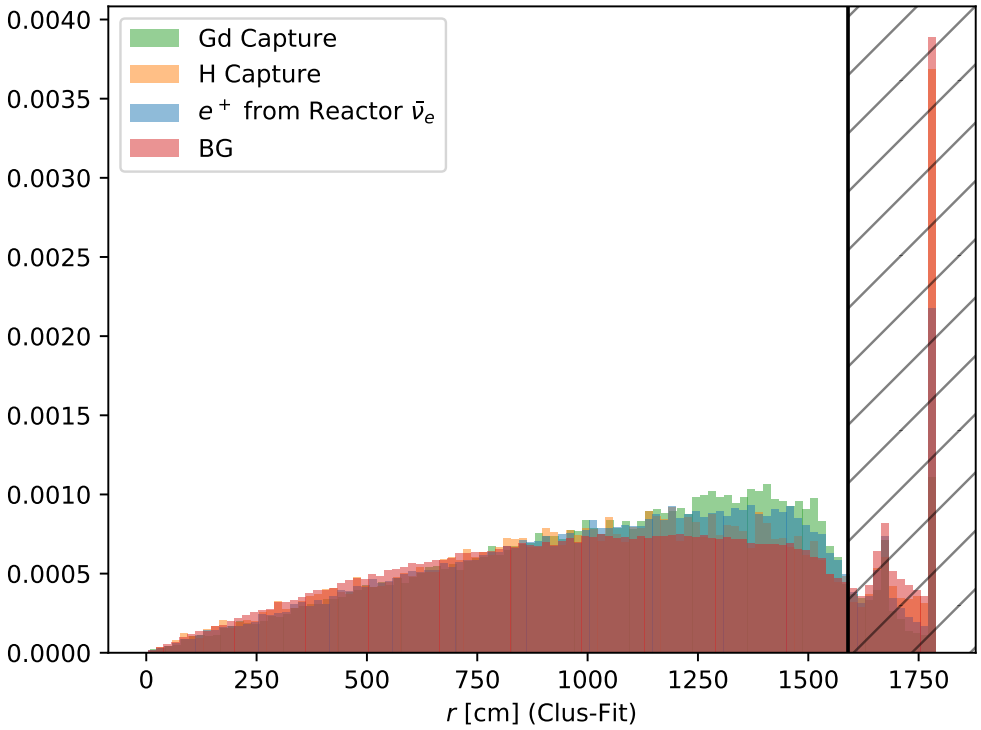


Figure 3.9: Radial position for simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by Clus-Fit. Events above the black line (in the hatched region) are rejected in WIT.

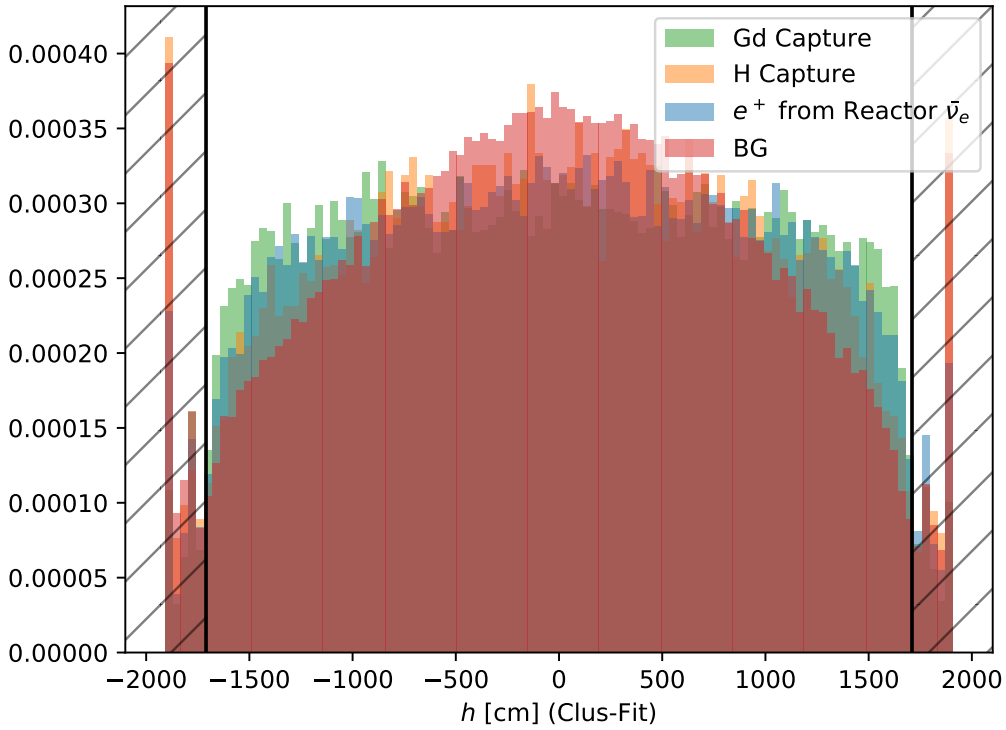


Figure 3.10: Height of simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by Clus-Fit. Events outside of the two black lines (in the hatched region) are rejected in WIT.

4. BONSAI Fiducial Volume

A lightweight version of the BONSAI fit with faster annealing is then performed, using the test vertices generated as part of the STORE step as the starting points. This fit considers all PMT hits from 0.4 μs before to 1 μs after the Clus-Fit fitted vertex. Given the higher quality of fit in this step, the full fiducial cut is applied, removing events which reconstruct within 2 m of the PMT plane. Figures 3.11 and 3.12 show how a significant number of events which pass the Clus-Fit fiducial volume cut reconstruct within 1 m of the PMT wall with the more accurate BONSAI reconstruction.

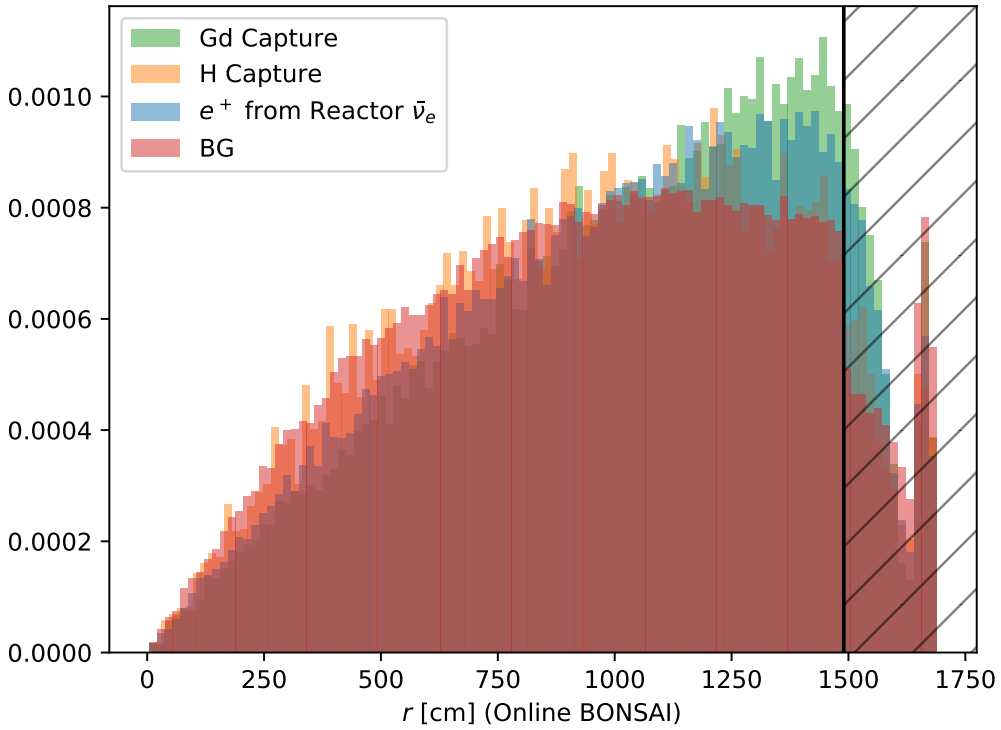


Figure 3.11: Radial position for simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by BONSAI. The disconnect in the BG distribution at the trigger threshold is a symptom of the WIT process applying FV cuts to some events, and is not physical. Events above the black line (in the hatched region) are rejected in WIT.

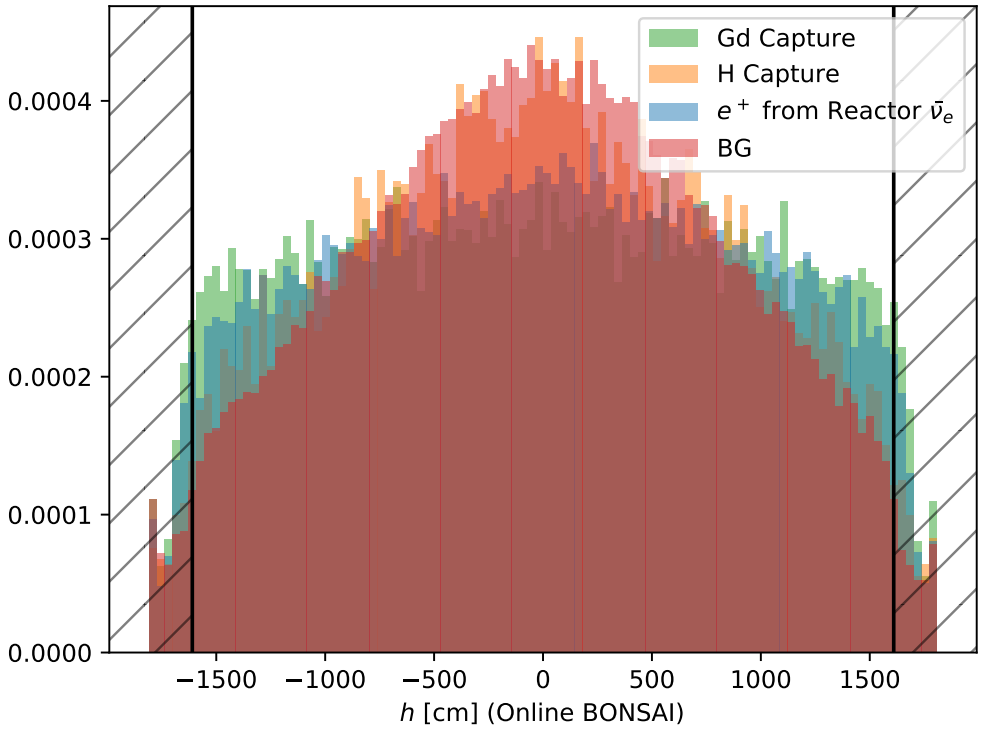


Figure 3.12: Height of simulated events and raw WIT data (BG) that pass the preceding stages of WIT, reconstructed by BONSAI. The disconnect in the BG distribution at the trigger threshold is a symptom of the WIT process applying FV cuts to some events, and is not physical. Events outside of the two black lines (in the hatched region) are rejected in WIT.

5. BONSAI Goodness

In response to the high trigger rate of SK-V an additional cut was introduced which removes events with a BONSAI goodness below some threshold. This is an evaluation of the timing residuals of PMT hits following TOF-correction, discussed in more detail in section 3.5. This threshold was set to 0.3 in early SK-V, shown in Figure 3.13 to be a relatively loose cut, removing the tail of the BG distribution.

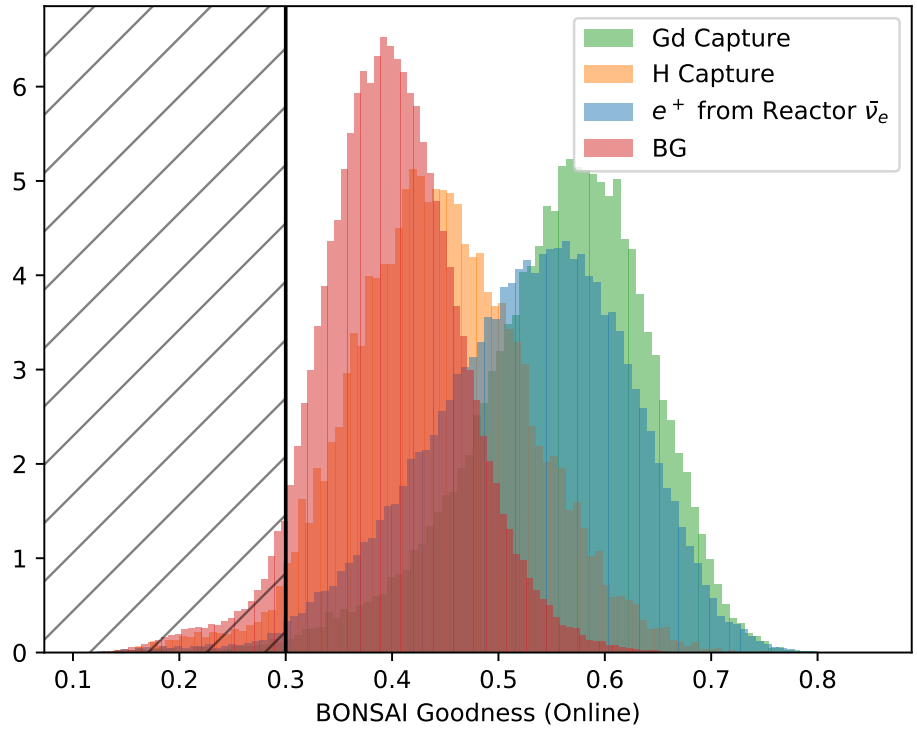


Figure 3.13: BONSAI goodness of simulated events and raw WIT data (BG) that pass the preceding stages of WIT. Events below the black line (in the hatched region) are rejected by WIT.

6. N18

   Finally, the number of hits in an 18 ns window around the BONSAI vertex is counted, acting as a final hit coincidence metric. Despite its simplicity, N18 is a powerful discriminator, with the majority of BG events having a particularly low N18, as seen in figure 3.14. The N18 threshold is currently set at 9.

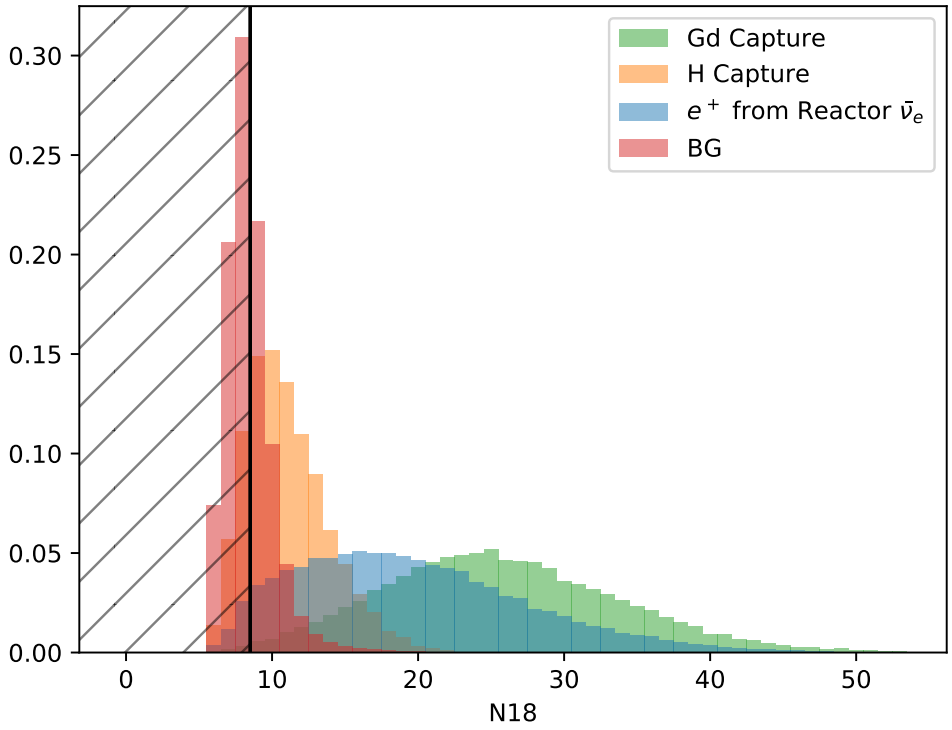


Figure 3.14: N18 of simulated events and raw WIT data (BG) that pass the preceding stages of WIT. Events below the black line (in the hatched region) are rejected by WIT.

Given its computationally intensive nature compared to the other SK triggers, WIT requires dedicated hardware, independent of the primary SK DAQ dataflow. WIT consists of ten computers. Eight are from the original WIT installation, each consisting of two 2.7 GHz Intel Xeon E5-2680 CPUs (8 physical cores, 16 hyper-threaded). Seven of these machines have 62 GB of DDRIII RAM and are tasked with running the main WIT execution; they are worker machines processing the reconstruction and triggering described above. The eighth machine has 256 GB of RAM. This is the organiser machine, tasked with receiving the data from the other WIT machines and organising the events in time order. Two additional worker machines were installed after WIT's initial deployment, each with two 2.6 GHz Intel Xeon Gold 6132 CPUs (14 physical cores, 28 hyper-threaded) with 125 GB of DDRIII RAM. New WIT machines are expected to be installed sometime in 2022.

The organiser is the main bottleneck of the WIT data flow. Data does not arrive at the WIT machines in time order, instead arriving in data blocks as described in section 3.3. While the events within a given block will be time-ordered, the blocks themselves will be processed across the main WIT worker machines asynchronously. This is the reason for the larger memory allocation to the organiser machine. All data must be collected by this machine, organised, then transferred to long term storage. If the organisation takes too long, the machine's memory will become full and the WIT worker machines must hold onto their data until space is available. If they themselves run out of memory, the main DAQ is instructed to stop sending data to WIT until the system can recover. In rare cases, the automatic processes dictating the flow of data may not catch such a backlog, resulting in memory overflow errors and the crashing of the WIT machines.

To prevent such issues, the various trigger values are tuned to produce a data rate a little below the maximum the organiser can handle. This means in the case of an increased event rate, WIT will still be able to continue as normal or, in extreme cases, temporarily keep data on the worker machines until the rate decreases and the organiser recovers.

The WIT machines hook into the SN watch system, which would read out all events from SK in the case of a supernova. Data is directly handed to SN watch, with no triggering, in a data stream shared with WIT. This gives WIT access to SK's raw data before any kind of reduction is applied. The drawback of this approach is that, in the case of a supernova, data is no longer sent to WIT as priority is given to the SN system. There are efforts, however, to instead maintain the data flow to WIT during an SN neutrino burst in an effort to improve the efficiency of detecting neutron captures and other low energy events during the burst.

Once data has been saved by WIT and transferred to the SK computing cluster, the complete BONSAI reconstruction is run on each event, providing more accurately reconstructed event information with which to apply a first reduction to the data. Given the large number of events passing through WIT, even with the online cuts, the first reduction process is crucial to retain a manageable amount of data. In

the typical WIT first reduction, some minor, energy-dependent cuts are applied, namely on OVAQ and effWall which removes the majority of all events from the data. Since March 2020, WIT can identify pairs of temporally and spatially coincident events in the online data (discussed in section 4.5) and apply relaxed trigger requirements in these cases. Such events have specific trigger ID numbers and so in parallel to the main WIT first reduction, coincident pairs are fully reconstructed and saved regardless of their event quality. These do not make up a significant portion of the raw data, and so it is reasonable from a storage space perspective to retain all pairs of events.

3.7 Monte-Carlo

The vast majority of analyses in SK take a Monte Carlo (MC) approach to predicting detector response. This process involves simulating the intricacies of particle interactions in the materials throughout SK as well as their Cherenkov light production. The PMT responses and their dark noise signals must also be simulated.

SK's standard MC production is primarily handled by the SKDetSim detector simulation software. This is built on the GEANT-3 particle simulation platform [64], an extensive suite of FORTRAN-based MC libraries. The long operational history of SK has allowed for a large amount of development and validation of SKDetSim, crucial to understanding the detector's response in support of a wide range of analyses. GEANT-3 has been superseded by Geant4 (now stylised without all-caps), primarily built in C++ [65]. A new version of SKDetSim called SKG4 which utilises this is in limited use, motivated by GEANT-3's inability to accurately simulate neutron capture on gadolinium [66], a requirement of the SK-Gd phase of the experiment (chapter 4). Analyses not utilising SKG4 typically model neutron capture externally, and the produced gamma cascade vertex information is used as input to SKDetSim. SKG4 is still in its early adoption phase, and so all MC production in this work was handled by SKDetSim.

The dark noise rate of the PMTs in SK is frequently measured and SKDetSim is updated to account for this; however, correctly modelling the pattern of dark

noise "hits" throughout the detector is of particular importance at energies where these dominate over true hits of Cherenkov light. To minimise inaccuracies in the reconstruction quality of events due to this effect, MC production for WIT studies involves injecting simulated events into real data. The events are simulated using SKDetSim, but with the dark noise rate of the PMTs set to zero. This results in events with exclusively true hits. A modified version of the standard WIT data processor takes these events and overlays them onto raw data taken after the DAQ but before WIT processes the file. The program then processes the file essentially as if it were a real data file, which means that MC events will not be saved naturally if they do not pass WIT's requirements. To track the MC which was injected into the file, an additional process forces the reconstruction of all of the injected MC events, no matter the event quality, using the true timing as the centre of the event window.

This often results in MC events being duplicated in the final file: one version which was forced to save with the event window centred on the true event time, and another which was saved naturally if WIT deemed it a good event based on its trigger values. The former version of the event should be avoided in most cases, as it will generally have superior reconstruction quality to the latter given that the starting point of the reconstruction is at the correct time vertex. These events are trivial to remove as they are saved with a unique trigger identification number which indicates that they are MC events forced to save.

The issue comes when identifying the MC events which saved naturally, as there is no clear identifier separating them from events which WIT naturally identified in the data file. Such events cannot necessarily be tagged as MC on reconstruction, as the concept of what deems whether it truly is an MC event is somewhat abstract.

The data file used for MC injection will contain numerous events that WIT deems an eligible event. When injecting a large number of MC events into a data file, there is a chance that such an event would overlap with an injected MC event. In these cases, the MC event itself may not pass the WIT thresholds; however, it may contribute a small number of hits to the event from the file which may or may not factor into the reconstruction. The simple presence of MC hits in an event is

therefore not a good indicator of an MC event that naturally saved. If instead the MC would've produced a good event in isolation, the overlap with another event can skew the reconstruction to produce a vertex not indicative of the injected event.

Practically, the vast majority of events do not experience an overlap with events in the real data; however, such events must be avoided so as to not skew the sample selections, especially as events which overlap will inherently produce more hits and will likely form a higher quality event. To identify the MC events in the file, the reconstructed time of events is compared to that of the nearest simulated event time, stored in the file for all events produced in SKDetSim. Events from the data file will be randomly distributed in time, so only a small number will fall close to the true emission vertex. The reconstructed vertex of MC events, however, will be temporally close in the case of at least reasonably accurate reconstruction. Events greater than 150 ns from the MC vertex are therefore rejected, resulting in an MC sample with $> 99\%$ purity.

4
A Gadolinium Loaded Super-K

Contents

4.1 Motivation **56**
4.2 Water Systems **57**
4.3 EGADS **59**
4.4 Deployment and Validation **61**
4.5 WIT Coincidence Trigger **63**

4.1 Motivation

As discussed in section 2.4, inverse beta decay is the dominant interaction channel for electron antineutrinos of a few MeV, producing a positron and a neutron. These neutrinos originate from numerous neutrino sources of interest to Super-Kamiokande. Given the high amount of background in this energy regime, IBD searches in SK typically require the detection of a delayed neutron capture in coincidence with a prompt positron event. In SK's pure water phase, the neutron would capture on a hydrogen atom after $\sim$200 μs, producing a 2.2 MeV photon. Such a photon goes on to scatter atomic electrons that in turn produce Cherenkov light. Given the low energy, very few PMTs detect light in these events; on average, only 7 will register a hit. Neutron capture events are therefore usually lost to the main SK

trigger system, and so either the WIT or 500 μs after trigger systems (see section 3.6) are employed in neutron capture searches. According to Monte Carlo studies, WIT is approximately 17% efficient at detecting fiducial 2.2 MeV photons.

To greatly improve the detectability of neutron captures in SK, in 2003 Beacom and Vagins proposed GADZOOKS! (Gadolinium Antineutrino Detector Zealously Outperforming Old Kamiokande, Super!), an upgrade to Super-Kamiokande through the doping of its water with 0.2% gadolinium trichloride by mass [44]. Natural gadolinium's neutron capture cross section is 48,800 b, $\sim 100,000\times$ larger than that of free protons. At the proposed concentration, this results in a mean capture time of $\sim$30 μs, 10$\times$ shorter than in pure water. Additionally, when a neutron captures on gadolinium it releases an 8 MeV gamma cascade. These cascades result in around 23 PMT hits, a far more distinct signal in comparison to hydrogen capture and one which WIT is approximately 100% efficient at triggering on.

A long process of research and development supported this project, including the move away from gadolinium chloride hexahydrate to using gadolinium sulphate octahydrate[1] $Gd_2(SO_4)_3 \cdot 8\,H_2O$, thanks to its lower reactivity while retaining easy dissolution and Cherenkov light transparency. GADZOOKS! was realised as the SK-Gd phase with two planned stages: a first deployment 0.02% gadolinium sulphate by mass, ensuring approximately half of all neutrons are captured by gadolinium, before later increasing to the GADZOOKS! proposed loading of 0.2% with a $\sim$90% Gd neutron capture rate. The first phase successfully completed its deployment in July of 2020 [1] and loading for an interim phase where the concentration was increased to 0.06% finished in July of 2022.

4.2 Water Systems

In its pure water phase, SK's water systems strived to remove any and all impurities from the water, resulting in exceptional water purity. If the water were doped with gadolinium sulphate in this configuration, it would quickly be removed by these

[1] *Gadolinium sulphate* or more simply *gadolinium* is often used to refer to the full molecular form as a short hand.

systems which continuously circulate and purify the water. A challenge is therefore presented: retain the level of water purity and hence transparency of SK's pure water phases without removing the gadolinium sulphate. This required development of new, bespoke water systems specific to the SK-Gd phase, first to dissolve and pre-treat the $Gd_2(SO_4)_3 \cdot 8\,H_2O$ and then to continuously circulate and purify the gadolinium-loaded water. A schematic of these systems is shown in figure 4.1.

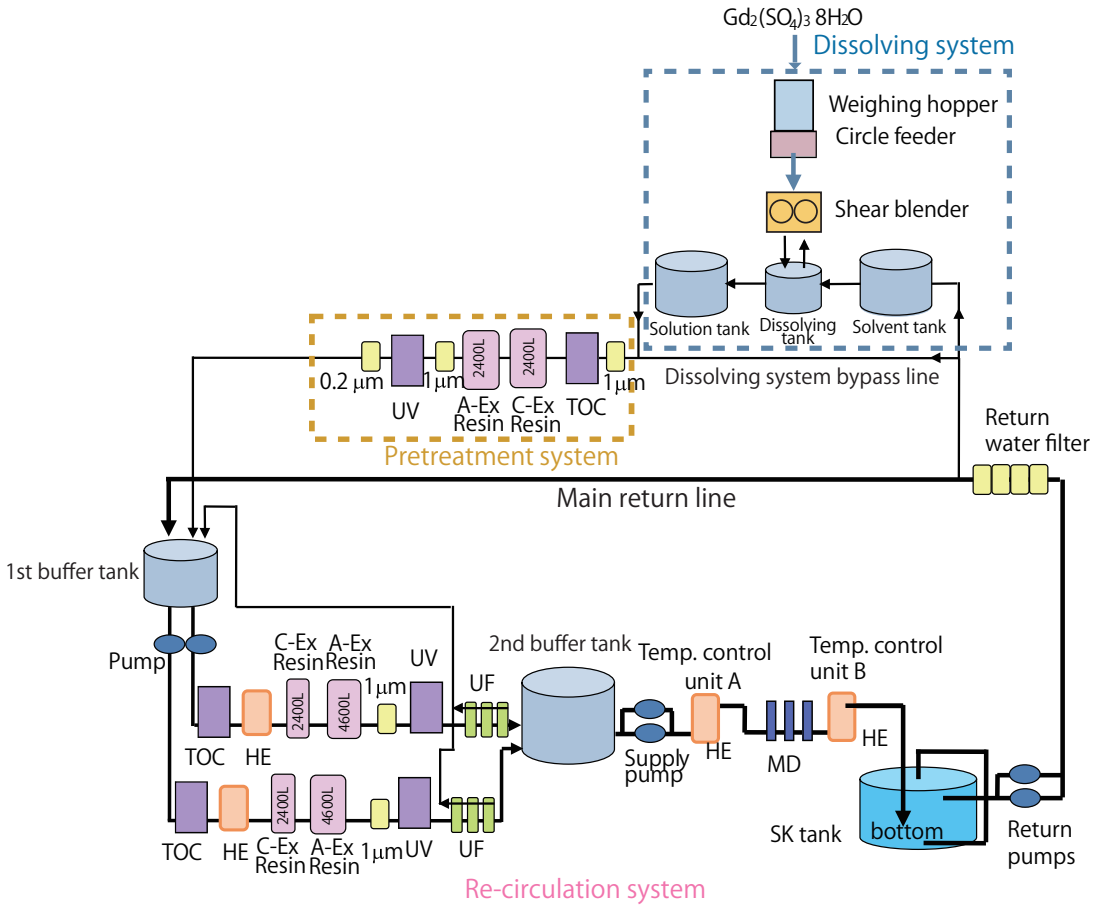


Figure 4.1: Schematic of the SK-Gd water systems [1].

First, a measured amount of $Gd_2(SO_4)_3 \cdot 8\,H_2O$ powder is fed into a buffer tank via a circle feeder where it is dissolved in water taken directly from the SK water's recirculation feed. The solution is then transferred to the pretreatment system, responsible for removing all substances from the water other than gadolinium and sulphate ions. Pretreatment begins with a passive 1 μm filter before irradiation with ultraviolet light designed to remove any organic impurities by oxidising carbon

and other compounds. The light is delivered using a total organic carbon reduction (TOC) lamp. To remove inorganic, ionised impurities, the specially prepared AMBERJET™1020 [67] and AMBERJET™4400 [68] ion exchange resins remove positive and negatively charged impurities respectively. These resins are specifically designed to retain the gadolinium sulphate. The solution is then sent through a 1 μm filter, further UV sterilisation and finally a 0.2 μm filter, all designed to remove bacterial contamination.

The pretreatment system is only used during initial Gd loading; however, the recirculation system which continuously cycles and treats the water largely mirrors pretreatment with some modifications. The final, 0.2 μm filter is replaced by a set of ultra filtration (UF) modules and a heat exchanger is included to carefully control the temperature of the water. The water is kept at 12°, primarily to limit bacterial growth that may reduce the water purity and dirty the detector components. Additionally, a membrane degasifier is included, primarily to remove dissolved radon gas from the water. The recirculation system is duplicated, with two lines running in parallel supporting a flow rate of 60 $\mathrm{m^3h^{-1}}$ through each line. This allows for maintenance on one line while the system continues to treat the water, albeit at a lower flow rate. Any water which is rejected by the recirculation system is returned to a buffer tank upstream so it can be reprocessed.

4.3 EGADS

Before full-scale deployment for SK-Gd, the unique water treatment systems were developed and tested first in small scale at the University of California, Irvine (UCI), before scaling up to a test bed known as EGADS (Evaluating Gadolinium's Action on Detector Systems) [69]. Shown in figure 4.2, EGADS is a 200 t detector utilising the same components and construction techniques as SK, serving primarily as a test platform for SK-Gd, based in the Kamioka mine. Though lacking an OD, it otherwise acts as a scale version of SK, instrumented with 240 PMTs, 224 of which are the 50 cm Hamamatsu R3600-05 design used in SK's ID. The remaining 16 are of various designs that will potentially be used in Hyper-Kamiokande.

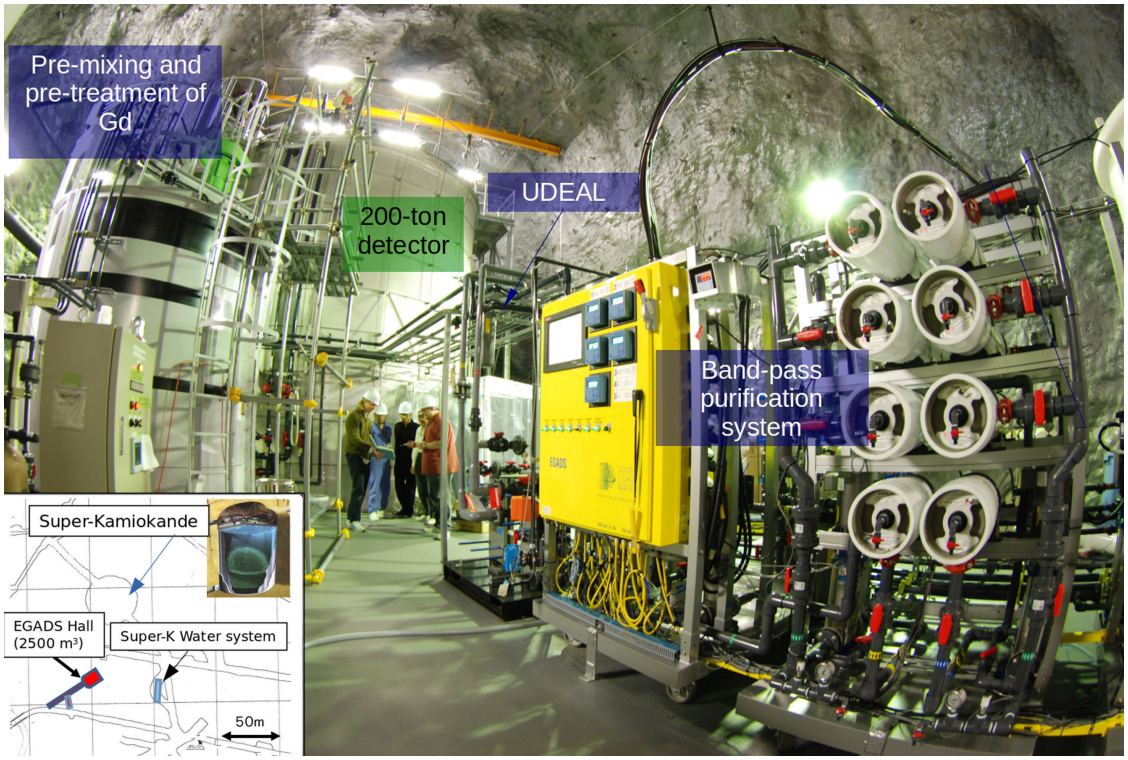


Figure 4.2: The EGADS hall.

As well as serving as a proof of the water systems' capabilities, EGADS also addressed two main challenges regarding Gd loading.

First, SK's size introduces particular importance in retaining the Cherenkov light transparency; a significant loss of light attenuation length would affect all analyses utilising SK. To monitor this, the UDEAL (Underground Device Evaluating Attenuation Length) system was deployed to measure the light attenuation length in water from EGADS for a range of wavelengths. UDEAL showed that, outside of short reductions when first loading gadolinium, the light attenuation length remained within the range of SK-III and SK-IV values.

Another challenge with Gd loading arises with a property of gadolinium in that it promotes rusting in untreated metals. SK's tank and structure are both constructed of stainless steel which should resist this; however, understanding the effect of gadolinium-doped water on its numerous other submerged components as well as long term exposure to stainless steel is crucial. EGADS sought out to be as close an analogue of SK as possible, even down to its stainless steel bolts.

Gadolinium was first loaded into EGADS in 2014 and after two and a half years of running in SK-Gd conditions with full Gd loading, the tank was emptied and the inner structure was inspected. No deterioration was found.

4.4 Deployment and Validation

To house the extensive SK-Gd water systems, a new cavern named Hall G was excavated in the Kamioka mine in close proximity to the SK experimental hall. The existing water systems were also modified to accommodate the new water treatment process.

In 2018, the SK tank was opened for maintenance for the first time in 12 years. This maintenance largely consisted of: cleaning the structure and tank of SK of any rust and debris; replacing dead PMTs in both the ID and OD; replacing the Tyvek sheets in the OD; fixing any leaks from the tank; and installing upgrades to the water system. SK is a powerful supernova neutrino burst detector and as such any downtime such as this maintenance is minimised to maximise the chance of catching a galactic supernova. To ensure a speedy maintenance period, the cleaning of the structure and Tyvek and PMT replacement was largely a joint effort of SK and T2K collaborators. Maintenance of the tank and support structure was primarily carried out by employees of the Mitsui Engineering and Shipbuilding company. It is crucial that, when gadolinium-loaded, SK's water must not leak from the tank into the surrounding environment. Following the 2018 maintenance, measurements indicated no detectable leak, with a limit of $< 15\,\mathrm{kgd}^{-1}$.

After refilling, a minor modification of the ID inlet pipes was required to prevent a convection pattern forming, disturbing the water in the tank. These pipes are found at the bottom of the ID and required the fitment of a diffuser. In lieu of the long process of draining the majority of the tank to fit the diffusers by hand, a remotely operated underwater drone was employed, provided by the Mitsui company. The drone had a diffuser attached to its underside and was manoeuvred above the inlet pipes and gently lowered. Once the diffuser was seated, the drone would separate from the diffuser. Now lighter, the drone would float to the surface of

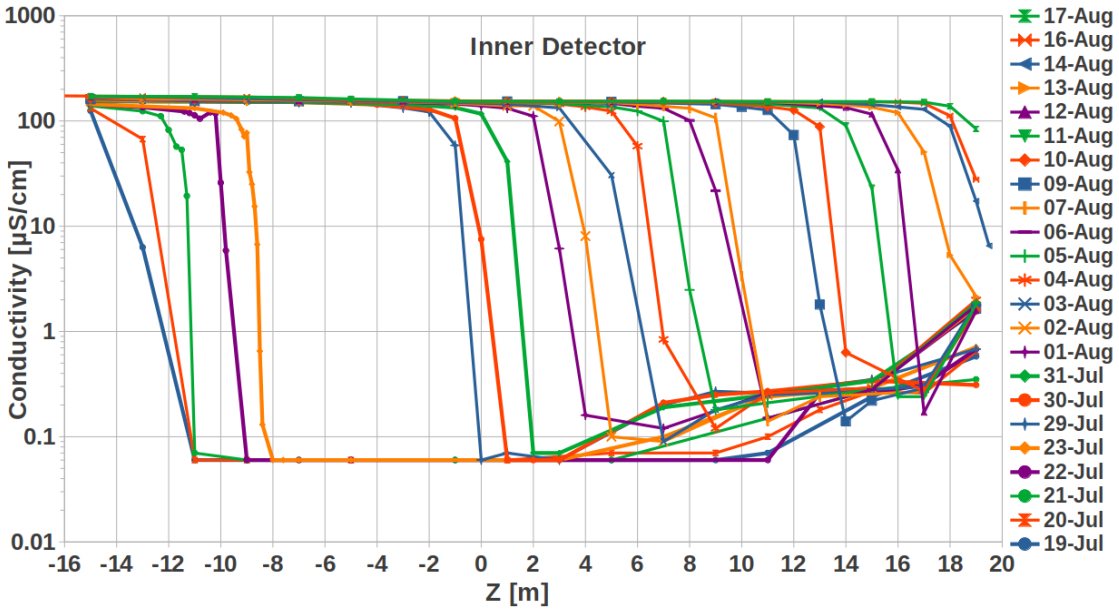


Figure 4.3: The conductivity of water samples taken at a range of heights throughout the SK ID over the course of first Gd loading, in 2020. The sampling system encountered issues between the 23$^{\text{rd}}$ and 29$^{\text{th}}$ of July.

the water where it would be recovered and fitted with another diffuser to repeat the process. A second drone was provided by the Canadian national laboratory TRIUMF to help guide the drone and monitor the drone's tether. This operation was the first use of drones to carry out maintenance in SK and prevented a long downtime that would be needed to drain the tank.

Initially planned for March 2020, the COVID-19 pandemic delayed the first Gd loading to July 14th, 2020 at 10:29. Gd-loaded water was pumped to the bottom of the tank at a rate of $60\,\text{t}\,\text{h}^{-1}$. To monitor the rise of the Gd-water front, water was sampled each day at a range of heights throughout the detector and the conductivity was measured. Ultrapure water has low conductivity, but the presence of the charged ions of gadolinium sulphate significantly increases this, so the conductivity of the samples can be used to differentiate between pure and Gd-loaded water. As shown in figure 4.3, a clear Gd front was observed during loading, advancing at a rate of approximately $1.5\,\text{m}\,\text{d}^{-1}$, as expected. Gd loading concluded on the 17th of August, 2020 with no major issues.

Both during and after the loading, the concentration of $Gd_2(SO_4)_3 \cdot 8\,H_2O$

was accurately measured with an atomic absorption spectrometer (AAS) using water sampled at various positions throughout the tank, alongside the conductivity measurements. Once loaded, the average $Gd_2(SO_4)_3 \cdot 8\,H_2O$ concentration was found to be (271 ± 4) ppm, corresponding to an ionic gadolinium concentration of (114 ± 2) ppm, with uniformity throughout the tank. This confirmed the desired homogeneity and concentration of the Gd loading.

To further verify the presence of gadolinium sulphate at the expected concentration, the neutron capture time was measured using an Americium-Beryllium (AmBe) source embedded in a bismuth germanate (BGO) crystal. AmBe produces a 4.4 MeV gamma and a neutron when it decays. The gamma scintillates in the BGO crystal and the resulting scintillation light can be used to fix the neutron's emission time and tag the decay event. Neutron capture times were measured in November of 2020 using AmBe in various positions throughout the detector, an example of which is shown in figure 4.4. Fitting provides a neutron capture time of $(115 \pm 1)\,\mu$s which corresponds to a gadolinium ion concentration of (111 ± 2) ppm, in agreement with the AAS measurements.

More gadolinium was loaded into SK's water in the summer of 2022, from June 1st to July 5th to increase the gadolinium concentration up to 0.06% by mass, acting as an interim phase before loading to the full targeted concentration of 0.2%. No date is currently set for this final Gd loading.

A wide range of analyses utilising neutron capture are now underway at SK with vastly improved sensitivity. Though the main motivation for SK-Gd was to measure the DSNB, many other analyses benefit from the gadolinium loading. In particular, the identification of reactor neutrinos for the first time in SK would be a momentous achievement only made possible through gadolinium loading.

4.5 WIT Coincidence Trigger

In searches for relatively low rate events like reactor neutrino interactions, efficiency must be maximised where possible. While the events that fail WIT's requirements are inherently of a lower reconstruction quality, it is crucial in an IBD search

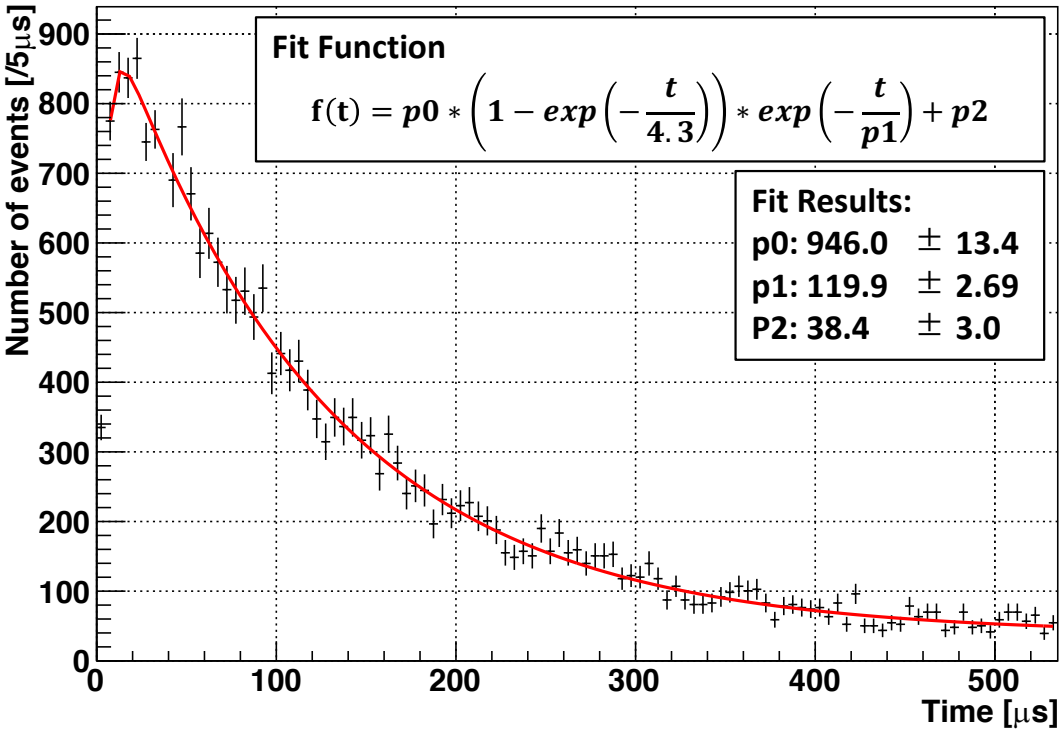


Figure 4.4: An example of neutron capture times and the resultant fit for 30 mins of AmBe calibration data with the source placed in the centre of the detector. Measurements were carried out for a range of source placements.

that both the positron and neutron capture are identified to correctly tag the interaction. In SK-Gd the efficiency of tagging neutron captures has significantly improved and to fully capitalise on this WIT must ensure that even low quality positron events are saved.

WIT's original configuration considered events in isolation. A relaxation of trigger requirements would therefore apply to all events that reach WIT. In the particularly low-energy, high-background regime close to the Cherenkov threshold, this results in a significant increase in background rate. WIT's trigger thresholds are effectively set at a point where minor fluctuations in rates can be handled, but loosening these cut criteria quickly reaches a point where WIT, or more specifically the WIT organiser, cannot keep up with the data rate.

The dominant sources of background in WIT produce events which are randomly time-distributed with no expectation of there being any spatial coincidence between successive events. Coincidence requirements on spatial and temporal distance

between events therefore act as a powerful discriminator between signal and background. Given the frequency of background events in WIT there will still be significant impurities; however, a coincidence requirement within the WIT system would allow an increased efficiency of IBD pair detection with a significantly lower total rate increase compared to reducing the standard trigger thresholds.

To realise this, I modified the logic of WIT's main trigger process to impose a coincidence cut on all successive events. Events are processed largely the same as in its original configuration; however, events must initially pass a relaxed set of trigger requirements before being temporarily saved in a buffer. These requirements are merely lowered thresholds of the same event variables as in the usual WIT process. The event is tagged with a trigger ID number which represents the highest trigger threshold it passed: 0 for relaxed, 1 for the usual trigger requirements or 2 for a high-level STORE goodness, described in section 3.6.

The next event is then processed, and its vertex is compared to that of the event in the buffer. If the events pass the preset positional (Δr) and temporal (Δt) coincidence thresholds, the event in the buffer is saved to file regardless of its trigger ID and the new event takes its place, flagged as being a coincident event.

The event now stored in the buffer will be saved regardless of its coincidence with the subsequent event; however, it must still be temporarily stored in memory to check the next event's coincidence. Given the trigger rate of WIT, it is entirely possible that for three successive events, any two may form an IBD pair with the third being background. It is therefore important to check the coincidence of all events, even if the buffered event is already in coincidence with its previous event. Practically, this means that all events are evaluated for coincidence with the buffered event before replacing it, regardless of the outcome of the coincidence check.

If the incoming and buffered events are not in coincidence, the buffered event will only be saved if it is flagged with a trigger ID of 1 or 2, i.e. it successfully passed the normal WIT trigger requirements.

As well as the advantage of increased online pair efficiency, the coincidence trigger also simplifies the identification of coincidence pairs in the offline data,

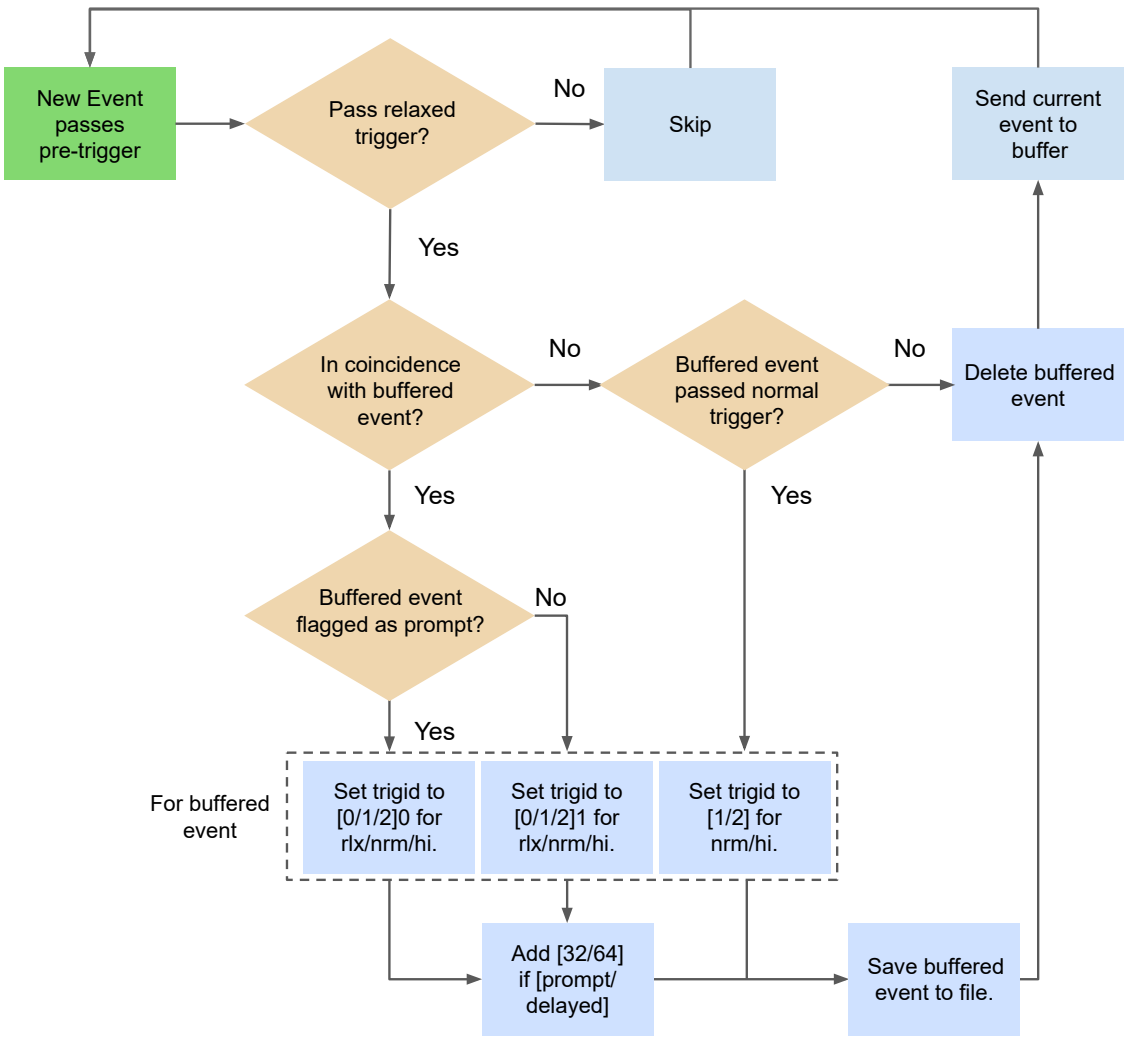


Figure 4.5: The logic of the WIT coincidence trigger, including the trigger identification numbering scheme.

as coincident pairs can be tagged in WIT itself. For events which are first in their coincidence pairs, 32 is added on to their existing trigger ID and saved with the event. These prompt events therefore have a trigger ID of 32, 33 or 34 if they passed the most stringent out of the relaxed, normal or high levels of trigger respectively. Delayed events have 64 added to their trigger ID, regardless of if subsequent events are in coincidence with them. The logic for this trigger ID scheme and the coincidence trigger as a whole is shown in figure 4.5. Additionally, a unique identifier of the event's partner is saved. Pair events can therefore be easily selected in the data through their trigger ID and then matched up with their partner through the identifier stored in the event.

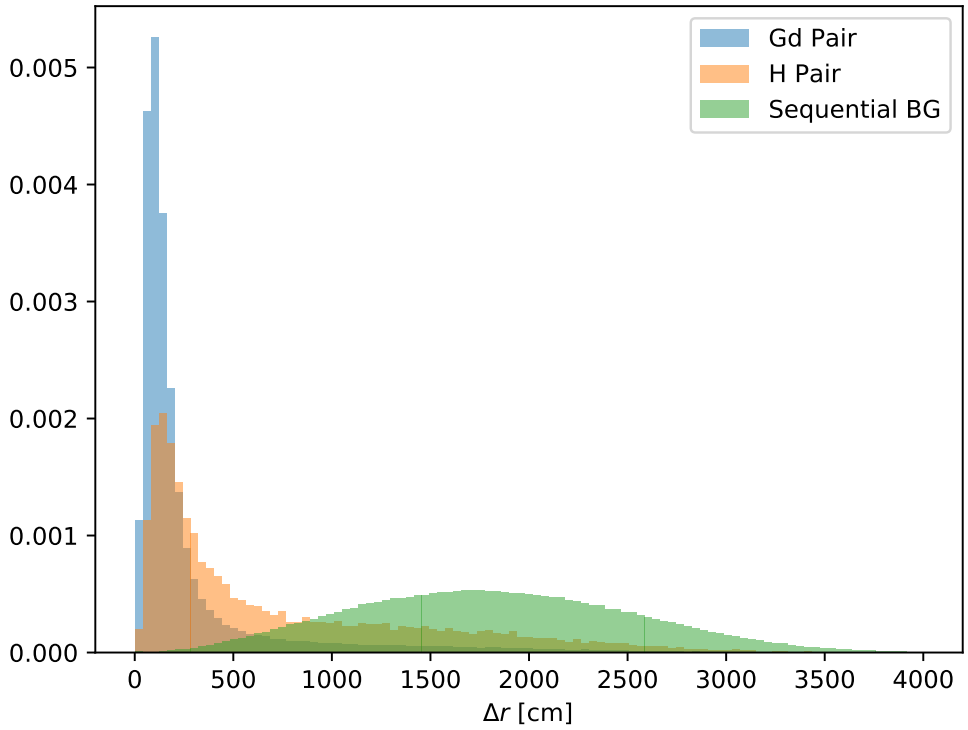


Figure 4.6: The distance between prompt and delayed event vertices for simulated gadolinium and hydrogen capture, reconstructed by online BONSAI, compared to the distance between sequential events in a sample of WIT data considered to be background.Gadolinium concentration is taken to be that after first loading (0.02% by mass).

With the relaxed thresholds set to that of the normal, the processing time of WIT with the added coincidence trigger logic was measured to be less than 1% longer of that of the original WIT configuration, within typical fluctuations in processing time.

The coincidence requirements for two sequential events are currently set to be $\Delta r < 5\,\mathrm{m}$ and $\Delta t < 512\,\mu\mathrm{s}$. The Δr requirement was based on an existing spallation-related neutron cloud analysis using WIT data, and removes a significant amount of background pairs as seen in figure 4.6 when compared to MC reactor IBD pairs.

The Δt requirement is minimal, requiring that two events must be within 30 hardware triggers of one another. It can be gathered from figure 4.7 that this is a less powerful discriminator than the Δr cut, and was chosen initially to retain as many hydrogen capture events as possible. Especially with an increased

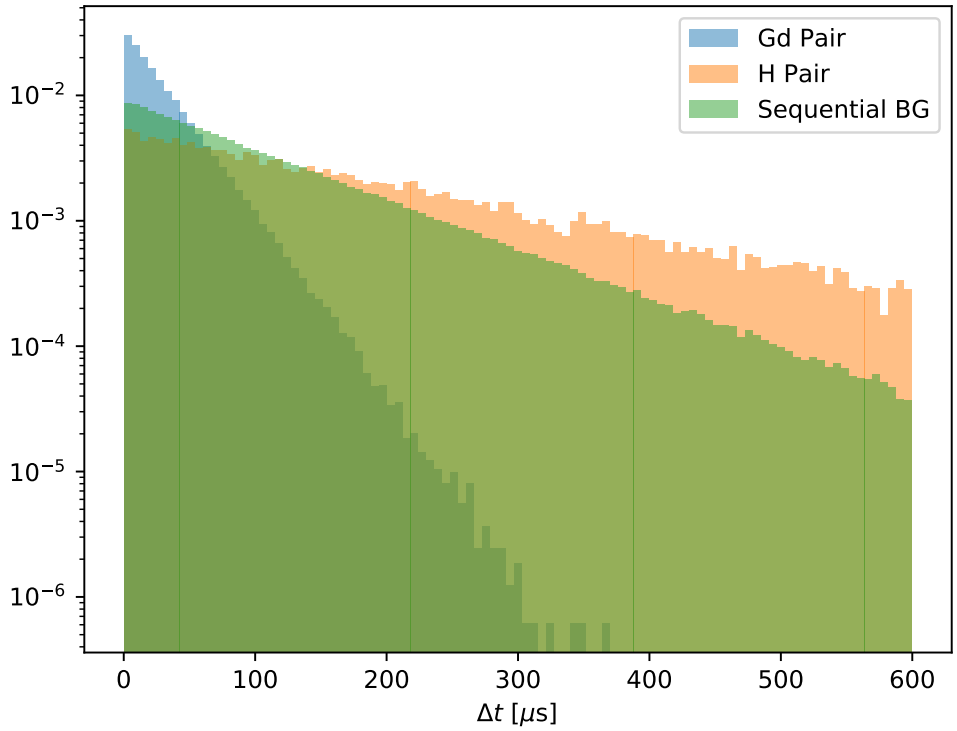


Figure 4.7: The time difference between prompt and delayed events for simulated gadolinium and hydrogen capture, reconstructed by online BONSAI, compared to the time between sequential events in a sample of WIT data considered to be background. Gadolinium concentration is taken to be that after first loading (0.02% by mass).

Gd concentration following future loading, there is scope to impose a tighter temporal coincidence requirement.

The WIT organiser is the main bottleneck of the WIT system, with a more-than-linear processing time dependence on the number of events. The main impact of the coincidence trigger on processing time can therefore be measured simply in the produced data rate. To be effective, the coincidence trigger must allow a significant increase in IBD rate without significantly increasing the background rate. Figures 4.8 and 4.9 show that imposing coincidence requirements of $\Delta t < 512\,\mu$s and $\Delta r < 5\,$m results in no significant increase in background efficiency when relaxing the STORE goodness and N18 trigger thresholds, while significantly increasing the efficiency for simulated reactor pairs for both gadolinium and hydrogen capture. Here, the background efficiency is evaluated using sequential events in real WIT

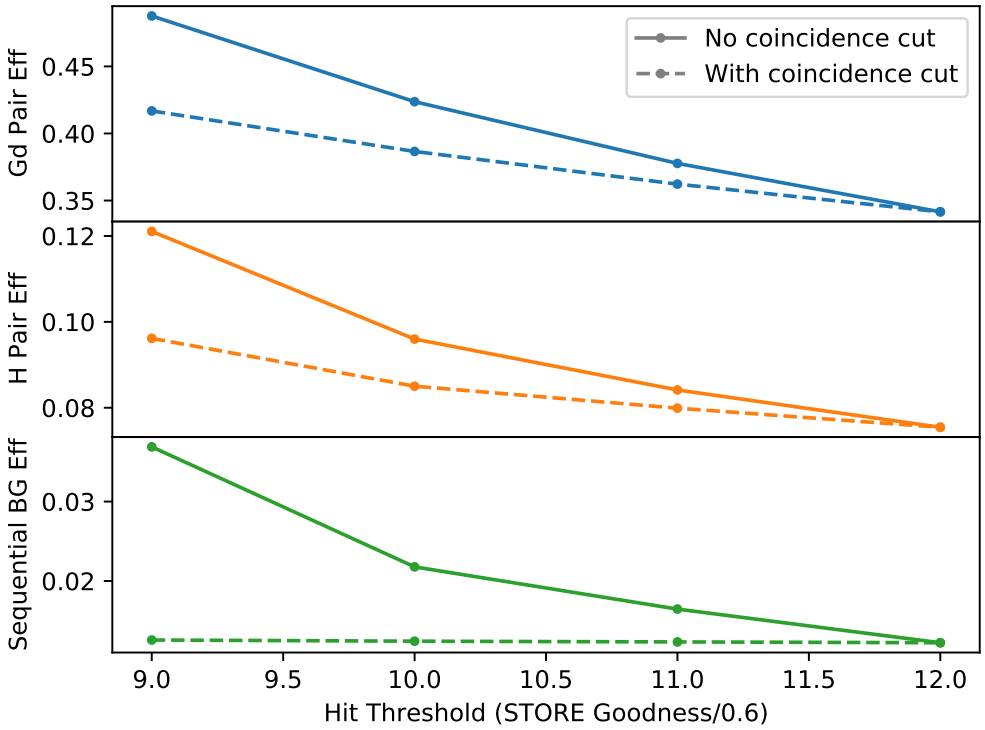


Figure 4.8: WIT detection efficiencies when relaxing the hit threshold (and therefore STORE goodness threshold) for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs.

data, i.e. the efficiency of the nth and $(n+1)$th events passing the cuts when paired. The BONSAI goodness threshold of > 0.3 could be considered a relatively minor cut, so figure 4.10 takes the standard trigger requirement to be a tighter, but still minor, 0.35 to more clearly demonstrate the effect of the coincidence requirement on background. Especially when combined, these relaxed trigger thresholds can provide a significant increase in IBD pair detection efficiency with minimal increase in background trigger rates.

The coincidence trigger came online in March of 2020, ahead of Gd loading. Given the high trigger rate following the 2018 tank open when compared to SK-IV levels, the trigger thresholds for coincident pairs were initially set to the same as the normal thresholds, to verify that the coincidence trigger did not interfere with any part of the WIT data stream. In this configuration, events in coincidence would

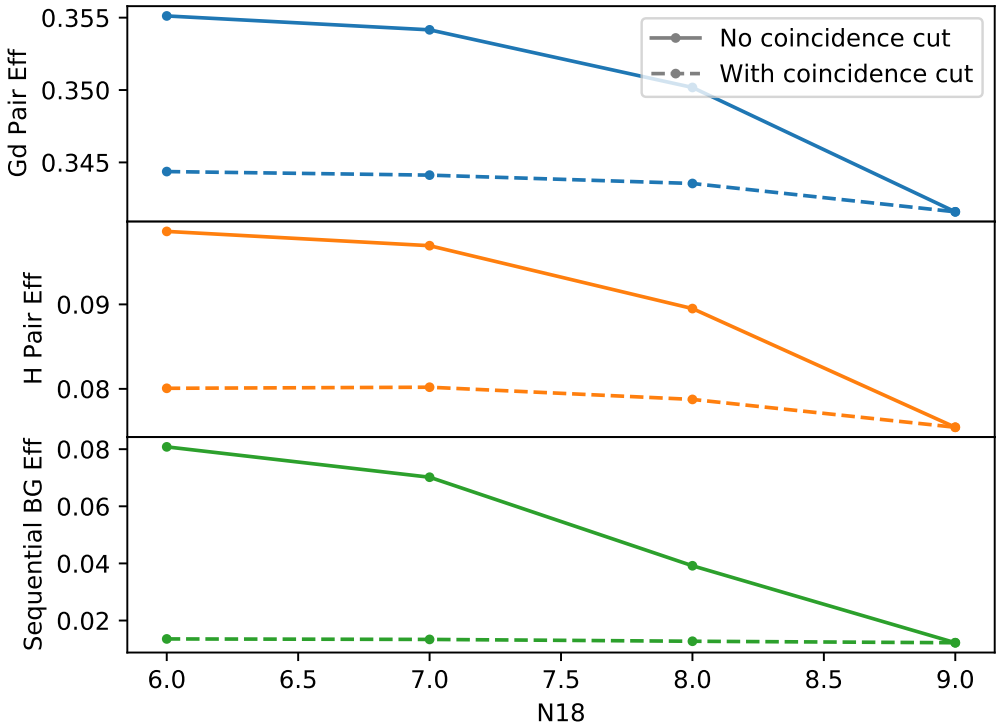


Figure 4.9: WIT detection efficiencies when relaxing the N18 threshold for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs.

still be identified as pairs, but would have to pass the usual WIT requirements.

During Gd loading, the trigger rate increased and in response the normal trigger values were tightened: the hit threshold was increased from 11 to 12 (therefore increasing the STORE cut by 0.6), the minimum N18 requirement was increased from 9 to 10, and a BONSAI goodness requirement of > 0.3 was imposed. The hit threshold and N18 requirement was also increased for coincident pairs, but the BONSAI goodness requirement was not introduced. This configuration meant the coincidence trigger was saving extra events in addition to the normal trigger scheme; paired events would not need to meet any BONSAI goodness requirement. No issues were found while WIT ran in this configuration and the extra events were seemingly saved correctly.

The trigger rate has effectively stabilised for the SK-Gd period. N18 was

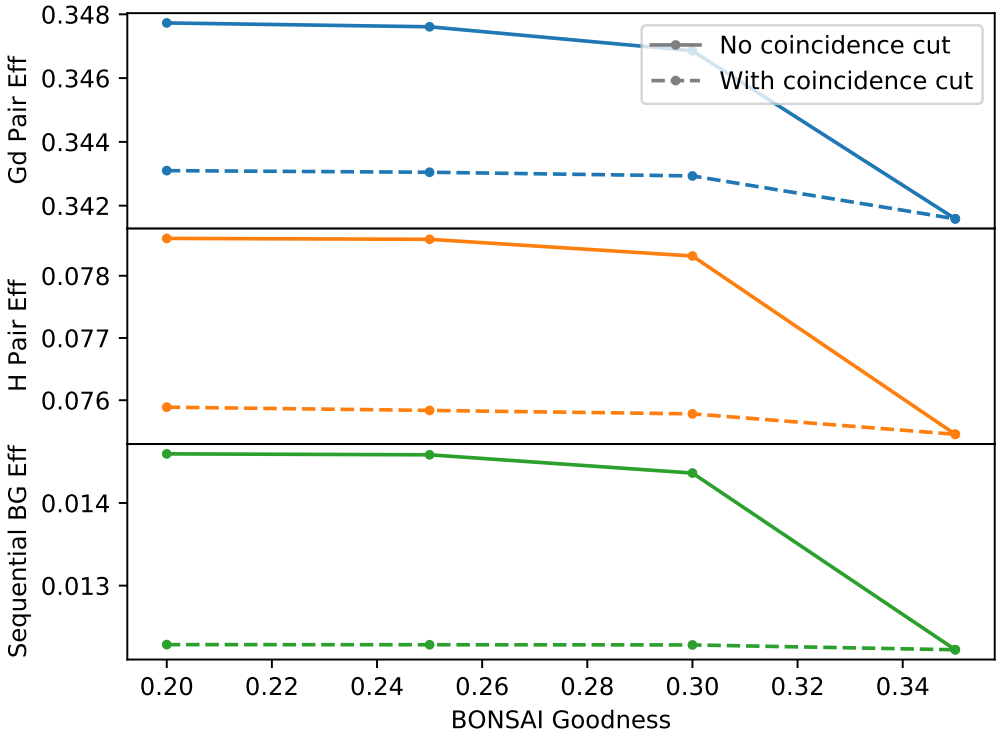


Figure 4.10: WIT detection efficiencies when relaxing the BONSAI goodness threshold for all events, and for coincident pairs only. Sequential events in WIT data (BG) are compared to simulated reactor IBD pairs.

therefore reduced to 9, though the hit threshold was left at 12. There is scope to reduce the relaxed trigger requirements further, based on the potential gains shown in figures 4.8, 4.9 and 4.10, especially with the future expansion of the WIT machines. However, in the interest of simplifying ongoing analyses it is generally preferable to minimise disruption in the available data, so this has not yet been fully explored. The reactor analysis specifically assumes the tightest online trigger thresholds in the period to provide a stable background rate and avoid the need to consider different signal efficiencies for different periods in the data. As such, WIT continues to operate such that paired events need not pass a BONSAI goodness threshold but must otherwise meet the same requirements as singles.

5
Reactor Neutrinos

Contents

5.1 Nuclear Reactors 72
5.2 Nuclear Power in Japan 79
5.3 Neutrino Production 82
5.4 Reactor Neutrino Experiments 85
5.5 Monitoring Nuclear Reactors 89
5.6 SK-Gd and Reactor Neutrinos 91
5.7 Simulating Reactor Neutrinos: SKReact 93

5.1 Nuclear Reactors

In 1938, expanding on work by Enrico Fermi and his colleagues [70], the team of Otto Hahn, Fritz Strassman, Lise Meitner and Otto Robert Frisch discovered that bombarding uranium with neutrons resulted in the splitting of atomic nuclei [71]. This process was dubbed nuclear fission in reference to binary fission, the name at the time for the splitting of living cells. Fission was observed to produce a large amount of energy, both through electromagnetic radiation and in the kinetic energy of the products. This energy stems from the difference in binding energy between the reactants and products. Following the discovery that more than one neutron is produced in fission, Leó Szilárd theorised that these could induce further

fission, triggering a chain reaction that may yield a vast amount of energy. Given the advent of World War II, Szilárd advocated for restricting the publication of these findings in the interest of preventing development of weapons utilising this chain reaction by foreign powers. He and Fermi then proposed the construction of a nuclear reactor, using natural uranium as fuel and graphite as a neutron moderator. The graphite would slow the neutrons to thermal energies, increasing the cross section of their capture on the uranium, thereby increasing the rate at which they would induce fission.

Fermi and Szilárd's proposal was approved and on 2 December 1942, Chicago Pile-1 (CP-1) became the world's first artificial, self-sustaining nuclear reactor [72]. It consisted of a wooden frame supporting piles of graphite blocks, assembled by Fermi and his team. The uranium was fashioned into round-ended rods and embedded in alternating layers of graphite blocks. CP-1 made use of control rods composed of cadmium, a strong neutron absorber, attached to wooden strips that could be dropped into the pile to halt the chain reaction. Given the low power CP-1 produced, there was no radiation shielding and the pile was not actively cooled.

The work on CP-1 was absorbed into the newly-formed Manhattan Project, the Allied efforts to develop nuclear weapons during the Second World War. The Manhattan Project subsequently produced a range of nuclear reactors with the primary aim of producing plutonium to be used in nuclear weapons. After the war, the reactor technologies developed were shared worldwide in the interest of promoting the use of nuclear power. The first non-research reactor was installed into the USS Nautilus submarine in January of 1954 followed in June by the first civil nuclear power plant, the Obninsk-1 reactor in the Soviet Union. The first commercial plant reached criticality in 1956 in the United Kingdom.

The fuel typically used in modern nuclear reactors has a much higher ^{235}U concentration than that used in CP-1. CP-1 used natural uranium, comprised of 0.7% ^{235}U and over 99.2% ^{238}U. ^{238}U is non-fissile in that slow (thermal) neutrons do not induce fission. Under neutron bombardment, however, ^{238}U can produce ^{239}Pu in the reaction

$$^{238}\mathrm{U} + n \rightarrow {}^{239}\mathrm{U} \rightarrow {}^{239}\mathrm{Np} + e^{-} + \bar{\nu}_e \rightarrow {}^{239}\mathrm{Pu} + e^{-} + \bar{\nu}_e. \tag{5.1}$$

Though ^{239}Pu is fissile, it is a relatively slow process to convert the ^{238}U to ^{239}Pu. ^{235}U is therefore the dominant source of fission in natural uranium. A process known as enrichment increases the percentage of ^{235}U by separating the isotopes in gaseous form in a centrifuge. For use in power reactors, low-enriched uranium (LEU) is sufficient, with about 3–5% ^{235}U, though the definition of LEU extends to 20% enrichment. Highly enriched uranium (HEU) composed of $> 20\%$ ^{235}U is used in nuclear weapons; however, weapons-grade uranium is typically used in these applications, with a ^{235}U concentration of $> 85\%$.

Though commercial reactors all use steam-driven turbines to generate electricity, the method of extraction of this steam varies as well as the general reactor design. These designs differ in many ways, crucially in the materials used as their coolant and neutron moderator. There are 443 operable commercial nuclear power plants in the world [73], each taking one of six forms:

- Pressurised water reactor (PWR)

  The most common design, PWRs were originally designed for use on submarines and are still used in this capacity. PWRs contain their enriched uranium fuel and control rods within a pressurised vessel full of ordinary water. The water acts as both the neutron moderator and coolant. The pressure is kept at a level such that the water heated by fission does not form steam in the main vessel but instead is pumped to an external vessel. Here the heat is exchanged between the reactor water and an external water supply without cross contamination. The heated water is allowed to boil to steam which is used to drive a turbine, generating electricity, before being condensed and pumped back into the steam vessel. The reactor water is also then pumped back into the reactor vessel for re-heating. An inherent safety feature emerges in that any failure in the pressure vessel would release the water as steam, thereby depriving the reactor of its neutron mediator, reducing

the fission rate. Another advantage to this design is that the reactor water is kept in a closed loop, separate from the water used to drive the steam turbine. As the reactor water is contaminated with radionuclides, it is preferable to keep this separate to the water used in power generation, removing the need for radiation shielding of the generators and easing maintenance.

- Boiling water reactor (BWR)

  BWRs remove the secondary vessel and hence separate water loop from PWRs and instead allow the reactor water to form steam in the top of the main reactor vessel where it is extracted for power generation, condensation and recirculation. Though this is a simpler design, the reduced neutron moderation in the reactor vessel results in an overall lower efficiency. Additionally, the steam turbine must be shielded to prevent damage from the radioactively contaminated water and maintenance requires radiological protection. An upgraded design of BWRs known as the advanced boiling water reactor (ABWR) improves many aspects of the design while retaining the core operating principle and is now in limited use globally. The first ABWR to reach criticality was unit 6 of the Kashiwazaki-Kariwa plant in Japan in 1995 [74].

- Pressurised heavy water reactor (PHWR)

  PHWRs operate under the same principle as a PWR, but use natural uranium as their fuel. This small fissile fraction warrants heavy water to act as a more efficient moderator to effectively sustain the chain reaction. Given the low fissile fuel density, more frequent replacement of the fuel rods is required. Fuel rods are placed in separate horizontal pressure tubes through which the heavy water flows, in contrast to the large pressure vessel in PWRs. This allows for the replacement of individual rods without disrupting the operation of the reactor as a whole. Though the energy produced per unit fuel is lower, there are great cost savings made through skipping the enrichment process.

- Advanced gas-cooled reactor (AGR)

  These reactors do not use water in their design outside of the steam generators. Graphite is used as a neutron moderator and carbon dioxide as the coolant. The reactor is contained within a steel and concrete pressure vessel, into which the steam heat exchangers penetrate. A development of older gas cooled reactors (GCRs), specifically the British Magnox reactor designed primarily to generate weapons-grade plutonium, AGRs offer the option of online refuelling, similar to PHWRs. This reactor design forms the majority of nuclear reactors deployed in the UK.

- Light water graphite-moderated reactor (LWGR)

  As the name suggests, graphite is used as a neutron moderator in LWGRs. Ordinary water is the coolant and, like in BWRs, is allowed to boil in the vertical pressure tubes surrounding the fuel rods. This design presents a risk in that a positive feedback loop can occur. If the core temperature rises, more water boils to steam in the vessel. Given the graphite acts as a neutron moderator, this simply reduces the neutron absorption from the water without reducing the neutron moderation, increasing the rate of fission. The No. 4 reactor at the Chernobyl Nuclear Power Plant used the Soviet RBMK design of an LWGR and, through a series of missteps, initiated such a loop. A further flaw in the RBMK design, specifically the graphite tips of the control rods, meant that initiating an emergency shutdown significantly increased the fission rate, which led to the meltdown at Chernobyl. The design flaws were addressed and a number of RBMK reactors have since been in operation for over thirty years with no serious incidents.

- Fast neutron reactor (FNR)

  FNRs do not use a moderator and rely on fission induced by fast neutrons. These high-energy ($\sim$1 MeV) neutrons are sufficiently energetic to induce fission in the non-fissile ^{238}U. ^{239}Pu produces three neutrons in its most common fission reaction, compared to two or three for ^{235}U. This neutron yield is

sufficient to sustain a chain reaction if the fuel is sufficiently enriched, offsetting the reduced neutron capture cross section of fast neutrons. Additionally, the extra neutrons can be used to breed ^{239}Pu in ^{238}U blanketing the reactor core. The result is a reactor which can make efficient use of ^{238}U and therefore increase the energy output per unit fuel compared to other designs, with the added advantage of greatly reducing the long-lived radioactive isotopes in spent fuel. A typical coolant in such a design is liquid metal. This has the advantage of superior heat transfer over water and carbon dioxide gas, as well as a far higher boiling point. The disadvantage is that liquid metal is typically more toxic and reactive than water, requiring more robust and therefore expensive cooling systems. If an FNR produces more ^{239}Pu than they consume, they are referred to as a fast breeder reactor (FBR).

The breakdown of reactor types and their commercial usage worldwide as of September 2021 is given in table 5.1.

Type	Main Countries	Number	GWe	Fuel	Coolant	Moderator
PWR	USA, France, Japan, Russia, China, S. Korea	304	289.1	Enriched UO_2	Water	Water
BWR	USA, Japan, Sweden	62	63.1	Enriched UO_2	Water	Water
PHWR	Canada, India	49	24.5	Natural UO_2	Heavy water	Heavy Water
AGR	UK	14	7.7	Natural U, enriched UO_2	CO_2	Graphite
LWGR	Russia	12	8.4	Enriched UO_2	Water	Graphite
FNR	Russia	2	1.4	UO_2, PuO_2	Liquid Sodium	None

Table 5.1: Operable commercial nuclear power plant designs throughout the world [73]. The column "GWe" refers to the total electrical power capacity, measured in GW.

Nuclear power far surpasses the fuel efficiency provided by fossil fuels and is often cited as a viable alternative to gas, coal and oil stations in the face of

climate change. Despite this, a small number of high-profile incidents, especially the Chernobyl disaster, as well as the association with nuclear weapons has resulted in an often negative public perception of nuclear power. One common criticism of nuclear power is in regard to its spent fuels. Though nuclear power produces a minute amount of greenhouse gases in comparison to fossil fuels, spent fuel rods present environmental risks of their own.

After the immediate halting of fission in a reactor, the flow of coolant must continue due to decay heat in the fuel. This is the significant heat produced by radioactive decay of short-lived isotopes, amounting to 6–7% of the core's total heat produced during fission. While this quickly decreases, spent fuel rods must still be stored with some form of cooling once removed from the core, often for many years kept in pools of cycling water which also acts as a shield to the radiation. This radioactivity is the primary concern with regards to long-term storage of spent fuel. The half life of elements producing harmful radiation can be $\sim 10,000$ y, requiring careful, effectively permanent storage of the hazardous waste. Currently, spent fuel is disposed of in dedicated facilities, often in disused mines to minimise exposure to the environment. FNRs seek to combat this issue of disposal by inducing fission in the long-lived, fissionable radioactive actinides, leaving only the fission products. These fission products (primarily strontium-90 and caesium-137) have half lives $\sim$30 y, far reducing the length of time spent fuel poses risk due to radioactivity.

Another complaint from nuclear power critics is regarding the risk of a meltdown. A meltdown occurs when, through loss of coolant or runaway chain reactions, the fuel rods are heated beyond their melting point. The separate fuel rods may then combine after melting their housings, increasing the fission rate further and releasing a dangerous amount of energy. A number of meltdowns and criticality incidents have occurred in the history of commercial nuclear power, most notably the Chernobyl disaster in Soviet Ukraine of 1986 and the meltdown at the Fukushima Daiichi plant in Japan following the 2011 Tohoku earthquake. Despite these events, the scientific and statistical consensus is that the human impact of nuclear power is one of the most beneficial of all energy sources, rivalling that of renewables like

solar and wind power. Nuclear power plants are, however, particularly expensive to build and maintain, with long construction times.

5.2 Nuclear Power in Japan

Japan was a strong adopter of nuclear power since its first reactor, Tokai 1, came online in July of 1966. By the year 2011, approximately 30% of Japan's electric power supply came from its 33 nuclear plants and Japan is home to the largest capacity nuclear power plant in the world: Kashiwazaki-Kariwa, which also houses the world's first advanced boiling water reactor. Japan aimed to provide 40% of its power through nuclear by 2017. These plans were halted, however, when the magnitude 9.1 Tohoku earthquake struck off of the east coast of Japan's main island of Honshu in 2011. Though all of Japan's nuclear plants survived the initial quake with no failures, the resulting tsunami triggered the meltdown of cores 1, 2 and 3 of the Fukushima Daiichi plant.

Fukushima's reactor cores correctly shut down, ceasing fission, upon detecting the earthquake. Despite this, the flow of coolant is still required due to the decay heat of the fuel. In the absence of power generated by the reactor, emergency diesel generators are used to power the coolant pumps. These correctly activated; however, the tsunami was high enough to flow over the sea wall and into the power stations, flooding the basements where the generators were located. The flooding disabled the generators and hence the coolant pumps. Backup generator sites were similarly flooded. Three of the six reactor cores then overheated and underwent a meltdown; the other three cores were offline for maintenance at the time of the earthquake. Overheating also led to hydrogen gas to form in three of the reactors, eventually triggering a hydrogen-air explosion in each.

This incident led to a shift in public perception of nuclear power in Japan, as well as a re-evaluation of Japanese reactor safety standards. In response, all of Japan's reactors were shut down and each plant is now required to implement more stringent tsunami-related safety measures (including extended sea walls) as well as gain local government approval before they can restart operation. Japan's Nuclear

Regulation Authority (NRA) must also approve the restart. The first to restart was Sendai-1 in August 2015, followed in October by Sendai-2. A further 8 have since restarted and 15 have applied for approval. While some local governments and communities have resisted, Japan's severely reduced energy production capacity motivates the restarts, with the insurance of far more stringent safety standards.

There are a number of reactors in relatively close proximity to SK, most prominently the Takahama, Ohi, Kashiwazaki and Mihama plants. Extending further, there are a number of Japanese reactors at the extremities of Honshu and beyond, as well as a number of reactors on the east coast of South Korea that are both more powerful and closer to SK than some Japanese reactors. Table 5.2 lists all active Japanese and South Korean reactors, ordered by (thermal power P_{th})/(distance to SK R_{SK})^2 to approximately demonstrate the most relevant contributions to reactor neutrino flux at SK. The full list of Japanese and South Korean reactors is given in appendix A.1

Plant	Core	Type	P_{Th} [MW]	R_{SK} [km]	P_{Th}/R^2_{SK} [MW km^{-2}]
Mihama	3	PWR	2440	145.6	0.115
Ohi	4	PWR	3423	178.8	0.107
	3	PWR	3423	178.8	0.107
Takahama	4	PWR	2660	191.4	0.073
	3	PWR	2660	191.4	0.073
Ikata	3	PWR	2660	560.3	0.008
Shin-Kori*	4	PWR	3983	732.0	0.007
	3	PWR	3983	732.0	0.007
Genkai	3	PWR	3423	753.7	0.006
	4	PWR	3423	753.7	0.006
Shin-Wolsong*	1	PWR	2825	708.9	0.006
	2	PWR	2825	708.9	0.006
Hanul*	6	PWR	2825	709.0	0.006
	5	PWR	2825	709.0	0.006
	4	PWR	2825	709.1	0.006
	3	PWR	2825	709.2	0.006
	2	PWR	2775	709.3	0.006
	1	PWR	2775	709.4	0.006
Kori*	4	PWR	2912	732.5	0.005
	3	PWR	2912	732.6	0.005
Shin-Kori*	1	PWR	2825	732.0	0.005
	2	PWR	2825	732.0	0.005
Wolsong*	4	PHWR	2061	708.6	0.004
	3	PHWR	2061	708.7	0.004
	2	PHWR	2061	708.8	0.004
	1	PHWR	2061	708.9	0.004
Sendai	2	PWR	2660	829.8	0.004
	1	PWR	2660	829.8	0.004
Kori*	2	PWR	1882	732.7	0.004
Hanbit*	3	PWR	2825	987.0	0.003
	4	PWR	2825	987.0	0.003
	5	PWR	2825	987.0	0.003
	6	PWR	2825	987.0	0.003
	1	PWR	2787	987.0	0.003
	2	PWR	2787	987.0	0.003

Table 5.2: Nuclear power reactors currently operating in Japan and South Korea, ordered by $(P_{th})/(R_{SK})^2$ where P_{th} is thermal power output capacity of the reactor and R_{SK} is the distance from the core to the SK detector [75]. An asterisk by a plant's name denotes a South Korean plant.

5.3 Neutrino Production

Nuclear fission produces isotopes rich in neutrons. These fission fragments are naturally unstable, decaying to stable isotopes via beta decay ($n \rightarrow p + e^- + \bar{\nu}_e$). Given the high frequency of fission reactions in reactor cores, and that a single fission reaction yields around 6 neutrinos on average, nuclear reactors are particularly prominent sources of electron antineutrinos; the first neutrinos ever detected were from a nuclear reactor, in the Cowan-Reines experiment in 1956 [7].

The exact spectrum and flux of neutrinos produced by a reactor varies based on fuel makeup and thermal power. Different fuels produce different combinations of neutron rich isotopes following fission, each with their own neutrino emission spectra and decay branches.

Predicting the electron antineutrino flux from nuclear reactors is non-trivial thanks to the contribution of over 6,000 decay branches to each isotope's production [76], and four primary isotopes which contribute to over 99.9% of the fission occurring in a reactor core: ^{235}U, ^{238}U, ^{239}Pu and ^{241}Pu. Two main approaches to quantify these contributions exist. First is the *ab initio* or "summation" method, which individually calculates and directly sums the contributing decay branches based on their measured beta emission spectra [76]. The alternate approach is to generate virtual decay branches to match measurements of integrated electron emission spectra of a given isotope as a whole, dubbed the "conversion" method [77]. While the *ab initio* method introduces uncertainty from measurements of rare or otherwise not well understood decay branches as well as non-existent measurements, the conversion method's uncertainty comes from linking nuclear charge number and the endpoint energy of the beta decay branches when generating the virtual branches.

Mueller et al. developed a parameterisation of the antineutrino spectra for the four dominant fuels, an exponential fifth order polynomial defined as:

$$S_k(E_{\bar{\nu}}) = \exp(\sum_{p=1}^{6} \alpha_{pk} E_{\bar{\nu}}^{p-1}) \tag{5.2}$$

where α_{pk} is the pth coefficient for fuel k's emission spectra and $E_{\bar{\nu}}$ is the antineutrino energy [78]. An *ab initio* approach was taken for ^{238}U, whereas the spectra for ^{235}U, ^{239}Pu and ^{241}Pu were fit using a conversion approach to the integral beta spectra measured in the 1980s by the group of Walter Mampe at Institute Laue-Langevin (ILL), the most recent and complete example of such measurements [79]. The fit coefficients are given in appendix B.1.

To add to the complexity of producing a total reactor spectrum, determining the fuel makeup of a power reactor at any one time is difficult. Firstly, there are the variations in fuel makeup for different reactor designs, resulting in a different total spectrum. To further complicate this, fuels undergo fission at different rates resulting in an ever-changing fuel makeup, an effect known as burnup [80]. Power reactor operation typically involves the redistribution and replacement of fuel rods within the reactor core over time to balance the effect of burnup, maximising the power produced by the plant as well as ensuring spent fuel rods conform to safety standards.

While power reactor design is public information, burnup status and fuel distribution is typically a closely guarded secret due to the research and development companies put into maximising profits from their reactors through delicately balancing burnup and fuel restrictions. To account for burnup, measurements of the fuel fractions from different sampling times during the operation of a reactor can be used, approximating an average makeup over a core's lifetime.

Also publicly available is the monthly load factor (LF) and yearly reference power of a reactor. The load factor of a reactor for a given period is the net electrical energy produced in that period as a fraction (or percentage) of the electrical energy produced were the reactor running at reference power. The reference power is not strictly the maximum power the reactor can produce (load factors are often greater than one), but is the approximate maximum power the reactor aims to produce, typically given in mega watts (MW). The International Atomic Energy Agency (IAEA) publishes yearly reports on the world's power reactor load factors and reference power (modifications to existing reactors sometimes motivates a change of reference power) as part of their Power Reactor Information System (PRIS) [81].

The thermal power P_{th} of a reactor is the thermal energy released per second from the core when operating at reference power, ignoring any inefficiencies in electrical energy extraction. For fuel k, the total energy released per fission which contributes to P_{th} is calculated by:

$$Q_k = E^k_{tot} - \langle E_{\bar{\nu}} \rangle^k - \Delta E^k_{\beta\gamma} + E^k_{nc} \tag{5.3}$$

where E^k_{tot} is the total energy produced in the fission process including the beta decay of all the unstable fission fragments, $\langle E_{\bar{\nu}} \rangle^k$ is the mean total energy carried away by the antineutrinos produced in these decays, $\Delta E^k_{\beta\gamma}$ is the energy of beta electrons and photons from fission fragments that do not contribute to the reactor's energy (and subsequent heat) output and E^k_{nc} is the energy released from neutron capture (which does not induce fission) on materials within the core [82].

The contribution of fuel k to the total fission rate in a reactor running at P_{th} is therefore calculated by:

$$N^k_{fiss} = P_{th} \frac{p_k}{Q_k} \tag{5.4}$$

where p_k is the fuel's total fission contribution fraction. Various measurements of p_k have been made for different reactor designs, the results of some have been given in appendix C.1, including a collection of measurements for the most common reactor types from different sampling times to account for burnup.

Using equations 5.4 and 5.2, the resulting antineutrino spectrum produced by a given reactor of thermal power P_{th} running at a load factor of LF is then:

$$S(E_{\bar{\nu}}) = P_{th} LF \sum_{k=1}^{4} \frac{p_k}{Q_k} S_k(E_{\bar{\nu}}) \tag{5.5}$$

where $k = {}^{235}\mathrm{U}$, ${}^{238}\mathrm{U}$, ${}^{239}\mathrm{Pu}$ and ${}^{241}\mathrm{Pu}$. The spectrum for each fuel in an example PWR with their associated error is shown in figure 5.1.

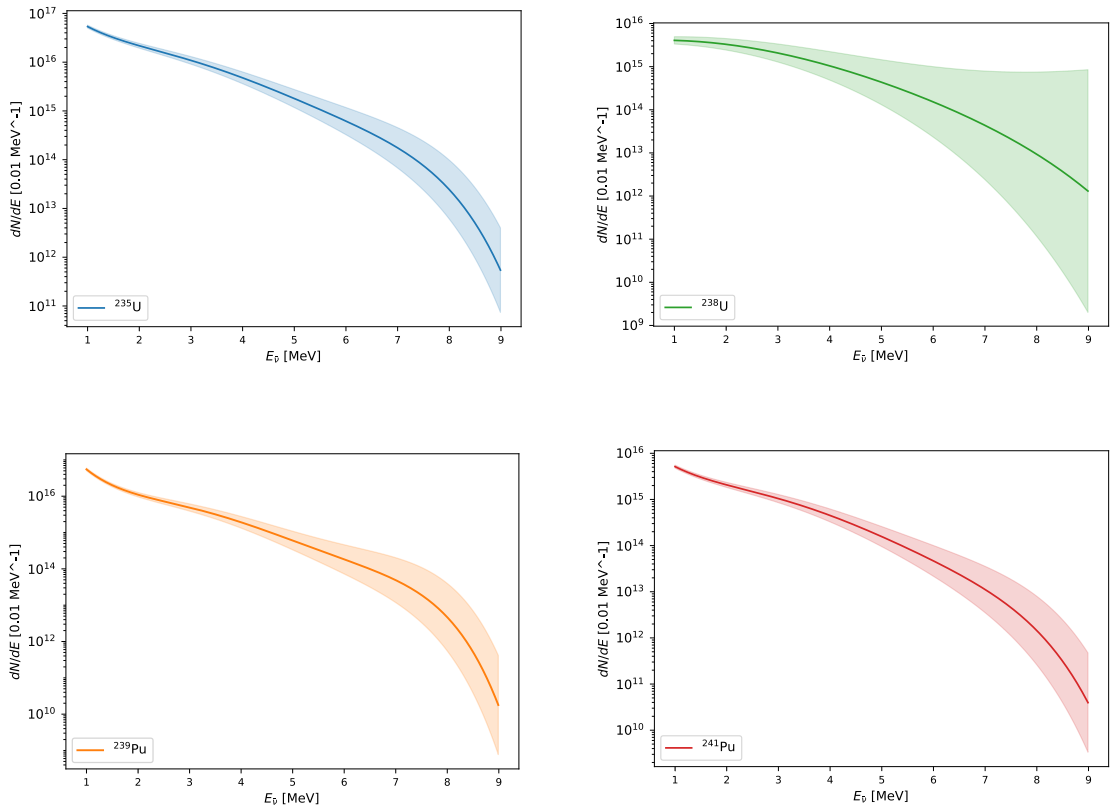


Figure 5.1: The simulated electron-antineutrino spectra with associated errors produced by individual fuels in a 3423 MW core of the Ohi 3 PWR reactor.

5.4 Reactor Neutrino Experiments

Given their abundance and flavour-purity, reactor neutrinos are a cornerstone of neutrino physics, especially in their historical significance in being the first neutrinos ever identified [7]. They provide an opportunity to measure electron antineutrino oscillation, $\sin^2 2\theta_{13}$ in particular. Three reactor experiments have been constructed primarily to provide a high-precision measurement of this oscillation parameter, each with similar detector designs. One, RENO, is still active; however, the other two, Daya Bay and Double Chooz, have now stopped data taking. There is also one large-scale next generation reactor experiment currently under construction in China: JUNO. The details of all these experiments are as follows:

- Daya Bay

The Daya Bay experiment, named for its location in China, was a combination of eight liquid scintillator detectors distributed across three locations investigating neutrinos from the Daya Bay and Ling Ao reactor complexes [83]. From August 2011 until December 2020 it monitored six reactor cores, two at Daya Bay and four at Ling Ao, and provided the first significant measurement of θ_{13} in 2012 [84]. Two of its experimental complexes acted as near detectors, each situated as close to their respective reactor cores as practically possible (360–370 m for Daya Bay, 470–500 m for Ling Ao). The third's location was determined through optimisation of oscillation sensitivity, with geological and practical considerations, acting as a far detector $\sim$2 km from the reactor cores.

The detectors themselves had a target volume of gadolinium doped liquid scintillator (GdLS) contained in a cylindrical acrylic vessel 3 m high and 3 m in diameter. This resulted in a target mass, per individual detector, of 20 t, 40 t for each near detector complex, 80 t for the far. This GdLS vessel was surrounded by a second, concentric vessel containing pure liquid scintillator (LS) which served to identify gammas. Finally, these volumes were surrounded by a third cylindrical vessel made of stainless steel. PMTs were mounted to a structure in this outer volume, filled with mineral oil for suppressing radioactive background from the PMT glass. PMTs were exclusively mounted on the walls of the detector, with circular reflectors employed to direct light from the endcaps to the PMTs. The individual detectors were immersed together in a large water bath surrounded by PMTs, one bath per detector complex. This detected Cherenkov light from cosmic muons, acting as a veto.

- RENO

  The only currently operating experiment in this list, RENO (Reactor Experiment for Neutrino Oscillation), based in Korea, uses a pair of GdLS detectors acting as near and far detectors with a target mass of 16 t per detector and has been running since August 2011 [85]. Using both a LS layer as a gamma catcher and mineral oil layer to suppress background, RENO uses a water

Cherenkov outer cosmic muon veto like Daya Bay. The inner detectors are fully instrumented with PMTs. Its detectors are 292 m and 1380 m from the six cores of the Hanbit nuclear power plant (previously known as Yonggwang) which has the second-largest thermal power output of any plant, second to Kashiwazaki-Kariwa in Japan.

- Double Chooz

  An evolution of the Chooz experiment, which published its results in 1993 [56], Double Chooz (DC) was a pair of GdLS detectors observing neutrinos from the two cores of the Chooz reactor complex in France [86]. The identical near and far detectors were placed 400 m and 1050 m respectively from the average reactor core position. Like Daya Bay, DC used GdLS in its central inner detector, albeit with a smaller target mass of 8 t per detector and with PMTs on the endcaps of the tank in addition to the wall. It also used a LS volume surrounding the target volume, itself surrounded with an oil volume to reduce radioactivity. DC used plastic scintillator as an outer muon veto. It stopped taking data in December 2017 after 6 years of data taking.

- JUNO (Under Construction)

  The Jiangmen Underground Neutrino Observatory (JUNO) is a 20 kt liquid scintillator detector currently under construction in the town of Jinji, China [87]. JUNO is of a far larger scale than existing reactor neutrino experiments, but aims to retain the powerful energy resolution offered by liquid scintillator. In contrast to the experiments described above, JUNO opts for a spherical design and only has a central acrylic vessel of un-doped LS used as the target material, suspended in a water bath. PMTs are mounted around the central vessel, searching both the target volume for scintillation light and the water bath for Cherenkov light from cosmic muons. JUNO will also make use of an extra plastic scintillator tracking detector mounted above the main tank, which will further support the cosmic muon veto. Placed 43 km from the Yangjiang and Taishin nuclear power plants, the detector is at a far longer

baseline than other reactor experiments. This is because JUNO's approach to measuring the mass ordering requires measurements of both θ_{12} and θ_{13}, which is possible at a 43 km baseline.

As well as their usefulness in measuring $\sin^2 2\theta_{13}$, in recent years reactor neutrinos have garnered further interest due to the reactor antineutrino anomaly (RAA) [88]. This is a discrepancy between the expected and detected IBD rates in reactor experiments, where the global average measures a deficit. In contrast, an excess of events was observed in the Liquid Scintillator Neutrino Detector (LSND) [89] and MiniBooNE [90] experiments. Shown in 5.2, recent reevaluations of the flux model have largely resolved the global deficit, reducing the difference to 2.5σ for the Huber-Mueller model of reactor antineutrino spectra, and lower for other models.

Though the flux anomaly is largely resolved, it originally prompted the suggestion that one or more additional neutrino states could be responsible for the difference. These so-called sterile neutrino states would interact only via gravity, explaining their as-yet lack of discovery, and would provide a disappearance channel, leading to a deficit in the flux. A single sterile state would not be able to explain both the global deficit and the LSND and MiniBooNE excess simultaneously. In addition to searches by existing reactor experiments, experiments such as STEREO (Sterile Reactor Neutrino Oscillations) [91], MicroBooNE [92] and the Short-Baseline Neutrino (SBN) program at Fermilab [93] have been established with the primary goal of searching for sterile neutrinos.

Another important reactor experiment is the KamLAND [95]. KamLAND was a 1 kt LS detector situated in the cavern once occupied by the KamiokaNDE experiment in the Kamioka mine, in close proximity to SK. The LS was contained in a 13m diameter balloon made of transparent film, in contrast to the acrylic vessels of modern LS experiments. The balloon was surrounded by oil, instrumented with PMTs and at the time of deployment was the largest target mass of any LS experiment.

While reactor experiments are typically placed as close as possible to reactor cores to maximise both neutrino flux and sensitivity to θ_{13}-driven oscillations, the

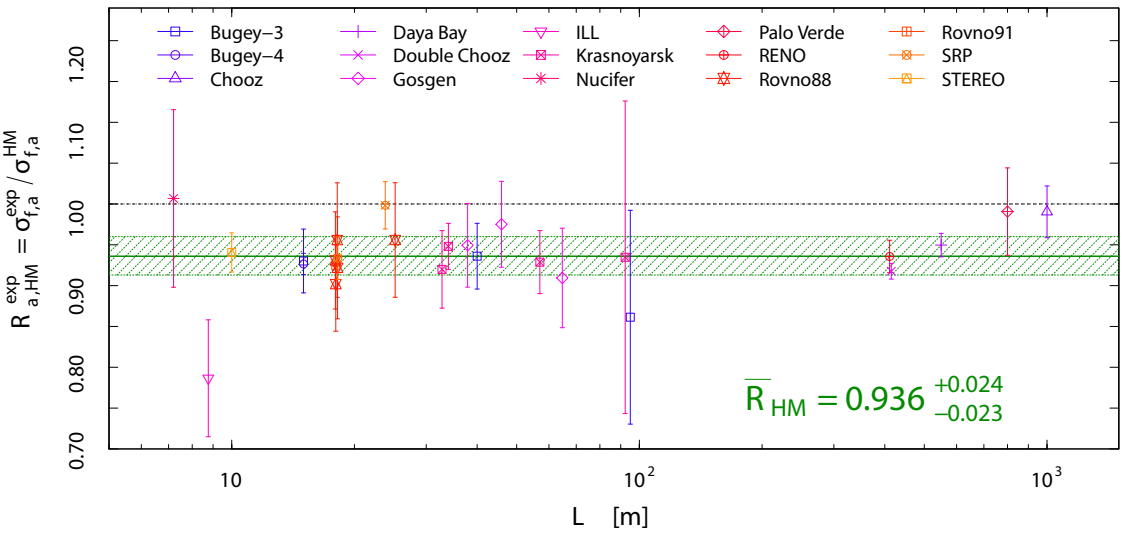


Figure 5.2: Ratio $R^{\text{exp}}_{a,HM}$ of measured and expected IBD yields for a selection of reactor experiments as a function of reactor-detector distance L for the Huber-Mueller model of reactor antineutrino spectra, taken from [94]. The horizontal green band shows the average ratio $\bar{R}_{HM}$ and its uncertainty, indicating a 2.5σ deviation from the expected.

KamLAND detector was uniquely placed with an average baseline of 180 km to its reactors of interest. This baseline allowed it to perform oscillation measurements previously limited to those of solar neutrino experiments, measuring the θ_{12} mixing angle. KamLAND provided the first evidence of neutrino oscillation from a terrestrial neutrino source and the first non-solar probe of θ_{12} and Δm^2_{21}. This was particularly important as there arose tension between solar and KamLAND measurements of $\sin^2\theta_{12}$ and Δm^2_{21}, shown in figure 5.3. This measurement has not been repeated by another experiment and KamLAND has now been repurposed as KamLAND-Zen, a neutrinoless double beta decay search.

5.5 Monitoring Nuclear Reactors

Due to their involvement in the production of nuclear weapons, the interest in reactor neutrinos is not limited to the academic. ^{239}Pu, created in nuclear reactors through the process described in equation 5.1, is a desirable fuel for nuclear weapons due to its high probability of induced fission but low rate of spontaneous fission. ^{240}Pu on the other hand is undesirable in nuclear weapons due to its tendency to emit neutrons which could induce unwanted chain reactions prior to a weapon's deployment. ^{240}Pu

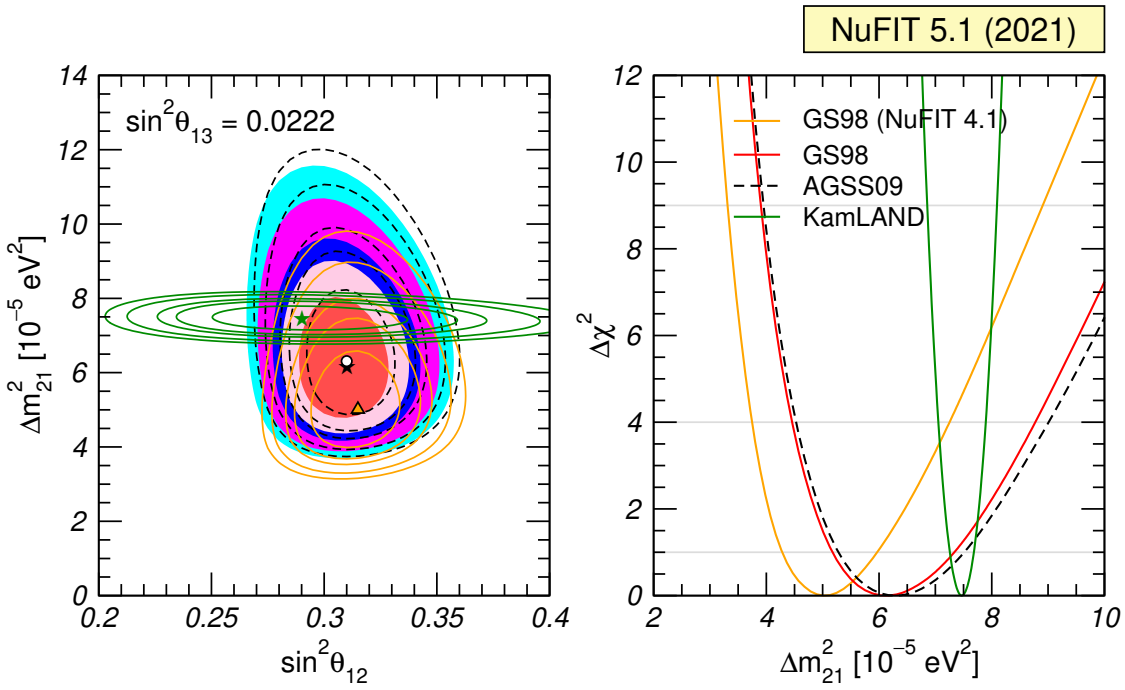


Figure 5.3: Left: Allowed parameter regions (at 1σ, 90%, 2σ, 99%, 3σ confidence limit for 2 degrees of freedom) for fixed $\theta_{13} = 8.6°$ from the analysis of KamLAND data (solid green contours with the best fit marked by a green star) and from the combined analysis of solar data for the GS98 model (full regions with the best fit marked by black star) and the AGSS09 model (dashed void contours with the best fit marked by a white dot). The orange contours are the results of a global analysis for the GS98 model using the previous SK4 neutrino data. Right: $\Delta\chi^2$ dependence on Δm^2_{21} for the same four analyses after marginalizing over θ_{12}. Taken from NuFit 5.1 (2021) [96].

is created when ^{239}Pu absorbs a neutron, and so to produce weapons-grade plutonium of 3-7% ^{240}Pu concentration a reactor must run long enough to produce ^{239}Pu from its ^{238}U fuel, but not so long that the fuel is significantly contaminated with ^{240}Pu. This typically happens on the timescale of a few days and results in a significantly different operating signature than that of power reactors that generally refuel once every 18-24 months and spend approximately 2 months ramping up to full power.

In 1970, the Treaty on the Non-Proliferation of Nuclear Weapons (NPT) came into effect, seeking to prevent countries without nuclear weapons from attaining them, encouraging the sharing of peaceful nuclear technology and promoting global nuclear disarmament [97]. Only a select few "nuclear-weapon states" are allowed access to nuclear weapons, and non-nuclear-weapon states are forbidden from producing nuclear weapons or weapons-grade fuel. Additionally, in January 2022,

the five 1970-era nuclear weapon states released a joint statement emphasising the importance of nuclear non-proliferation and disarmament [98].

To ensure non-nuclear-weapon states aren't breaching the NPT by producing weapons-grade plutonium in their plants, their activity must be publicly available to the IAEA, and frequent inspections must be permitted. If a state is secretly producing ^{239}Pu, however, it is difficult to detect. Almost all the products of fission can be shielded and if the reactor is in a remote location, disconnected from the power grid, there is no obvious external signature of its operation. Only the neutrinos cannot be shielded. It is therefore possible to monitor the flux of neutrinos to in turn monitor the fission rate of a reactor and use this to identify a refuel pattern inconsistent with power reactors. The majority of detector technologies in use today do not have the target material to provide an interaction rate great enough that such flux variations are detectable. The primary aim of the proposed Water Cherenkov Monitor for Antineutrinos (WATCHMAN) experiment is to prove that this approach to nuclear reactor monitoring is possible [99].

5.6 SK-Gd and Reactor Neutrinos

To date, only a single reactor neutrino interaction has been identified in a water Cherenkov detector, by the SNO+ experiment during its pure water phase [2]. Reactor neutrinos are very difficult to identify in a water Cherenkov detector due to the small amount of light produced at low energies a few MeV above the Cherenkov threshold. Despite this, the scalability of WC detectors can provide larger target materials than any other detector technology. With the vastly improved neutron sensitivity provided by gadolinium doping, SK is in a strong position to detect reactor neutrinos. The current situation in Japan in regard to its nuclear power stations also puts SK in a unique position where it may be the first experiment to detect the restart of a nuclear power plant at long range.

Since the completed loading of gadolinium into SK, there have been a total of nine Japanese reactor core restarts, listed in table 5.3. Cores 6 and 7 of the

Core	Restart Date [YYYY/MM/DD]	Interaction Rate [/d]	Total Rate [/d]
Base Rate	-	1.82	1.82
Sendai-1	2020/11/17	0.06	1.88
Genkai-3	2020/11/21	0.09	1.97
Sendai-2	2020/12/22	0.06	2.03
Ohi-4	2021/01/15	1.71	3.74
Takahama-3	2021/03/07	1.13	4.87
Genkai-4	2021/03/17	0.09	4.96
Takahama-4	2021/04/12	1.13	6.10
Mihama-3	2021/06/23	1.94	8.04
Ohi-3	2021/07/03	1.71	9.75

Table 5.3: Japanese reactor core restarts during the SK-Gd phase, ordered by date of restart. The expected increase in IBD events in the SK ID due to the restart is included, as well as the expected pre-restart rate, calculated from SKReact as detailed in section 5.7.

Kashiwazaki-Kariwa plant were slated for restart in late 2020; however, a security breach on the reactor premises has indefinitely delayed this.

The Ohi, Takahama and Mihama cores especially provided a significant increase in the reactor neutrino flux over the course of 2021. If an increase in IBD rate is detected to be consistent with the predicted increase in flux due to these restarts, it can be confidently said that SK-Gd has identified reactor neutrinos.

Detection of reactor neutrinos in SK-Gd has two primary motivations. First is the impact on non-proliferation of nuclear weaponry. While the activity of individual reactor cores would be lost due to a combination of the closeness in restart times and the long ramp up of power reactors, this would be the first detection of a change in reactor power output at long range, more than a few kilometres. This is the first step towards proving the effectiveness of large-scale, Gd-doped WC detectors as long range reactor monitors.

The second motivation is to measure θ_{13}. KamLAND is the only reactor experiment which performed an oscillation analysis with a baseline of $\sim$180 km. This was the only non-solar probe of θ_{13} and the results did not wholly agree with the solar fit. While beyond the scope of this work, if SK-Gd can identify reactor

neutrinos then this tension could be investigated further by performing a second measurement of θ_{13} using reactors at long range.

5.7 Simulating Reactor Neutrinos: SKReact

Unlike many sources of neutrinos (neutrino beams, a nearby supernova burst or the Sun, for example) which can be approximated as an almost point-like source, there are many individual sources of reactor neutrinos. Therefore, when considering an experiment like Super-Kamiokande where no one reactor is close enough for its flux to dominate, a mix of many baselines, reactor types and reactor operations must be considered when calculating the oscillated spectrum, in addition to converting this into what would be seen in the data.

To tackle the task of understanding what the reactor neutrino spectrum would look like in SK, I developed the SKReact software. The aim of this software is to simulate the production, oscillation, interaction and reconstruction of reactor neutrinos produced in all known power reactors worldwide, using the information provided by the PRIS (dating back to 2003), to support analyses involving reactor neutrinos in SK.

Interest in reactor neutrinos is not unique to a dedicated reactor analysis as the reactor neutrino spectrum overlaps with other interesting sources of electron antineutrinos that primarily interact via IBD, such as supernova relic neutrinos (SRN), geo-neutrinos and pre-supernova neutrinos. The relatively large flux of reactor neutrinos therefore acts as a significant background to these analyses, and so SKReact was designed with accessibility for other researchers in mind. Barna and Dye's "Web Application for Modelling Global Antineutrinos" [45] served as both inspiration and a form of validation for SKReact, and Baldoncini et al. provided the condensed information from the PRIS [100], as well as a reference throughout its development [101].

First, the produced spectrum (s^{-1}) is calculated for every reactor listed in the PRIS, using equation 5.5, assuming an LF of 100% and using the p_k/Q_k fractions applicable to each reactor type as listed in appendix C.1. The spectrum of the Ohi-3

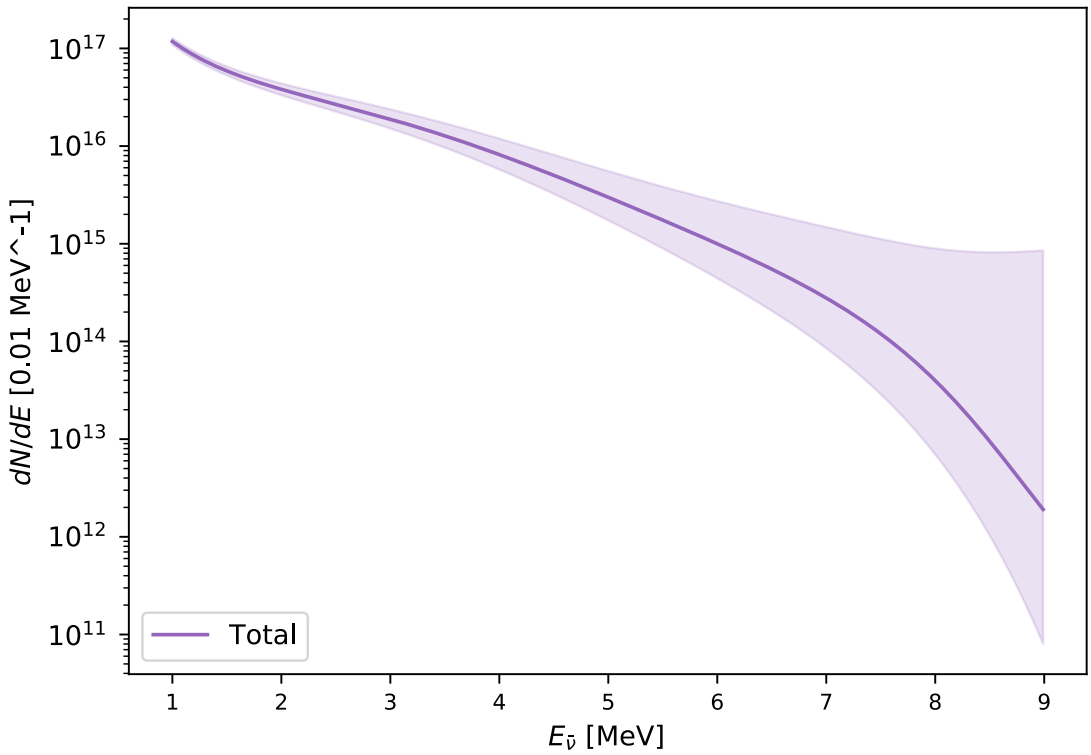


Figure 5.4: The simulated neutrino emission spectrum on production s^{-1} from the Ohi-3 PWR core.

pressurised water reactor (PWR) core is given as an example of this in figure 5.4. The energy range and number of points to generate is determined by the user (1-9 MeV, 800 points by default). This spectrum represents the neutrinos produced per second; however, of typical interest is the spectrum expected after some exposure time. Each spectrum is therefore scaled by the sum of monthly load factors in the period of interest (as selected by the user from the available PRIS information) multiplied by the number of seconds in the respective month. The integral of this spectrum provides the number of neutrinos produced by the reactor in the given period.

The survival probability after three-flavour oscillation of electron-antineutrinos of energy E at a distance of L is defined by:

$$P(\bar{\nu}_e \to \bar{\nu}_e) = 1 - \sin^2(2\theta_{12})\cos^4(\theta_{13})p_{21} - \sin^2(2\theta_{13})[\cos^2(\theta_{12})p_{31} + \sin^2(\theta_{12})p_{32}]$$

where

$$p_{ij} = \sin^2\left(1.267\frac{\Delta m^2_{ij} L[\mathrm{m}]}{E_{\bar{\nu}}[\mathrm{MeV}]}\right), \tag{5.6}$$

θ_{ij} are the PMNS matrix oscillation parameters and Δm^2_{ij} is the mass squared difference between the mass states [102]. The values of these parameters in SKReact are the global fit values as calculated by the NuFit group, from July 2019 (the most recent at the time SKReact was created) [103]; however, the user can redefine custom values.

The longitude and latitude of reactors in the PRIS is known, so the linear distance L to SK can be calculated through conversion from geodesic to geocentric coordinates, considering Earth's bulging at the equator. It is assumed that all reactors are at sea level as the altitude of the reactors is not necessarily known and the reactors relevant to SK are all based by the coast for easy access to water for cooling. L is calculated for all reactors, and each point on the spectrum is then scaled by $P(L/E_{\bar{\nu}})$ and $(4\pi L^2)^{-1}$ to produce the oscillated flux of reactor neutrinos at SK. Figure 5.5 shows the resultant spectrum s^{-1} for 2018 reactor information.

As discussed in section 2.4, inverse beta decay is the dominant interaction mode for electron-antineutrinos above the IBD threshold of $m_n + m_e - m_p \approx 1.806$ MeV up to a few tens of MeV. SKReact uses the Strumia and Vissani IBD cross section parameterisation to generate an interacted spectrum for all reactors. This is then scaled by the number of free protons (i.e. hydrogen atoms) in SK's ID ($\approx 2.14 \times 10^{33}$) to show the interacted spectrum in SK for a given exposure. Due to the kinematics of inverse beta decay, it is reasonable to approximate the outgoing positron spectrum as the neutrino spectrum offset by $\Delta = 1.293$ MeV.

From this interacted spectrum, SKReact can generate a file with initial positron vertex information which can be used as input into SKDetSim, SK's Geant3-based MC detector simulation software. For historical reasons, SKDetSim takes the file format produced by the NUANCE interaction generator as input [104], though NUANCE itself is no longer used.

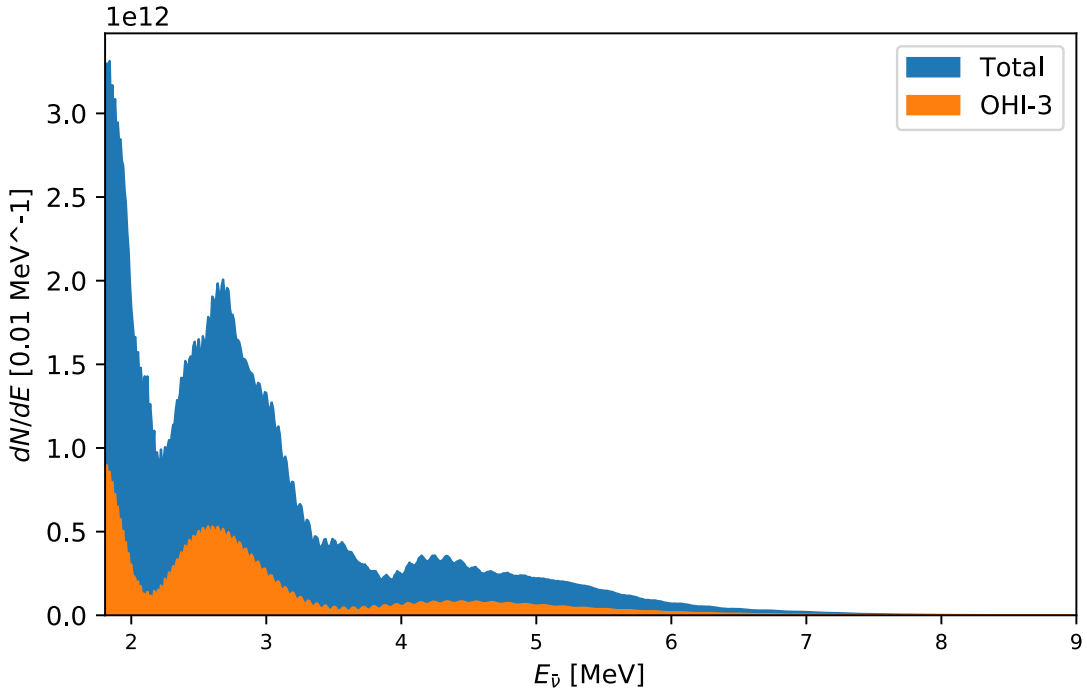


Figure 5.5: The simulated 2018 reactor spectrum after oscillation s^{-1} at SK, with the Ohi-3 PWR core highlighted.

It is often useful to modify reactor information, oscillation parameters and exposure periods, especially when simulating the restart of a currently-offline reactor. Regenerating an MC sample and passing it through the full simulation and reconstruction chain to evaluate the impact of these changes on the reconstructed spectrum becomes cumbersome when comparing between many different simulation parameters. To partially combat this, in addition to the true interacted spectrum, SKReact generates a "smeared" spectrum, indicative of what the reactor neutrino sample would look like after being reconstructed by BONSAI, discussed in 3.5. To achieve this, 90 MC samples of 10,000 positrons of a fixed energy were generated, with each sample increasing in energy from 1 to 9 MeV. The events were randomly distributed throughout SK's ID with random direction.

For each sample, events which would pass WIT's trigger requirements were selected, and the reconstructed energy spectrum of the resulting sample was fitted

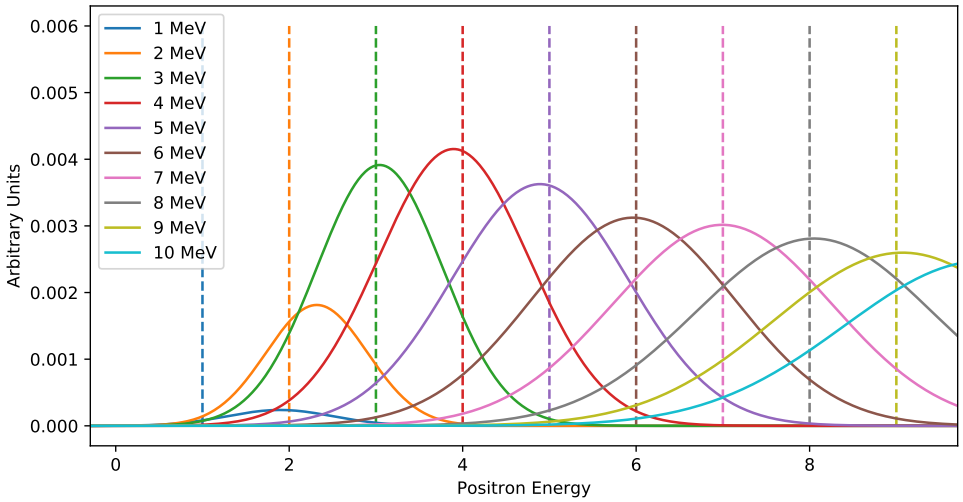


Figure 5.6: The Gaussian fit (solid lines) of BONSAI reconstructed energy for simulated positrons of energies from 1 to 10 MeV, area normalised to the efficiency of detection. Dotted lines at true energy given for reference.

with a Gaussian. The areas of these Gaussian fits were then scaled by the inverse of the sample size so that when a point on the interacted spectrum is multiplied by its corresponding Gaussian, it will be both smeared by the reconstruction resolution and effectively scaled by the efficiency of detection for that energy. Examples of these Gaussian fits are shown in figure 5.6 which demonstrates how BONSAI tends to reconstruct above the true energy in the particularly low energy regime. This is likely because the only events that reconstruct at very low energy are when an unusually large number of photons hit the PMTs, feigning a higher energy signal.

The Gaussian parameters are interpolated (if needed) to match the points in the interaction spectrum, and then sampled at a user-defined number of evenly-spaced points. These values are then put into matrix form, with each row containing a Gaussian fit/interpolation for a corresponding energy value in the interacted spectrum. The row vector of the interacted spectrum as a function of energy is then multiplied by this matrix, producing a smeared spectrum with the number of points matching the number of sample points of the Gaussian fits. A visualisation of this process is given in figure 5.7. The difference between the smeared and simulated neutrino interaction spectrum for the 2018 can be seen in figure 5.8. The smeared

spectrum simulates the reconstructed energy of the positrons so it is offset by Δ to demonstrate the incident neutrino energy.

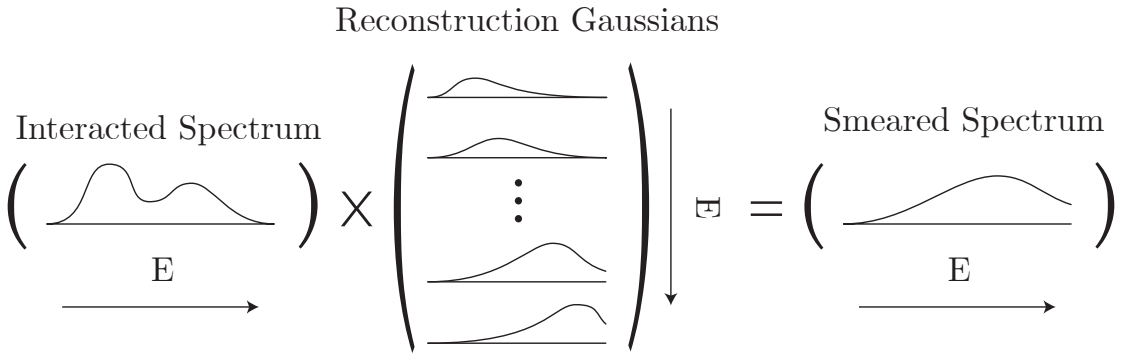


Figure 5.7: The process for converting interacted spectrum to reconstructed spectrum.

One advantage of generating a smeared spectrum directly from the calculated spectrum is that it can be used to measure oscillation parameters. By varying the parameters to minimise the χ^2 between the smeared and measured spectra, the best fit can be found. This χ^2 minimisation was implemented into SKReact partly in preparation for an oscillation analysis, but also as a simple validation of the smearing process by demonstrating that the measured oscillation parameters were consistent with those used as input to the smearing process. The current fitter crudely minimises using an iterative grid search of the fit parameter space and could be improved.

To increase SKReact's accessibility, each step of the process described here is displayed in a graphical user interface (GUI), with the option to vary and inspect many control parameters on the fly, shown in figure 5.9.

The exposure period selection (centre right) is accompanied by a plot of number of interactions in SK over the available data period to better visualise reactor activity. The spectrogram (top right) expands on this, showing the spectrum over time for the entire data period. A list of reactors contributing to the spectrum is also shown (top left), ordered by their integrated flux at SK over the exposure period. This allows for inspection and modification of a selected reactor's information (bottom left) as well as adding custom reactors.

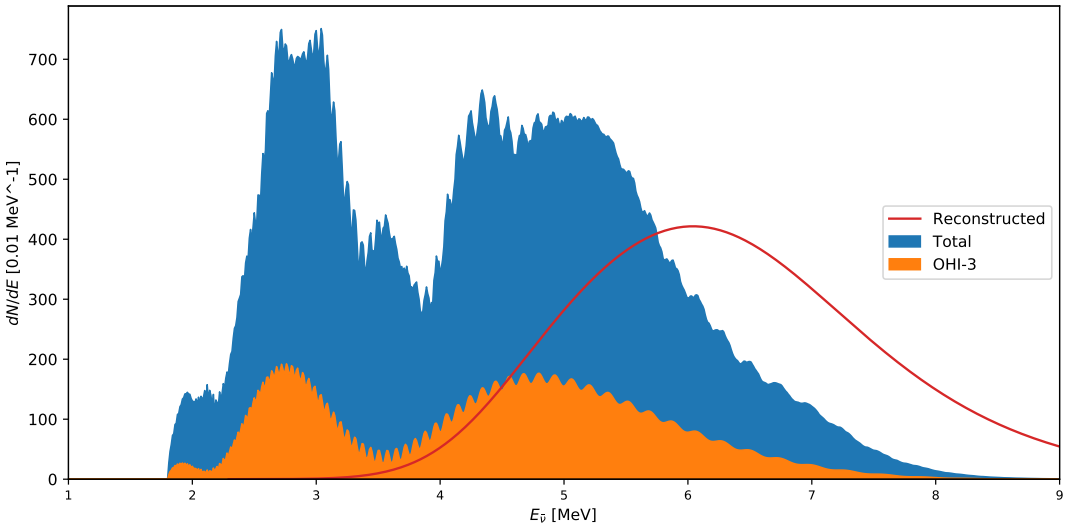


Figure 5.8: The interacted and smeared spectrum over the course of 2018, with the Ohi-3 core highlighted, simulated by SKReact.

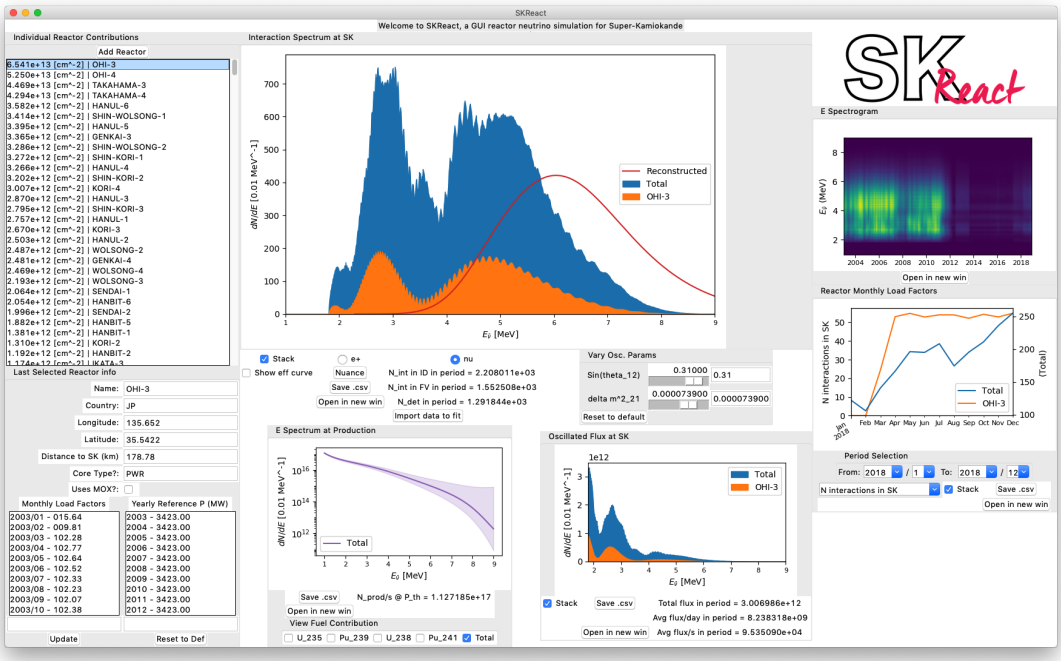


Figure 5.9: The SKReact GUI.

SKReact has been particularly useful in the development of a reactor sample selection and was used to calculate reactor background rates in SK's most recent DSNB study [43].

6
Restart Detection

Contents

6.1 Introduction and Objective 100
6.2 Sources of Background 105
6.3 Cut Optimisation with SPLASH 108
6.4 BDT Selection 117
6.5 Background Prediction and Data Availability 127
6.6 Results 136

6.1 Introduction and Objective

Following the first Gd loading of SK in 2020 there was a period of high event rates as the detector's water stabilised. This had recovered from the point of view of the WIT trigger rate by the end of 2020, opening up the opportunity to study reactor neutrinos in a Gd-loaded SK over the course of 2021. A caveat to this is that the SK computing systems that store the data and provide an analysis platform for SK researchers received a major upgrade in early 2022, updating its operating system and libraries. While this upgrade was largely met with no major issues, the WIT-specific offline reconstruction and reduction program was not trivially updated. As it must be manually initiated, it had been some time since the WIT reduction

had been processed. As of the writing of this thesis, uncut, fully-reconstructed WIT pair data is only available up to the 14th of October 2021.

While this would ideally provide around 10 months of data (December of 2020 is considered stable enough to include), upon analysing the 2021 WIT data for event rate stability it was discovered that between the 4th of March and the 31st of July there is some strange behaviour exhibited. A sudden, unexplained change in the trigger thresholds (specifically the hit counter and N18 requirements) resulted in an increase in the event rate. WIT experts recorded no manual changes to the WIT system at the time; however, there were issues in the power delivery to the Kamioka mine on the same day as the change. It is not clear how these issues could have caused a change in the compiled executables running on the WIT computers, but it is believed that this may be the cause. While it should be possible to retroactively counteract the changes to WIT by simply reapplying the tighter WIT trigger requirements as cuts to the offline data, doing so results in instead a drop in the event rate for this period relative to its surrounding periods. Applying loose quality cuts such as OVAQ or minimum energy requirements does not remove this drop, indicating a core issue with the WIT data.

This is confirmed when the individual file lengths are inspected, as calculated using the range of hardware triggers in each file. WIT data is partitioned into files which are usually cover up to 87.38 s of data as a result of the approach to processing data blocks issued by the DAQ. While the file lengths in this unstable period are generally matched to this cap, a relatively small number of files unexpectedly exhibit far larger lengths when calculated from the hardware triggers, as seen in figure 6.1, apparently not conforming to the cap. These large file lengths are not necessarily reflected in the file sizes, especially in the cases where the file length is extremely large. The majority of files are 87.38 s long, as expected; however, below this there is a large spread in the calculated file lengths as seen in figure 6.2, not generally observed in the rest of the data.

There are also a number of files of a significantly larger file size than the average, though this is observed elsewhere in the data. If the files with an unusually long

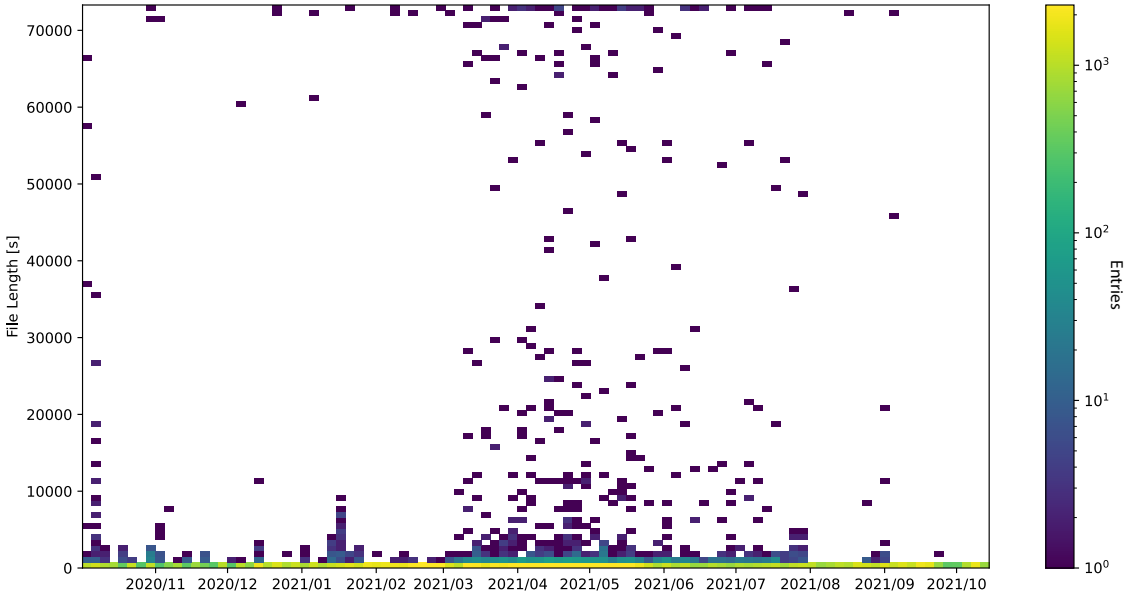


Figure 6.1: The length of the data files in the WIT pair data, calculated using the hardware trigger counter differences in each file.

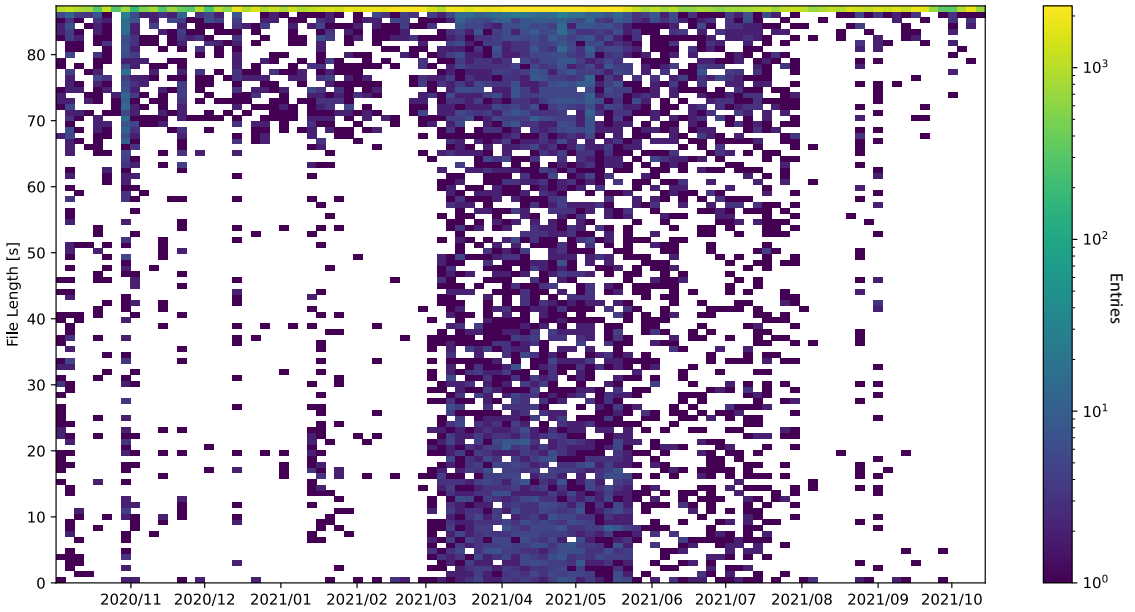


Figure 6.2: The length of the data files in the WIT pair data, calculated using the hardware trigger counter differences in each file, removing files of length greater than 87.38 s.

calculated length are removed, what remains in this period is an average file size over double that of before and after the period, as well as a far wider spread of file sizes, as seen in figure 6.3.

The beginning of this period was coincident with issues with the HV systems;

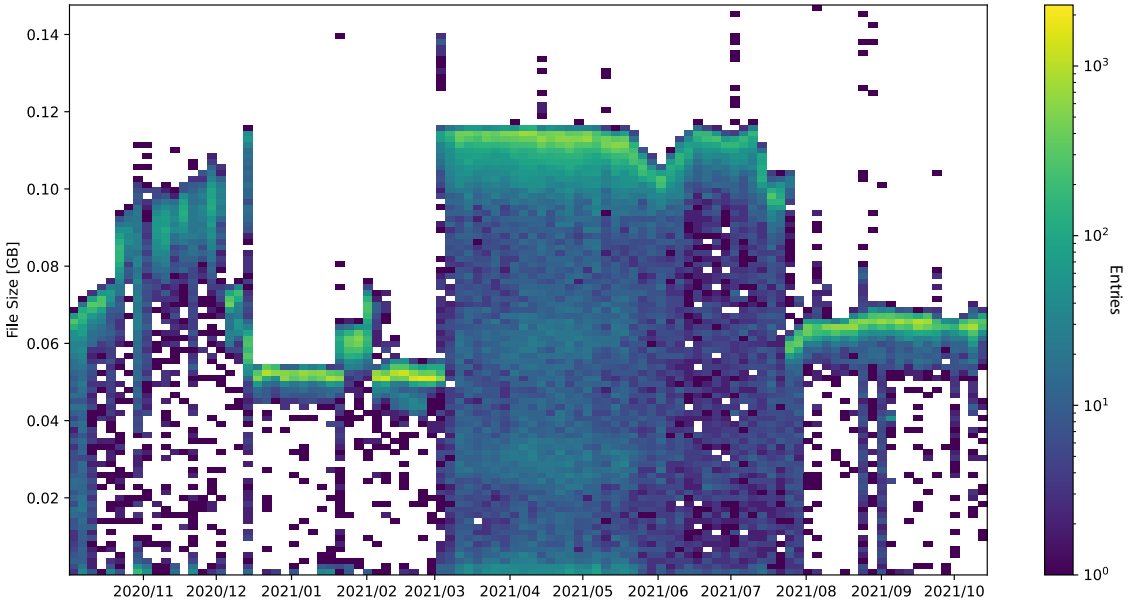


Figure 6.3: The size of the data files in the WIT pair data, removing files of length greater than 87.38 s.

however, this is not necessarily suspected as the cause, given that the WIT data seemed to return to normal following an unrelated update to the executables deployed on the WIT machines. It is not clear what the cause of the issues is, nor whether the data is recoverable through some form of re-processing. Due to storage requirements the raw data has since been deleted so if the core hardware information was somehow corrupt during or before the offline reconstruction step, the data cannot be recovered. It is currently unclear if this issue is reflected in the data which passed through the WIT singles data stream; however, use of this data would require new studies into the effect of the first reduction cuts on the sample selection discussed later in this chapter, developed before these issues were identified. To monitor the WIT data and prevent any such issues arising again, or at least identify them before the raw data is deleted, a number of plots of various event variables are now generated for each run. These plots are periodically inspected by SK collaborators within the low energy analysis group and runs that exhibit issues are flagged as requiring closer inspection.

Despite the issues with WIT, the data allows for comparison between the event rate at the start of 2021 with that at the end. This is a promising position to be

in as 2021 was a particularly active year in terms of Japanese reactor restarts. As mentioned in section 5.6, the SK-Gd phase has seen 9 such restarts, including two cores each of the nearby Ohi and Takahama plants following long maintenance shutdowns, as well as a core of the even-closer Mihama plant.

The restart information in table 5.3 is taken from public announcements; however, the monthly reactor power output is only available in the annual PRIS report, which is unpublished for 2021 data at the time of writing. From previous PRIS report data, it is understood that reactor cores typically take around 2 months to ramp up to full power. The time-dependent total interaction rate $I(t)$ can therefore be simulated by overlaying a series of sigmoid functions which approximate the restarts, in this case the logistic function given by

$$f(x) = \frac{1}{1 + e^{-x}} \tag{6.1}$$

which produces a sigmoid increasing from 0 to 1. If the ith reactor restart, with restart date t_i ramps up to its interaction rate contribution I_i over a period of Δt, $I(t)$ is given by

$$I(t) = \sum_i \frac{I_i}{1 + e^A} \tag{6.2}$$

where

$$A = \frac{-(t - t_i - \frac{\Delta t}{2})}{S}. \tag{6.3}$$

S scales the width of the sigmoid, i.e. the time it takes for the sigmoid function to increase from approximately 0 to its approximate maximum. For a pure sigmoid $f(x)$ (equation 6.1), this is taken to be $\Delta x = 10$ and so S is given by $\Delta t/10$. The $\Delta t/2$ in the exponent's numerator offsets the ramp up to begin on the restart date rather than being centred on it.

The continuous result of this approximation is given in figure 6.4, which can be averaged over a given period to give a predicted rate for each bin in whatever

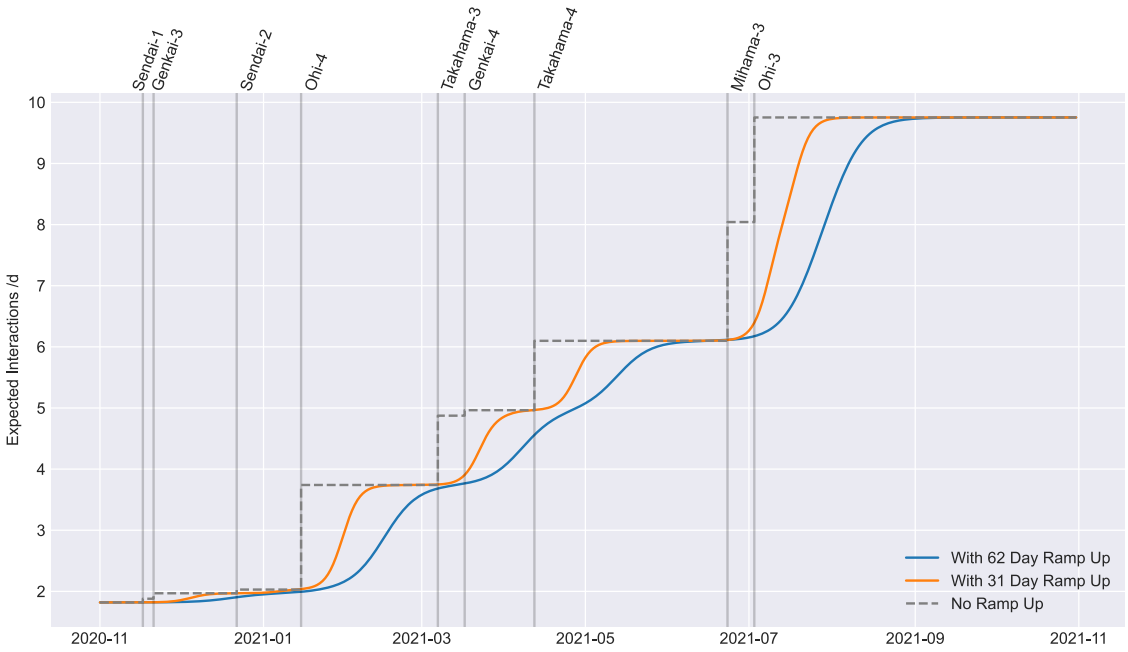


Figure 6.4: The simulated reactor neutrino interaction rate from late 2020 to late 2021, with the ramp up simulated using sigmoid functions. A ramp up of 62 days is used in this analysis.

scheme the data is presented in. A ramp up time of 62 days is used for this analysis, matching that of previous restarts reported in the PRIS.

Given the issues with the WIT data, two data periods are defined: the "pre" period (i.e. pre-2021 reactor restarts), beginning on the 1st of December 2020 and ending on the 4th of March 2021; and the "post" period, from the 31st of July to the 14th of October 2021. The aim of this work is to identify a significant increase in event rate in the post period when compared to the pre, consistent with the reactor neutrino interaction rate increase. With this in mind, no full reduction was applied to data beyond June 2021 during the development of a sample selection, so as to avoid biasing the selection to detect this feature.

6.2 Sources of Background

The main backgrounds to an IBD search on an energy scale close to the Cherenkov threshold search are radioactive decays, and dark noise hits registering as an event, both of which may happen in coincidence with other background events. For radioactive decays, the main contributor in the low energy regime is radon which

Chain	Isotope	Criterion [mBq/kg]	Physics target
^{238}U	^{238}U	<5	SRN
	^{226}Ra	<0.5	Solar
^{232}Th	^{232}Th	<0.05	Solar
	^{228}Ra	<0.05	Solar
^{235}U	^{235}U	<30	Solar
	^{227}Ac/^{227}Th	<30	Solar

Table 6.1: Criteria of radioactive impurities in $Gd_2(SO_4)_3 \cdot 8\,H_2O$ powder and the physics goals driving the requirements [1].

emanates from the walls of the detector, particularly from the PMT glass. Radon is carefully controlled and well understood in SK, and the fiducial volume cuts exist primarily to reduce this background [105].

In preparation for SK-Gd, extensive studies were performed to quantify the radioactive contamination of the gadolinium that would be introduced into the detector. This was an iterative process: receiving samples of powder from suppliers, measuring their radioactivity with high-purity germanium (HPGe) detectors, and feeding back to the suppliers to identify a source of gadolinium sulphate which conformed to a stringent set of criteria. The broad requirement for the cleanliness of the gadolinium powder was that radioactive background rates for SK low energy analyses should less than double. These criteria are given in table 6.1.

Neutrons are also produced in the process $^{18}O(\alpha, n)^{21}Ne^*$, and its equivalent with ^{17}O, triggered by αs emitted by ^{235}U. Spontaneous fission of ^{235}U is another source of neutrons. These neutrons may be misidentified in coincidence as an IBD pair.

As discussed in section 3.6, the dark noise in SK can be simulated in Monte-Carlo, but for particularly low energy events where the number of dark noise hits is comparable to those from signal, it is preferable to use real data. This is also true when considering background events entirely composed of dark noise hits. The dark noise rate for the individual PMTs in SK can fluctuate, especially during changes to the detector as with gadolinium loading. Increases to the dark noise rate increases the flat event rate due to the larger number of hits potentially triggering the DAQ and being saved by WIT. An increase in the flat event rate therefore increases the

pair rate, which can be quantified by monitoring the event rate through WIT itself, an advantage to its data collection independent of SK's main triggers.

In this work, backgrounds were not simulated and instead the events in pre-reactor restart WIT data were collectively treated as background, with a reactor neutrino "contamination" that can be quantified using SKReact. The signal rate for a given window of data in WIT is orders of magnitude lower than the total event rate and so cut optimisation and training of machine learning methods are not affected by potential reactor neutrino events. This allowed for collective measurements of all the aforementioned sources of background as an overall rate, including accidental coincidences of dark noise events.

The spallation rate was not explicitly quantified. Spallation is a process where high energy cosmic muons produce fast neutrons, and unstable nuclei which decay. The majority of the neutron interactions and nuclear decays occur a short time after the muon passes through the detector, short enough compared to the cosmic muon rate in SK that all events in some time window after the muon can be rejected with little dead time introduced into the data. This was not implemented in this analysis as the BDT selection alone (introduced in section 6.4) was found to be sufficient in reducing the collective background rate to acceptable levels. Gains may be made, however, by introducing a dedicated spallation cut, as the BDT does not currently consider the time since the most recent cosmic muon event, nor an event's proximity to the muon's path.

There also exists irreducible background in the form of geo-neutrinos. These are electron antineutrinos of a few MeV; however, as discussed in section 2.3.6, the energy of geo-neutrinos is limited to a little over 3 MeV, below the reactor neutrino peak. At Kamioka, the integrated flux is also far lower than that of reactor neutrinos which combines with their low energy to remove them as a concern for a reactor search. Pre-SN neutrinos would also form an irreducible background; however, these are only detectable in the final hours or days before a nearby supernova, an event which has not occurred since 1987. Were a supernova to occur in the period of

interest of a future reactor study, the pre-SN neutrinos would be trivially removed by correlating their timing with the supernova burst.

Ideally, the background rate would be measured using pre-reactor restart data, and this would act as a fixed baseline and a significant increase in flux above this baseline would indicate a reactor restart. While the sources of background are theoretically stable with time, the radioactive contamination of the water and dark noise rate of the PMTs can fluctuate, and it is important to ensure any increases above the pre-restart baseline are not due to such fluctuations. The approach to quantifying this included is described in section 6.5, utilising the BDT selection introduced in section 6.4.

6.3 Cut Optimisation with SPLASH

The MC sample used to develop the selection was formed by first simulating positron events throughout the ID volume, with an energy spectrum which approximates the predicted reactor spectrum at the end of 2021. At the time of MC generation and initial sample selection development, the 2019 PRIS report was the latest that was available. The energy spectrum was therefore simulated to match that of the 2019 reactors power outputs, modified so that the reactors described in table 5.3 were set to 100% load factor for the year if not already described as such in the 2019 PRIS report. The only significant changes to the Japanese reactor situation between 2019 and 2020 was that the Ohi and Takahama plants were shut down. In approximating the 2021 spectrum, these plants are simulated to be running at full power anyway, and so the 2019 spectrum serves as an equally appropriate baseline as the 2020 spectrum.

An equal number of delayed events were simulated alongside the positron MC sample, each uniquely matched to a positron vertex. These were equally split between a single 2.2 MeV photon and an 8 MeV gamma cascade, simulating hydrogen and gadolinium capture respectively. The number and energy distribution of the gammas in the Gd capture sample was pre-simulated using Geant4 by an SK collaborator. Given that the prompt and delayed events aren't truly generated

together and the neutron capture isn't itself simulated, once paired, the Δt between the events is generated using a random number weighted by the neutron capture time formula (exponential decay) with a neutron capture time of 203.7 μs for hydrogen capture [106], and 30 μs for gadolinium.

A 24 h run of real WIT data was used as the background to optimise against, and the signal rate was taken to be 9.75 d^{-1}, as it is predicted to be at the end of 2021. The BG data run began on the 18th of February 2021, before any major reactor restarts; however, there is still an expected reactor interaction rate of 1.82 d^{-1} due to more distant reactors. To calculate the expected background rate for a given cut or set of cuts, the MC and BG would be reduced in parallel. An efficiency for the MC can then be calculated, which is then used to find the expected number of reactor neutrino events appearing in the BG sample post-cut. This number is then subtracted off of the number of BG events surviving the cuts and this offset BG rate is used to calculate the BG efficiency without any (expected) signal contamination. This introduces a dependency on the reactor interaction rate prediction to the BG efficiency; however, without this offset, the BG efficiency would be overestimated.

The development of a reactor neutrino sample selection began with the cuts used in the SK-IV solar analysis [21]. These were chosen as a starting point as the solar analysis is particularly mature in SK and as such the variables used within it are well established as powerful discriminators in event identification. There is little information in low energy events in SK given the small number of hits and the general simplicity of the detector and there is therefore little scope to innovate on the established low energy cuts. The variables used in the SK-IV solar analysis sample selection are:

- Vertex position - A fiducial volume is defined by the minimum distance from the PMT plane to the event vertex, 2 m in most analyses.
- Energy - BONSAI reconstructed energy of the event.

- OVAQ - ($g^2 - \text{DirKS}^2$), where g is the BONSAI goodness and DirKS is a directional KS test of the PMT hit pattern. A general event quality parameter described in more detail in section 3.5.

- effWall - Distance to the detector wall in the reverse direction of the event, powerful at identifying radioactive events originating close to the PMT plane but emanating away from it. BONSAI has reduced vertex reconstruction accuracy along the direction of the event, so such events often reconstruct within the fiducial volume.

- Multiple Scattering Goodness (MSG) - Measure of directionality confidence, based on fitting multiple direction vectors assuming a 41° Cherenkov cone and evaluating their general alignment with one another.

- BONSAI Pattern Likelihood (PatLik) - A χ^2 goodness of fit for angles between pairs of hits to the expected opening angle of the Cherenkov ring.

Given the different energy regime and the addition of the pair selection, these cuts would need to be re-optimised to the reactor neutrino MC sample, separately for the prompt and delayed events, and the coincidence variables Δt and Δr, the time and distance between the prompt and delayed vertices respectively, would need to be added and optimised. The initial values of the cuts taken from the solar analysis were treated purely as a starting point to optimise from. To facilitate this, I developed the SPLASH framework.

SPLASH arose from the lack of a column-wise approach to cut-based analyses in ROOT, the C++ analysis framework most commonly used in the field of particle physics [107]. ROOT's approach to analysis is geared towards looping through a file event-by-event, saving events which pass cuts to a new structure and rejecting those which don't. While this is a reasonable approach to applying cuts, it becomes cumbersome to optimise an individual cut and especially a set of cuts. Such a process would require repeatedly looping through the dataset, re-applying the cut and evaluating its efficiencies with each iteration.

An alternative approach instead structures the data in a table-like Dataframe of rows and columns, with each row representing an event and the columns the event variables. In this format, events can be selected as in a database, with a single query returning the Dataframe which contain events which satisfy that query. To an extent, the per-event method is utilised in the backend of the code in this approach; however, a range of optimisations ensure far more efficient data extraction. While Dataframe structures have recently been introduced to ROOT's libraries, they are in their infancy and are therefore less well documented and established than other Dataframe libraries. They are also not present in the version of ROOT used in SK.

SPLASH at its core utilises pandas [108], a Python-based Dataframe library leveraging the NumPy [109] library's CPython arrays to provide fast data manipulation. To minimise its dependencies, SPLASH uses the Uproot [110] library to load ROOT files without requiring a local ROOT installation, before placing them into a pandas Dataframe container.

This Dataframe approach allows for easy per-cut optimisation as the cut can be repeatedly applied to the data effectively as a query, and the number of events which pass can then be counted. MC and background samples can be separately cut, and the number of each that survives, appropriately scaled to match real-world rates, can be passed to a function which acts as a figure of merit (FOM) for the cut. In SPLASH, this FOM takes the form

$$\text{FOM} = \frac{S}{\sqrt{S + BG}} \tag{6.4}$$

which, when maximised, provides the most certain measurement of S, though alternative FOMs can be used. In future work, it may be preferable to use FOM $= S/\sqrt{BG}$ to maximise the statistical difference between the signal S and background BG rates, depending on the intention of the measurement. The negative of this FOM is minimised using the SciPy [111] package's `minimize` function.

Optimising the established cuts for the reactor selection did not provide a significant signal rate above the background rate. The optimised selection criteria are given below. For each, the distribution is given for prompt and delayed events

from the 2021-spectrum MC pairs used in optimisation, both compared to prompt and delayed events for pairs identified in WIT data. Both have passed through the WIT online processing (or an analogue of it, in the case of the MC). Each distribution is given without any other cuts applied so as to better demonstrate the overall distribution for each; the statistics of the BG distribution quickly degrades as cuts are applied, as intended.

- Fiducial Volume

  Prompt $> 2\,\mathrm{m}$

  Delay $> 2\,\mathrm{m}$

  The majority of the background within the fiducial volume tends to reconstruct more centrally in the detector, most likely due the approximate isotropy of dark noise. This tendency is seen in figures 6.5 and 6.6. Optimisation beyond the standard SK FV of 2 m did not provide a significant improvement to the FOM.

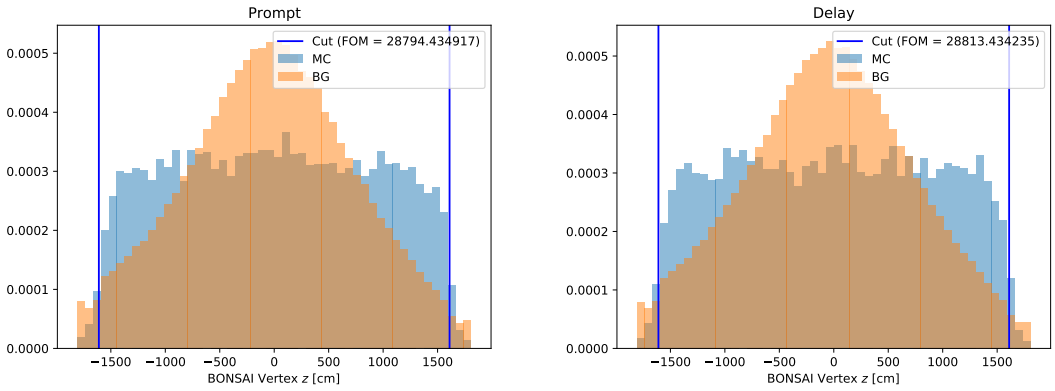


Figure 6.5: Vertex z as reconstructed by BONSAI for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data.

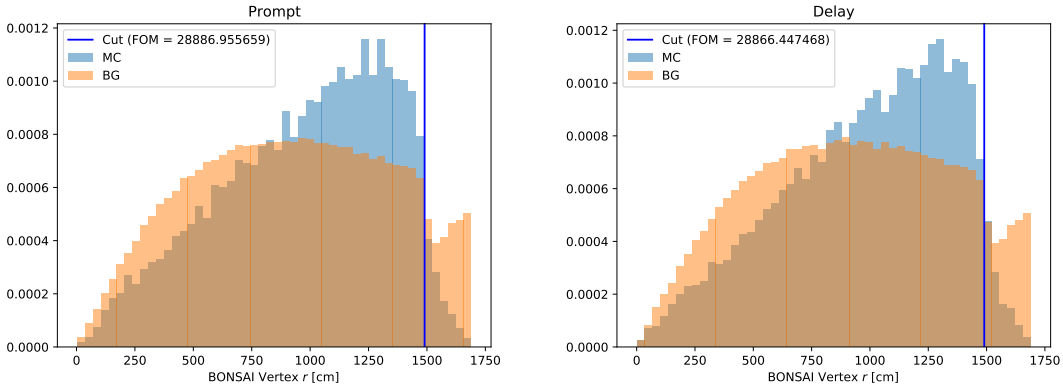


Figure 6.6: Vertex r as reconstructed by BONSAI for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data.

- Energy.

  Prompt $> 5.2\,\text{MeV}$

  Delay $> 5.3\,\text{MeV}$

  Shown in figure 6.7, the energy cut is the most aggressive, as it is the most significant measure of the quality of an event and has generally good separation between signal and background. This cut removes the most BG by far, leaving only the higher quality events to be removed by the more fine quality cuts.

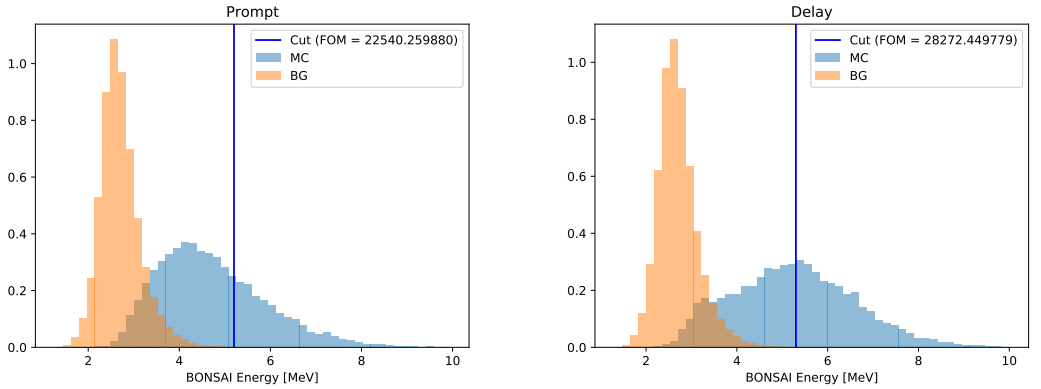


Figure 6.7: BONSAI reconstructed energy for 2021 reactor spectrum prompt positrons, and delayed neutron captures corresponding to 0.02% Gd concentration, both simulated (MC). BG is raw pre-reactor restart WIT data.

- OVAQ

  Prompt > 0.22

  Delay > 0.20

  The powerful background discrimination of OVAQ can be seen in figure 6.8, leading to another somewhat aggressive cut.

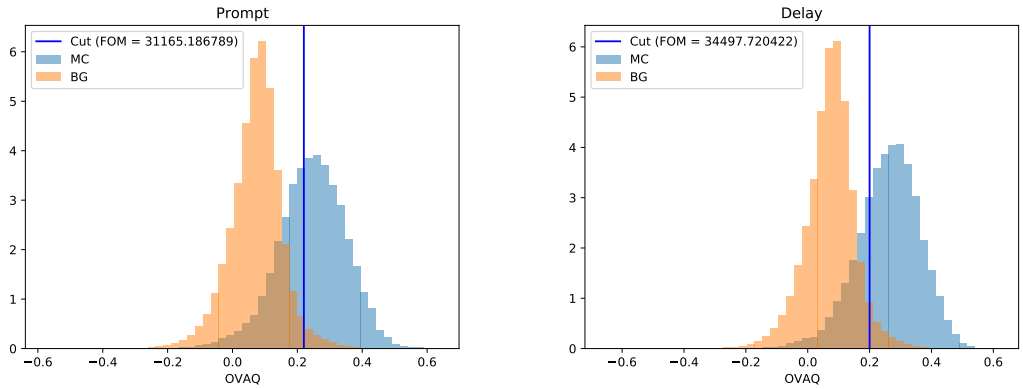


Figure 6.8: OVAQ goodness variable for simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data.

- effWall

  Prompt $> 4.3\,\mathrm{m}$

  Delay $> 4.4\,\mathrm{m}$

  EffWall shows some separation; however, the chosen FOM results in a fairly conservative cut, shown in 6.9.

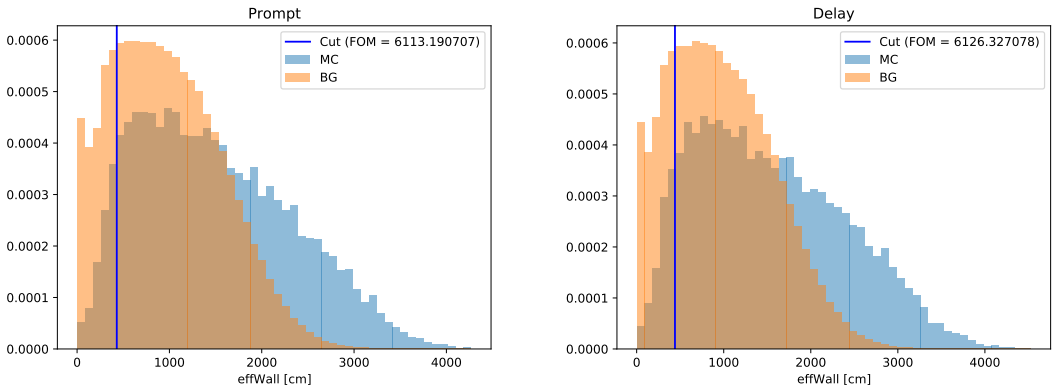


Figure 6.9: EffWall for simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data.

- $\Delta r < 1.52\,\mathrm{m}$

The Δr distributions shown in figure 6.10 demonstrate the power of the coincidence requirement in a gadolinium-loaded SK as there is clear separation between the MC and BG, a fact which the optimisation capitalises on with a particularly harsh cut.

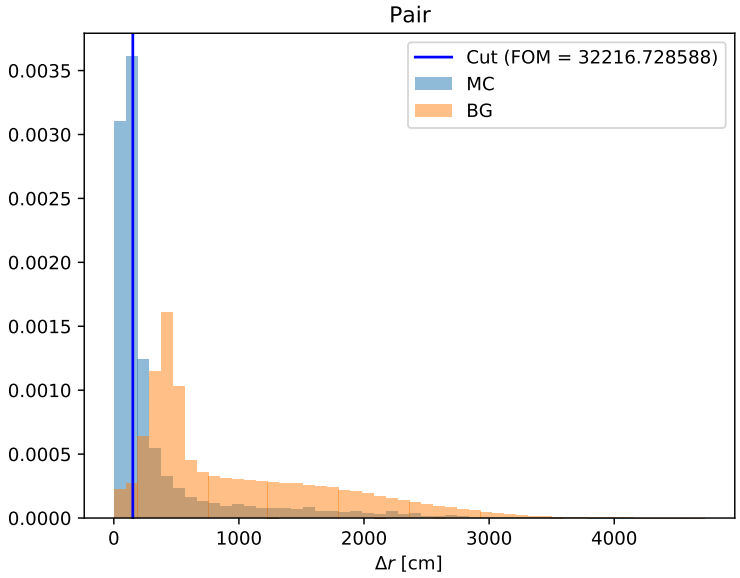


Figure 6.10: BONSAI reconstructed distance between simulated 2021 reactor spectrum prompt, and delays corresponding to 0.02% Gd concentration (MC). BG is raw pre-reactor restart WIT data.

- $\Delta t < 92\,\mu$s

Though less drastic a separation than Δr, the different Δt distributions for MC and BG can be seen in figure 6.11.

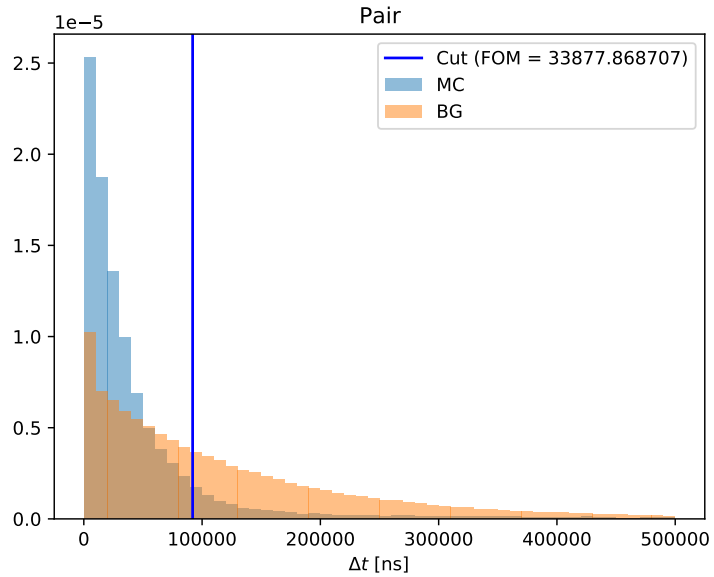


Figure 6.11: Simulated neutron capture times corresponding to 0.02% Gd concentration (MC). BG is time between prompt and delayed events in raw pre-reactor restart WIT data, identified by the WIT coincidence trigger.

Optimisation removed the MSG cut and PatLik was removed as, upon inspection, it was not clear that it was calculated correctly for MC and is often not calculated for very low energy events regardless, as a clear Cherenkov ring cannot be identified. The BONSAI energy cut was especially optimised to provide a particularly harsh rejection in an attempt to reduce the extremely high background rate. This cut alone removes around 53% of fiducial Gd captures to attain a background efficiency of just 0.03%.

The MC efficiency of the optimised cuts is calculated to be 1.5%, resulting in a signal rate of $0.14\,\mathrm{d}^{-1}$, while the BG rate was reduced to $36.0\,\mathrm{d}^{-1}$. The livetime in the data periods is not sufficient to provide a statistically significant difference between the two due to reactor restarts for this selection. In future work, this optimisation could be revisited, perhaps utilising an alternative FOM to focus on exclusively detecting a significant signal, and including a spallation cut.

6.4 BDT Selection

Manual optimisation of individual, independent cuts did not provide selection criteria that could significantly isolate reactor neutrino events. Increasing the complexity of such cuts to interweave different variables is possible, but it is a difficult and arguably unintuitive approach to such a sample selection. The application of machine learning (ML) techniques offers an alternative solution. Machine learning in general seeks to improve upon traditional multivariate analyses by leveraging the power of modern computing to generate complex models which can describe input data, be it through discrete classification or continuous output variables. Such complexity comes with the risk of over-optimising the model to match the data used to generate it, a phenomenon known as overfitting in the ML world. This is the primary concern when transferring from cut-based methods to ML methods; however, it is one which can be avoided with careful consideration, especially with the constant advancements in various ML implementations. Thanks in part to its ability to provide powerful models that are simple to implement, with the payoff of some model dependency, ML is used throughout particle physics and indeed the world of data analysis in general.

While cut optimisation was the original goal of SPLASH, a secondary focus was to provide a platform for ML tools which may be used in SK. The front end of ML libraries handling Boosted Decision Trees (BDTs), artificial neural networks (ANNs) and convolutional neural networks (CNNs), for example, are typically written in Python. It is therefore common for particle physics analyses to either use PyROOT, an officially-supported Python wrapper for the ROOT libraries, or produce intermediate files which can be read by Python programs when working with ML. Uproot has, however, more recently gained traction in this regard thanks to its independence on ROOT itself and comparable performance to native ROOT. SPLASH seeks to streamline the ML data preparation process, by providing the tools to load data from ROOT files and format the data ready for use in ML libraries without being bound to a specific analysis. Some work was done to implement the

framework for 2D and 3D CNN-based event classification in SPLASH; however, this analysis focused on the use of BDTs.

The pre-SN alarm recently implemented in SK uses BDTs to identify an increase in IBD rates consistent with a star progressing through its silicon-burning phase in WIT data [30]. To minimise the delay between a detectable rate increase and the issuing of an alarm, two BDTs are used: one running on the WIT system itself, using only variables calculated during online processing to intially reduce the data; and one which performs the final selection after the initial, partially-reduced sample has been fully reconstructed. A significant background to this alarm is the increase in rate due to reactor restarts given the similar energy regime. Though the predicted rates would not be significant enough to explicitly identify a restart, this begins to demonstrate the potential a BDT selection may have to identify a significant number of reactor neutrinos and helps to motivate their use in this study.

BDTs are a method of multivariate analysis, removing the need to optimise a number of cuts by using gradient boosting to assign a score to events. This score represents the BDT's predicted likelihood that the event is one of a number of classes. In the case of particle physics analysis this often means classifying an event as either signal or background. BDTs are in some ways comparable to neural networks given that they take a large, labelled dataset for training and seek to model the data to provide a score; however, the models take significantly different forms. While the more complex NNs can outperform BDTs in similarly complex, large scale datasets, BDTs are frequently used in multivariate analysis thanks to their quick implementation, minimal tuning requirements and strength in classifying less complex data where the input variables for a given row (event) is relatively small. For example, the Water Cherenkov Machine Learning (WatChMaL) group [112] developed an implementation of a ResNet CNN [113] to classify and reconstruct events in water Cherenkov detectors (the target network architecture for the 2D CNN formatter implemented into SPLASH). This CNN takes every single PMT hit's position and time as input for each event, intending to model the hit patterns.

To take the same approach with a BDT would require an unfeasibly complex model which would realistically never achieve the same accuracy as the CNN.

BDTs calculate an event's score using an ensemble of decision trees. Trees are a series of queries which subclassify a dataset into two or more "leaves". The depth of a tree is the maximum number of queries the dataset could be subject to, determining the maximum number of leaves in the tree (not all branches will necessarily reach the maximum depth). Events are assigned a per-tree score based on the importance of the leaf to which it is assigned. These scores are not necessarily symmetrical; being assigned to one leaf may provide a positive contribution to the tree score while the other won't necessarily negatively affect the event's score significantly. An example would be that events above a certain energy threshold have a good chance of being signal, but being below the threshold is not necessarily a good indication that the event isn't signal. The final BDT score for a given event is calculated by summing the score assigned by every leaf of every tree in the ensemble.

The reason these decision tree ensembles are considered "boosted" is because trees are generated iteratively, and the structure of each tree is influenced (or boosted) by the performance of the previous tree. The act of boosting defines a loss function that acts as an objective function: the metric for how well the algorithm predicted the class of an event. One example of such a function is the negative log likelihood function, given, for a prediction y_i of an event i of the class $\hat{y}_i$, by

$$L(\theta) = \sum_i [y_i \ln(1 + e^{-\hat{y}_i}) + (1 - y_i) \ln(1 + e^{\hat{y}i})] \tag{6.5}$$

where θ is the model's parameter space.

Beginning with a single tree, the BDT makes predictions for a subset of a large training dataset and the loss is calculated. The full tree ensemble is built one-by-one by iterating through subsets of the training set, the performance of the latest iteration determining the structure of the new tree. This is done through gradient descent, where the derivative of loss decides the direction in which the parameter space should be shifted to approach a minimum. A learning rate is defined, which determines the extent that each iteration moves the parameter space towards the

minimum. A lower learning rate avoids skipping over the minimum with each step, at the expense of a longer training time needed to reach the minimum.

An implementation of the XGBoost [114] BDT classifier algorithm was used in this work's sample selection. This algorithm is widely used in the world of multivariate analysis thanks to its portability and optimisations, both in the logical approach as well as in its hardware utilisation. XGBoost takes a generalised approach to tree ensemble optimisation to support a range of loss functions and in this work the negative log likelihood loss function was used.

In addition to the loss function, XGBoost introduces a regularisation term which limits the complexity of the model, linearly combining the two to form the objective function by which the model is judged. BDTs which do not include this term are prone to overfitting, a common issue in the world of machine learning whereby a model becomes overly complex during training. Such a model can predict the training data with a high degree of accuracy, but will struggle to classify new data which doesn't very closely match the training dataset.

To further reduce overfitting and to optimise the trees, XGBoost removes suboptimal leaves in a process known as pruning. After building a tree up to the pre-defined maximum depth, XGBoost then iteratively works upwards through the tree, evaluating whether each leaf's performance is good enough to warrant its existence. There are a number of other steps taken to optimise the trees and reduce overfitting in addition to this, as described in detail in [114].

The XGBoost model used in this work was trained using the same BG and MC sample used to optimise cuts as described in section 6.3, with $\sim$12 k MC events and $\sim$7 M BG events in the training set. This is a comparatively low number of MC examples to BG and would ideally be far higher; however, the issues with the WIT offline reconstruction in SK's new computing system prevented the generation of a large MC sample with the full set of offline reconstruction variables. A smaller 2021 MC sample generated prior to the upgrade was therefore used. An initial BDT using $\sim$1 M MC events matching the 2019 reactor spectrum was trained prior to the BDT used in this work for some initial investigations. The performance between this and

the 2021-spectrum BDT was comparable when cross-validated with a 2021-spectrum MC sample containing 10 k event pairs set aside for testing and so the newer BDT was used. The same 2021-spectrum testing sample was used to validate the BDT's effectiveness. Using a sample independent to the training data in such a way is commonplace in ML work as it verifies that the model is not overtrained if it can perform similarly when evaluating either dataset. If a BDT provides a classification accuracy far higher for the training set than for an independent sample, there is a chance the model is overtrained. Such a disparity was not observed in this work.

No strict conditions were placed on what variables were to be included in the BDT and the relatively small pool of information to pull from meant that initially all available variables from the offline reconstruction were included, both BONSAI and ClusFit, as well as those calculated after vertex reconstruction (OVAQ, DirKS, EffWall, MSG). PatLik was not included for the same reason it was removed from the traditional cut optimisation: it was not clear it was being calculated correctly for the MC and it is often not able to be calculated for low energy events regardless. The variables used in the traditional cut selection, given in section 6.3, composed the majority of the input variables, though the STORE goodness and N18 online variables used in the WIT system were included. Additionally, two more variables not frequently used in SK low energy analysis were added: fromWall, the closest distance between the event vertex and the PMT plane; and ClusFit goodness, the goodness of the ClusFit vertex, calculated similarly to the BONSAI goodness described in section 3.5. These aren't used in most analyses as they don't offer improved background rejection in most low energy analyses. EffWall is more effective than FromWall in reducing the most common radioactive backgrounds, and the BONSAI goodness is more relevant than ClusFit. When considering the input into a BDT, however, especially when investigating events with such little information, it is wise to include all available variables in case there is a correlation between seemingly irrelevant properties which has not been otherwise identified.

The importance of a variable is a metric used for evaluating whether a variable contributes significantly to a tree's decision. It is calculated for a single decision

tree by summing the classification performance of all splits in the tree which depend on the variable. This score is then averaged across all the trees in the ensemble. The XGBoost algorithm is sufficiently efficient that trimming which variables are included for the sake of performance was generally not required, though the BONSAI and Clus-Fit DirKS variables were removed due to being significantly less important in initial testing on the 2021 sample. OVAQ (which depends on the BONSAI DirKS value) is still included.

The BDT used in the pre-SN alarm sample selection includes a set of PMT hit isotropy parameters [31]. These are based on the Legendre polynomials of the angular separation between PMT hits, denoted β_{1-6} for each order of the polynomial. A BDT was trained including these variables (using and evaluated on 2019 reactor spectrum MC); however, all but the first order variable was found to have the lowest weight of all the inputs, not counting DirKS, as shown in the feature importance plot in 6.12. Given that these β variables are computationally intensive to calculate and must be manually added after full offline reconstruction, they were deemed unnecessary and were omitted from the final BDT.

The BDT ensemble was capped at 10 k trees of maximum depth 10. With a learning rate of 0.01, this took approximately 9 hours to train on the SK computing system. The feature importance for this model is given in figure 6.13, and the final tree in the generated ensemble is shown in figure 6.14.

The uncut distribution of the BDT score for MC and two weeks of pre-reactor WIT data is shown in figure 6.15. The vast majority of the background is easily identifiable as such thanks to the particularly low quality of the events, resulting in almost all of these events being assigned a score close to 0. MC events aren't as easily identified; however, only a small portion of events are clearly wrongly categorised, likely given the small amount of information in particularly low quality events.

SPLASH was again used to optimise the BDT score cut. Given the effectiveness of the BDT in comparison to the traditional cuts, two weeks of SK-Gd data was used as the BG sample rather than just 24 h. This data, with a >0.2 BDT cut to remove the obvious BG events, is shown in figure 6.16. From this distribution it

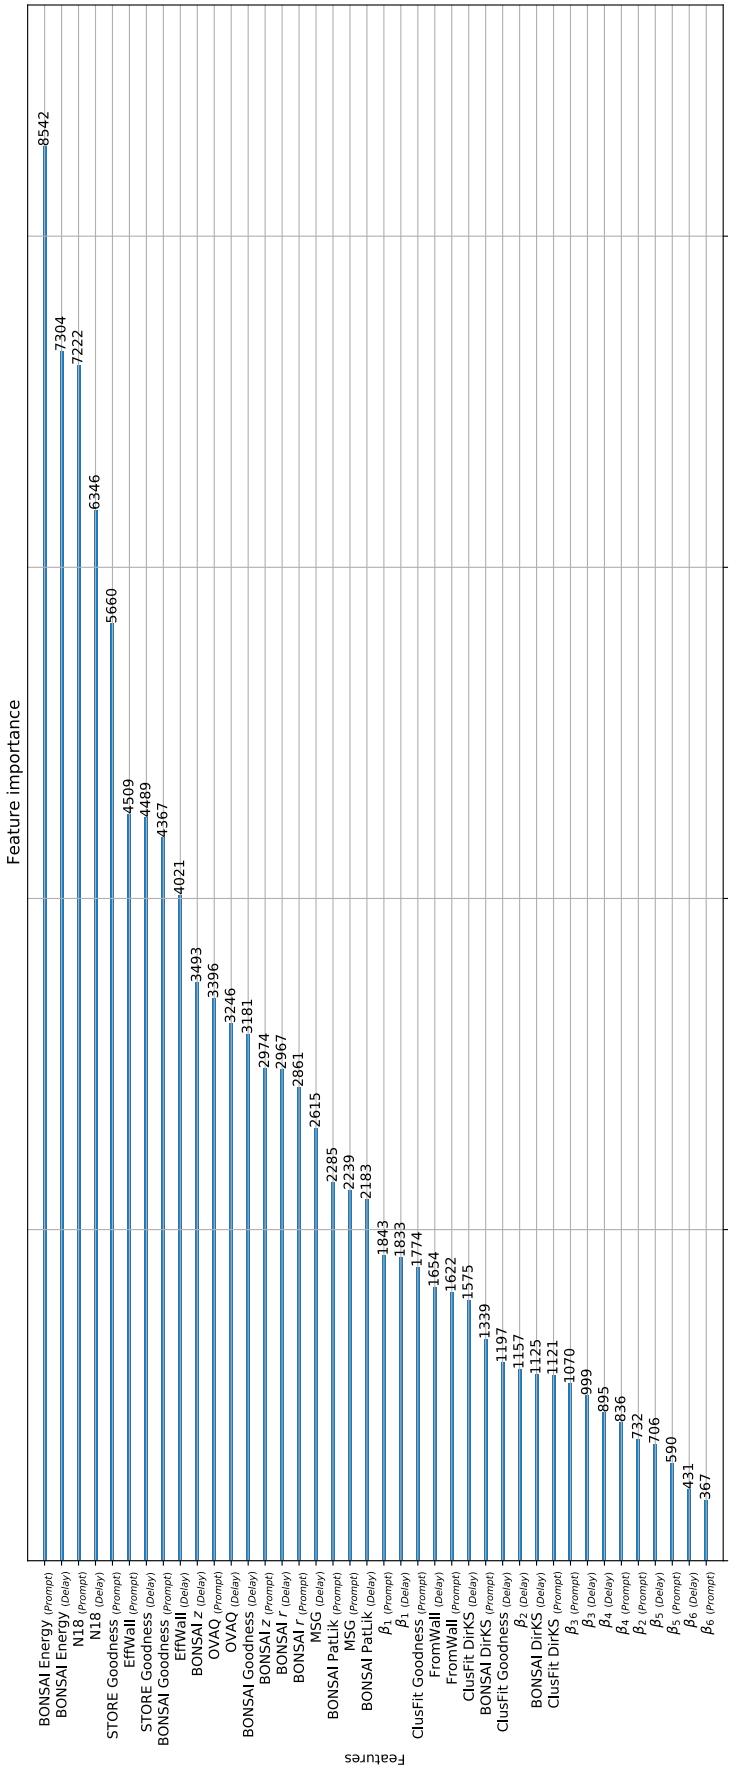


Figure 6.12: The input variable importance for a BDT trained on the 2019 reactor spectrum, including the β isotropy parameters.

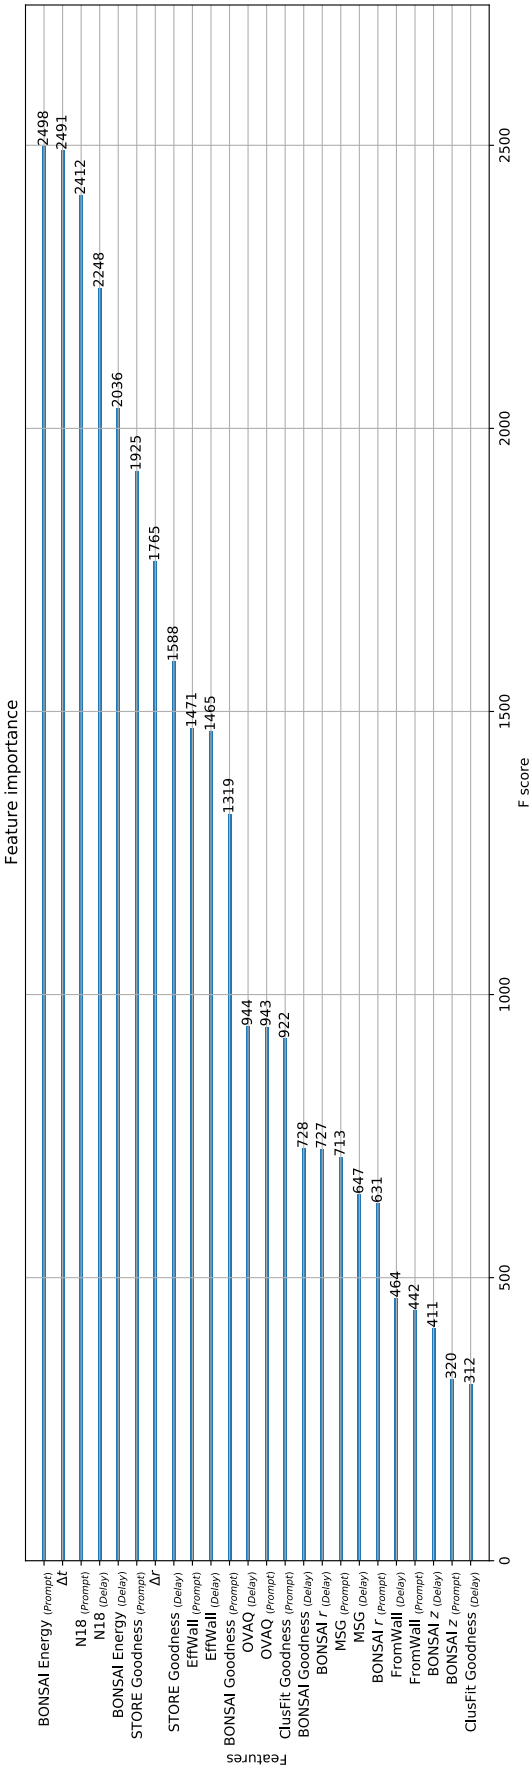


Figure 6.13: The input variable importance for a BDT trained on the simulated 2021 reactor spectrum, the BDT which formed the final reactor sample selection.

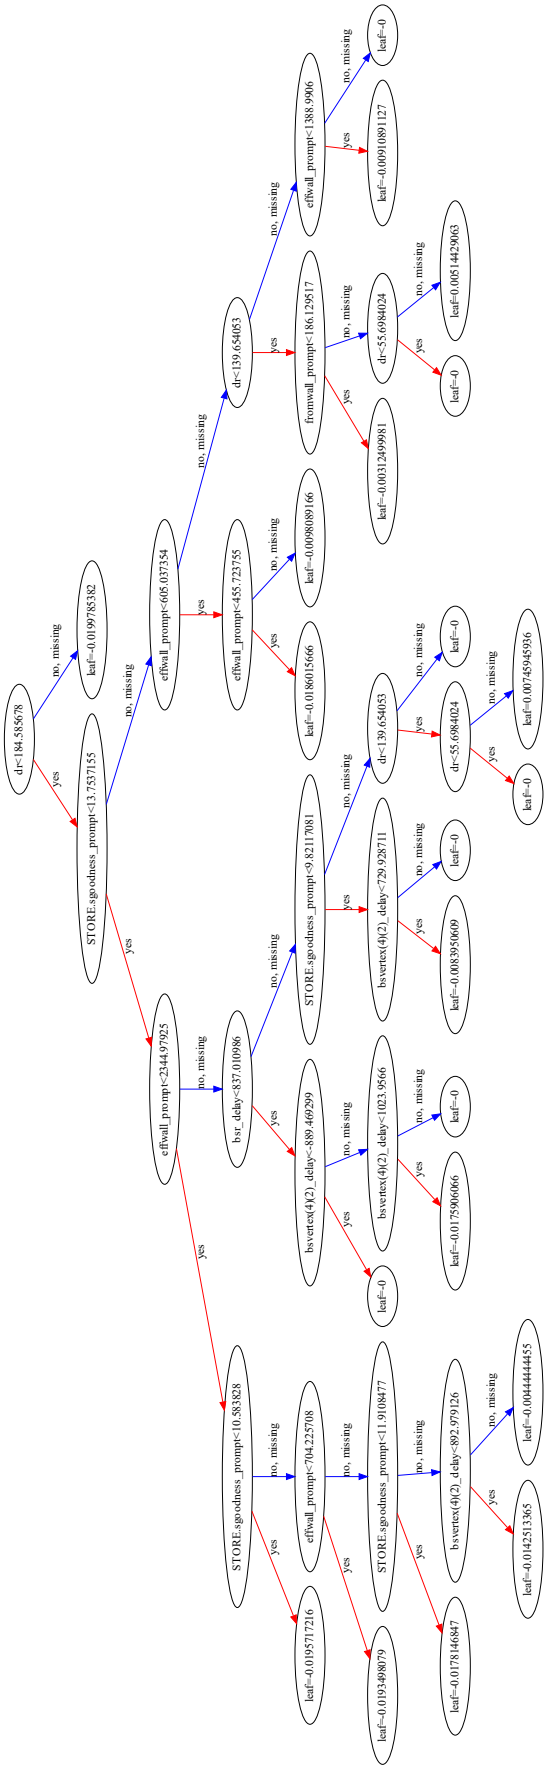


Figure 6.14: The final tree in the ensemble BDT used in this analysis.

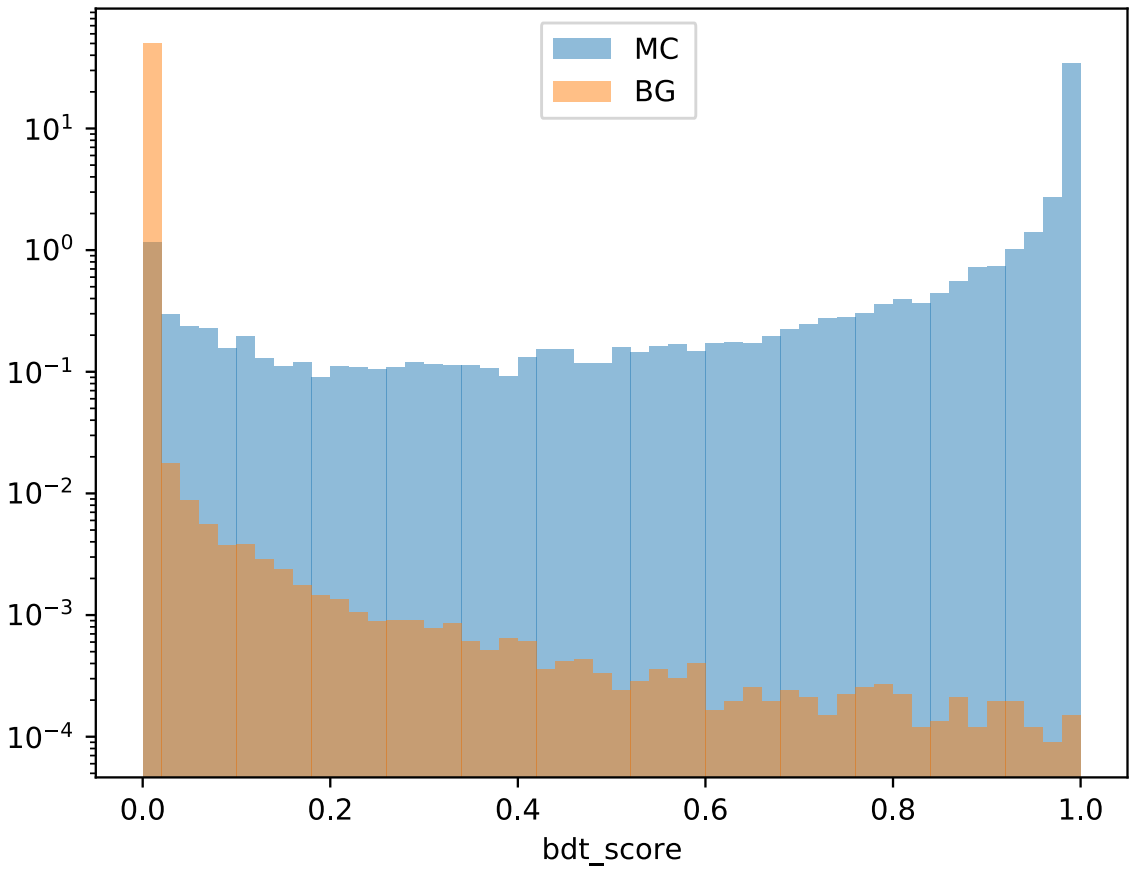


Figure 6.15: The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG).

can be seen there is no obvious cut off, as there are a number of events still present with scores close to 1, shown in more detail in the BDT score > 0.95 plot in 6.17.

It is unclear what the source of these high-score events are; however, they are likely accidental coincidences where the individual event qualities are unusually high and as such no specific non-BDT cuts would trivially remove them.

Optimisation resulted in a BDT score requirement of >0.991, producing an MC efficiency of 10.3% which translates to a signal rate of $0.97\,\mathrm{d}^{-1}$ by the end of 2021. In some cases, such a high and precise BDT cut would be indicative of an overtrained BDT; however in this case it is purely due to the extremely large BG rate compared to MC, requiring a very tight cut to remove a significant amount of BG. An alternative approach would be to apply some minor quality cuts on the event variables before the BDT is trained, to remove the vast majority of the BG first and only leave the questionable events. This would still result in basically the same

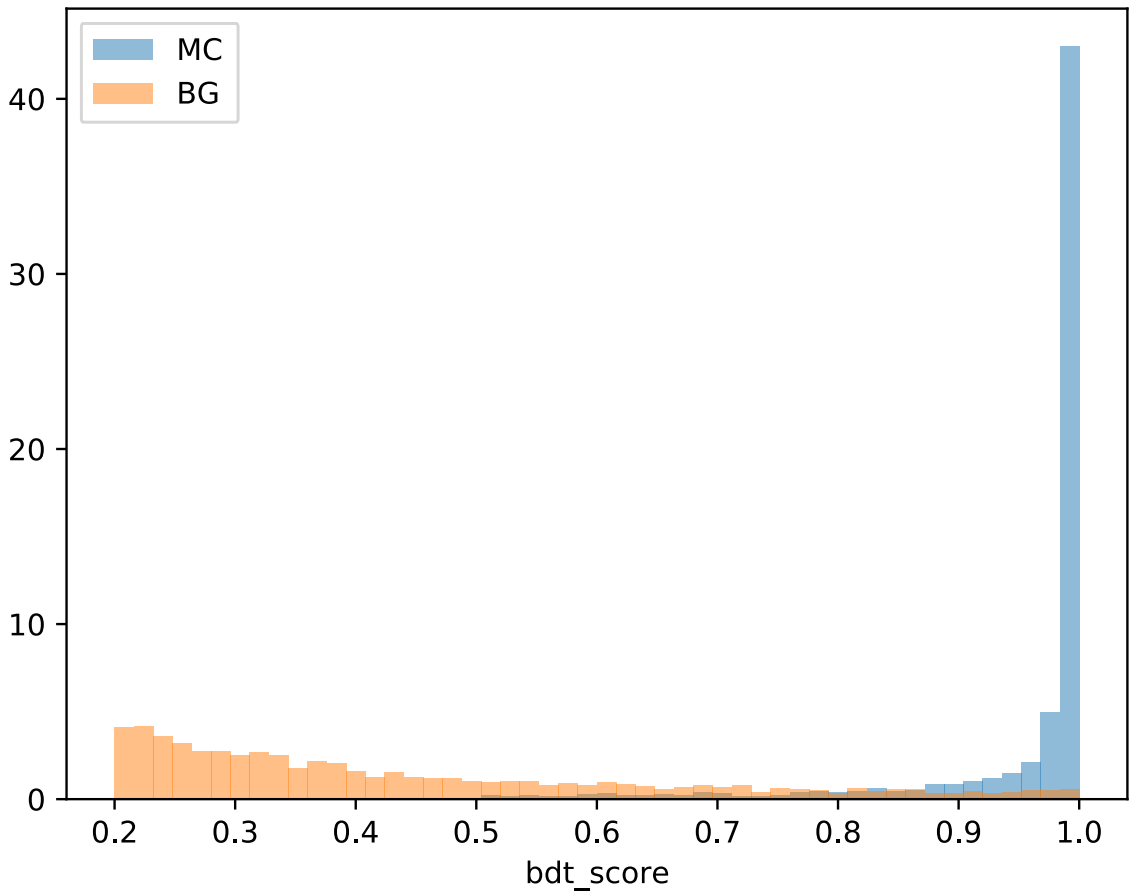


Figure 6.16: The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG), with a >0.2 BDT score cut.

issue, albeit with a less precise BDT score threshold. This approach wasn't taken so as to simplify the final sample selection somewhat, as the initial cuts would have to be taken into account when calculating efficiencies as well as the BDT cut itself.

6.5 Background Prediction and Data Availability

An attempt was made to predict the background rate as a function of time in the post-restart period, without (fully) unblinding the data. Before the data undergoes full reduction ready for analysis, a softer set of cuts was applied, known as the "online" cuts, named as such as they originally mimicked the WIT online thresholds, but later saw the inclusion of a soft BDT score cut of >0.5. The full >0.991 BDT cut is therefore referred to as the offline cut. The intention of the online cuts is to provide a dataset which does not reveal a significant number of reactor events while

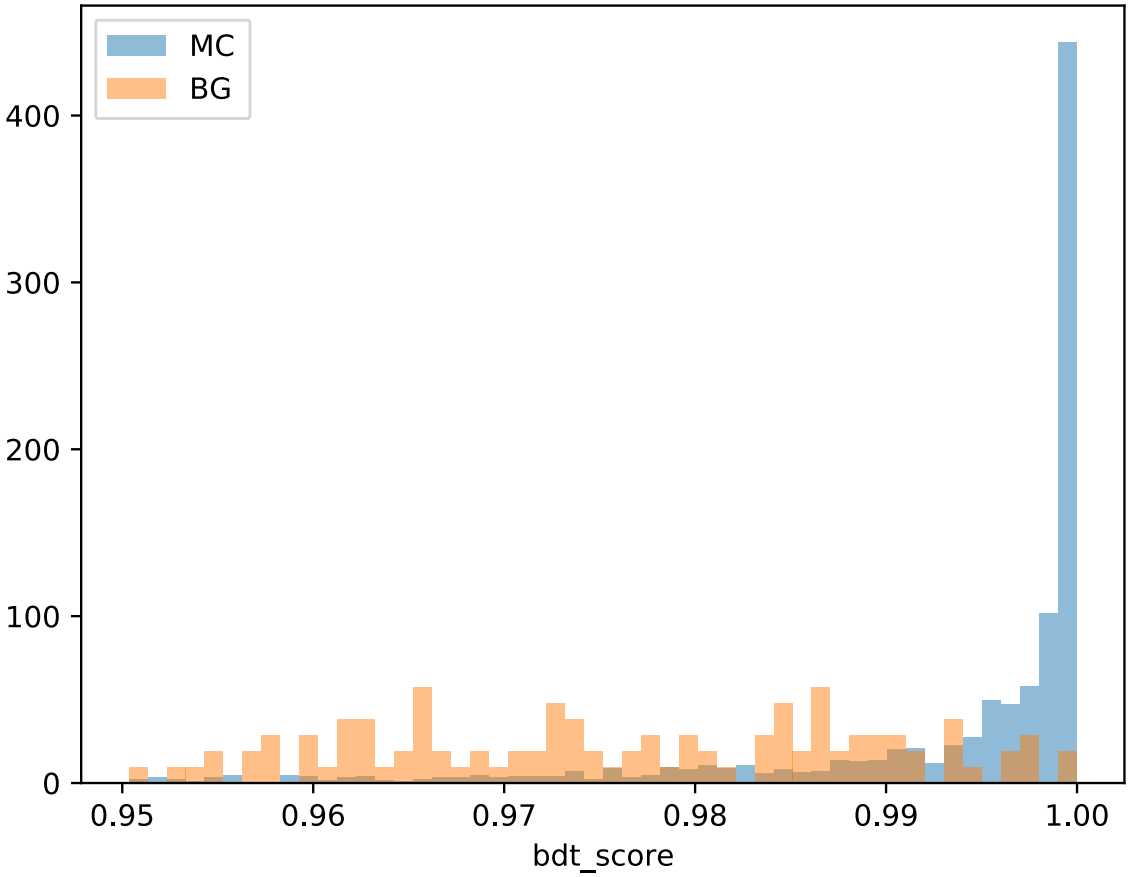


Figure 6.17: The BDT score for 10 k reactor MC events and two weeks of SK-Gd WIT data (BG), with a >0.95 BDT score cut.

providing a metric for the activity of the background which will influence the final dataset. The assumption is that if there is an increase in the rate of events passing a >0.5 BDT cut, the final BDT selection will likely see a similar increase.

To predict the background rate in each dataset, a subset of the pre data has both the online and offline cuts applied. The ratio of offline to online events is then effectively an online→offline efficiency. This is used to scale the number of events in a data period which pass the pre-selection (the online rate) to predict the final background rates. This allows for more direct background stability measurements, beyond using the pre dataset to predict the rates in the post.

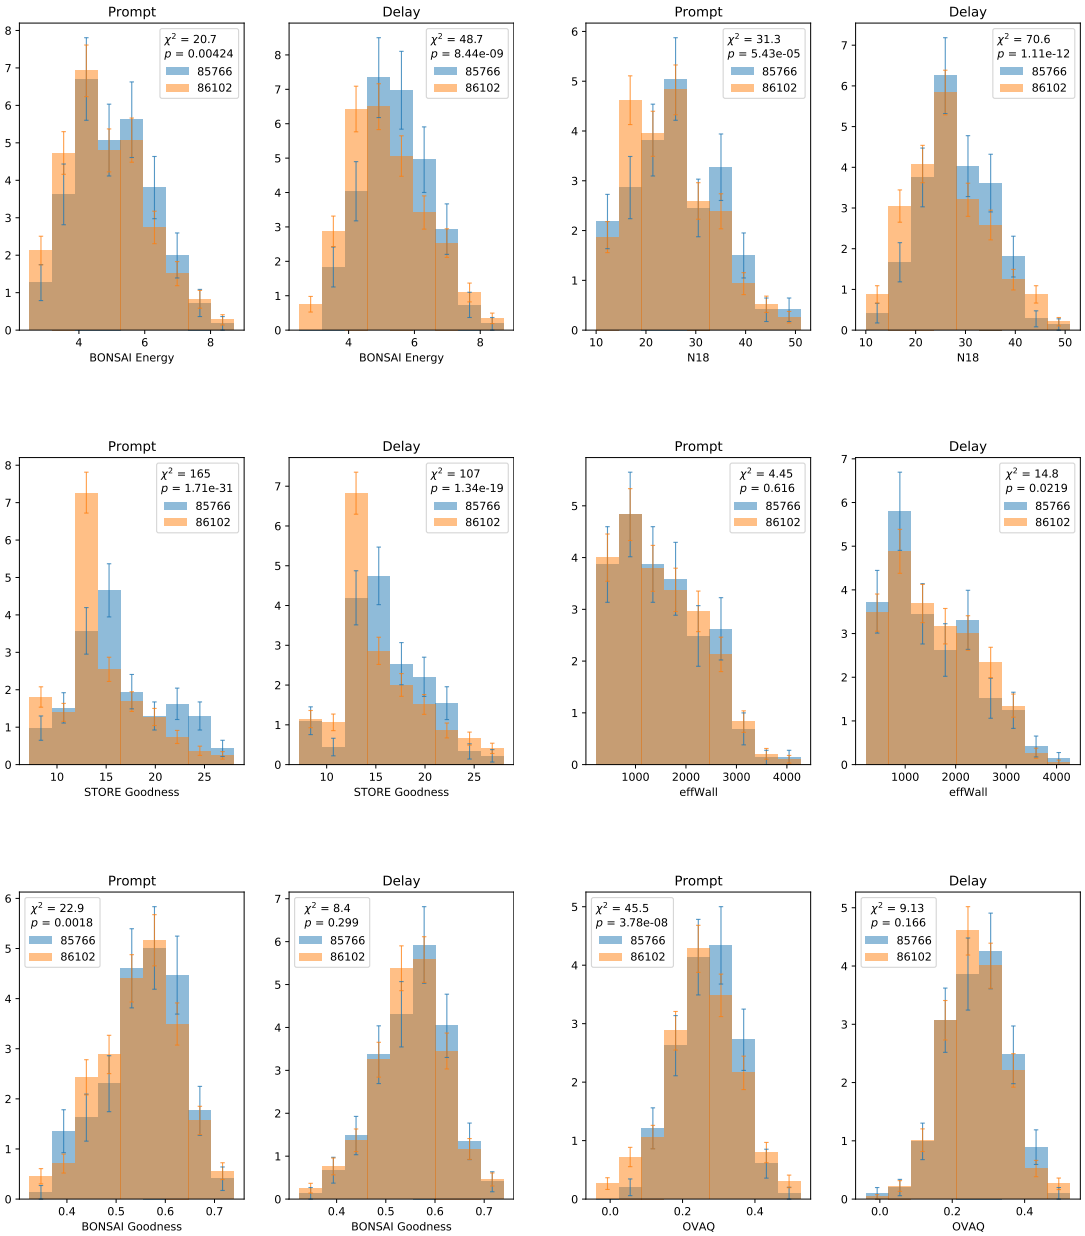
One side effect of this approach is that the data will inherently fit the predicted rates particularly well, as the two are somewhat coupled. Therefore, strict conclusions can not be drawn from the agreement between measured and predicted rates using this approach until the correlation can be somehow quantified. This does not

affect the comparison between data points, however, as is the case in detecting an increase in rates due to reactor restarts. Were it the case that SK's pair rates were stable over time, the subset of data used to calculate the background efficiency could instead be used to simply estimate a flat background rate; however, the dynamic background calculation acts as a probe of the stability of the pair rate in SK.

The online pair rate scaling method was applied to the entire post dataset, which revealed a concern in that the post had significantly higher background rates when compared to the pre; the online rate approximately doubled in the post period. This is not trivially explained, as there are few sources of time-dependant background in SK beyond the increase in activity introduced during Gd loading, which had stabilised by 2021. A soft >0.2 BDT score cut was applied to two 24 h runs in the pre and post periods and the BDT variable distributions were compared to shed some light on the nature of these extra events. Though little can be drawn from such a small sample size, there are some variables which demonstrate a shift in distribution from the pre run (85766) to the post (86102) as shown in 6.18. There are also a far larger number of events in the post (428) compared to the pre (161), more so than predicted by the online pair rate scaling. χ^2 was therefore calculated by scaling the pre data, taken to be the expected, to match that of the post, which was treated as the observed. This is not an exhaustive statistical test given the sample sizes but does indicate which distributions demonstrate a more significant shift.

While not statistically significant, the downward shift in Δt could be indicative of an increase in the dark rate driving the flat event rate up and reducing the time between accidental coincidence events. This would cause a shift upwards in Δr as these events won't necessarily be spatially correlated, a shift which is observed but that is, again, most probably not statistically significant. An increase in dark rate may also explain, for example, the increase in events with a lower STORE goodness, as low-quality events would then have more hits and therefore coincidence hits, feigning a higher quality event by this metric. Given the small sample size, it is difficult to confirm that this is indeed the cause. Regardless, the result is a noticeable shift in BDT score is shown in figure 6.19. There is not enough data in a single

run to determine whether the assumption of a constant online→offline background efficiency holds, and to determine this would result in necessarily unblinding the data somewhat, introducing bias in the BG efficiency calculations.



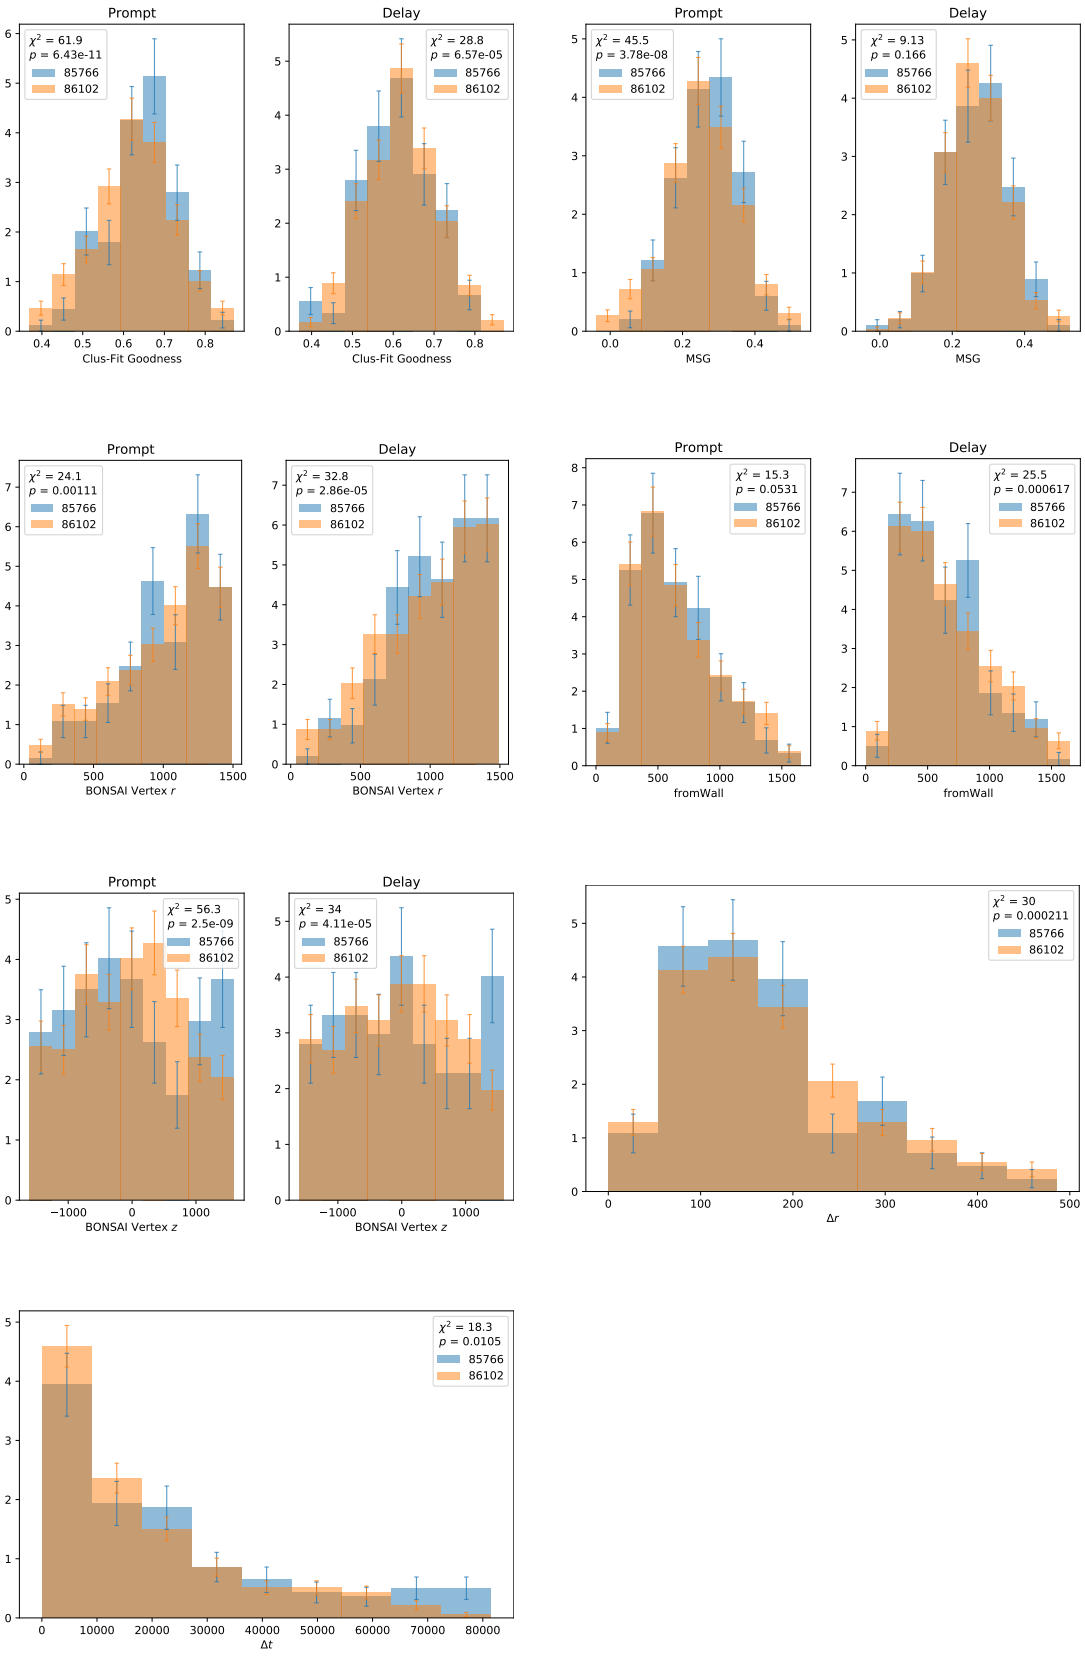


Figure 6.18: A comparison of all BDT variables for a 24h run in the pre (85766) and post (86102) datasets.

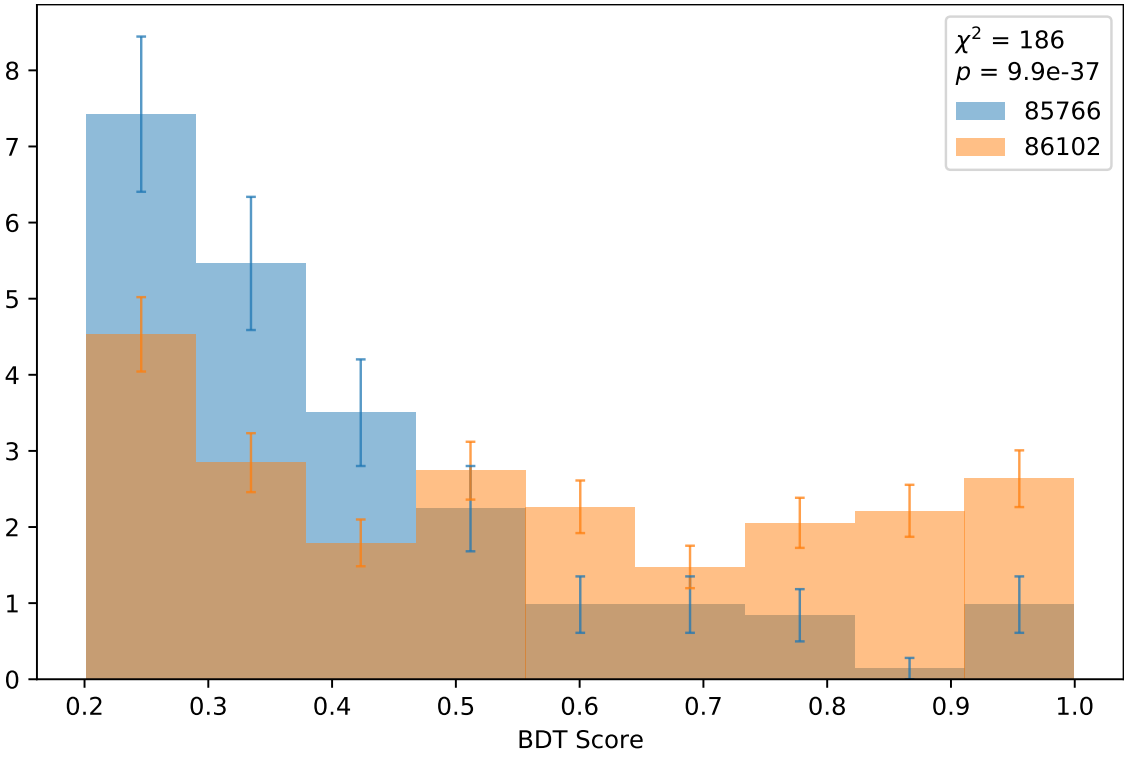


Figure 6.19: A comparison of the BDT score between events in a 24h run in the pre (run 85766) and post (run 86102) datasets.

In this case, it would be useful to view the online rates with more temporal granularity to determine if it was a sudden or gradual increase in the rate, and if the increase is coincident with some detector-related activities. The issue with the WIT data between the pre and post periods inherently prohibits this. The weekly-binned data up to June of 2021 in figure 6.20 shows that the selected data exceeds what is expected based on the online data rates, as well as being inconsistent week-on-week, confirming this data's unsuitability.

A number of files in the entire dataset showed signs of failed reconstruction or some form of corruption which may cause some of the more erratic online rates, especially those seen in the post period. The following requirements were imposed to ensure these erroneous files were removed:

1. N online pairs < 50 and file length $< 1000\,\mathrm{s}$.

   While the typical file length cap is 87.38 s, a number of files were calculated to have a longer length, which was sometimes reflected in the number of events

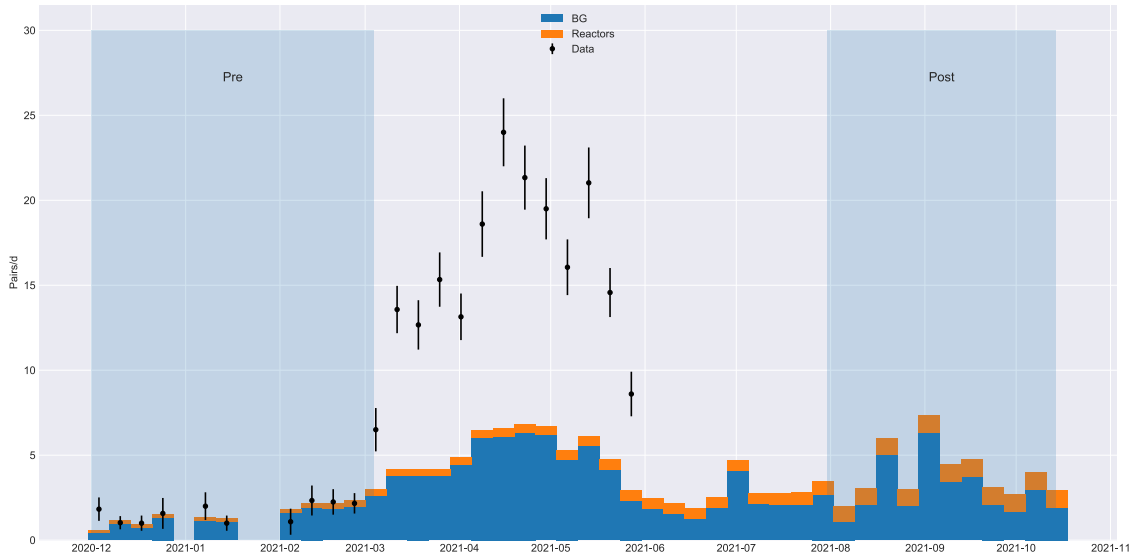


Figure 6.20: The predicted BG and reactor rate over the entire dataset, binned weekly, with no file quality requirements. Errors are statistical only.

they contained. Such files may therefore have an accurate file length and event rate and should be retained.

2. N offline pairs < 20.

   Similar to the previous cut, but ensuring a reasonable number of offline pairs.

3. File size $> 1\,\text{MB}$.

   WIT file sizes are typically between 40–60 MB; however, a number of smaller files seemed to have a strangely large number of online events.

4. File length $> 100\,\text{s}$ for files of size $> 70\,\text{MB}$.

   Ensures that files which are larger in size are so because they cover a longer period of data. This is the most impactful requirement, removing a large portion of the intermediate period between the pre and post data where there appears to be issues with the file length calculations and/or hardware trigger counters.

Additionally, runs are removed if their total file size is more than halved following these cuts. This is a relatively loose requirement given that a number of otherwise good-quality runs are skewed by a single, extremely large file, and the majority

of runs lose only 1–2 % of their total file size, if any at all. This requirement, coupled with requirement 4 in the list above, removes the entirety of the data in the pre-post intermediate period.

Given the relative immaturity of SK-Gd, a number of calibration efforts were undertaken over the course of 2021, including calibration campaigns around a week in length in both the pre and post periods. The runs where calibration is active must be discarded. Ideally the runs in between calibration runs would produce viable data, when work is completed for the day and normal data taking resumes overnight before calibration continues the next day. Inspecting the online pair rate in such runs revealed a significant increase in events not localised to the calibration ports, likely due to activity in the PMTs thanks to HV restarting or as a by-product of opening the tank to insert sources. Each entire calibration period, including a day following, is therefore discarded in this work.

As well as removing these calibration periods, two weeks of data is set aside in the pre period to calculate the background efficiency. Given the low rates involved with this analysis, it is important to properly understand the expected background rate. Two weeks of data straddling the 2021 new year was used to calculate the background efficiency. This subset was selected specifically to give the most accurate prediction in the rest of the pre dataset. The separate two weeks of data used to optimise the BDT cut wasn't removed, as it doesn't bias the BG rate calculations and merely provides a general distribution to optimise against. The 24h run of data used as BG in the BDT training was excluded, however.

The trimmed-down rates are shown weekly-binned in figure 6.21, showing good agreement between the prediction and the data in the pre period and an increased and still somewhat erratic set of rates in the post.

Between figures 6.20 and 6.21, there is reason to believe the WIT issue itself is the cause of the increased rates. Despite the file length issues being resolved at the start of the post period and the overall improvement of the file quality, the erratic, high online event rate seems to persist. This is discussed in more detail in the context of the unblinded data in section 6.6.

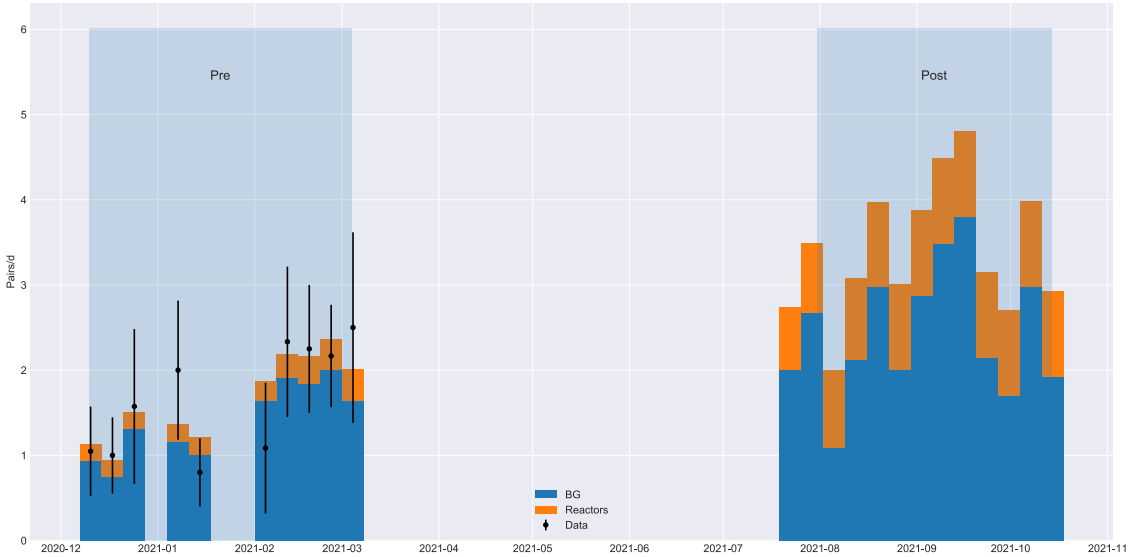


Figure 6.21: The predicted BG and reactor rate over the entire dataset, weekly binned, with file quality requirements. Errors are statistical only.

	Livetime [d]	Total Rate [/d]	Total Events	BG Events	Frac. Err.
Pre	35.6	1.6	58	50*	0.13
Post	51.8	3.7*	189*	140*	0.073*

Table 6.2: The predicted pre/post results based on pre offline data and post online. An asterisk denotes a predicted value.

With 100% livetime, the pre/post approach would provide 93 and 75 days of data in each dataset respectively; however, the trimming described above significantly reduces this. The true total livetimes, calculated using the known run lengths, are 35.6 d for the pre period and 51.8 d for the post. While the post period can be expanded once the offline processing systems have been recovered, there is little scope to expand the pre period, limiting the statistical strength of any conclusion that could be gained from this analysis. Section 6.6 discusses an alternative approach to livetime calculations, using the files themselves, which may retain more data otherwise lost by removing whole runs. Despite this, between background efficiency calculations, high data rates in 2020, and the abundance of calibration periods, the pre period will always be somewhat limited in its livetime. The expected pre/post results using the reduced livetimes are given in table 6.2.

6.6 Results

In the absence of a more in-depth systematic error evaluation, all errors shown in these results are statistical only, calculated assuming Poissonian statistics using $\sigma = \sqrt{N}$ for N counts in each bin. The comparison between the rates measured in the pre and post periods is given in figure 6.22, showing an increase in the rate between the two periods far in excess of what is predicted. The measured rate in the post period is $(8.8 \pm 0.4)\,\mathrm{d}^{-1}$ while the predicted is $3.7\,\mathrm{d}^{-1}$.

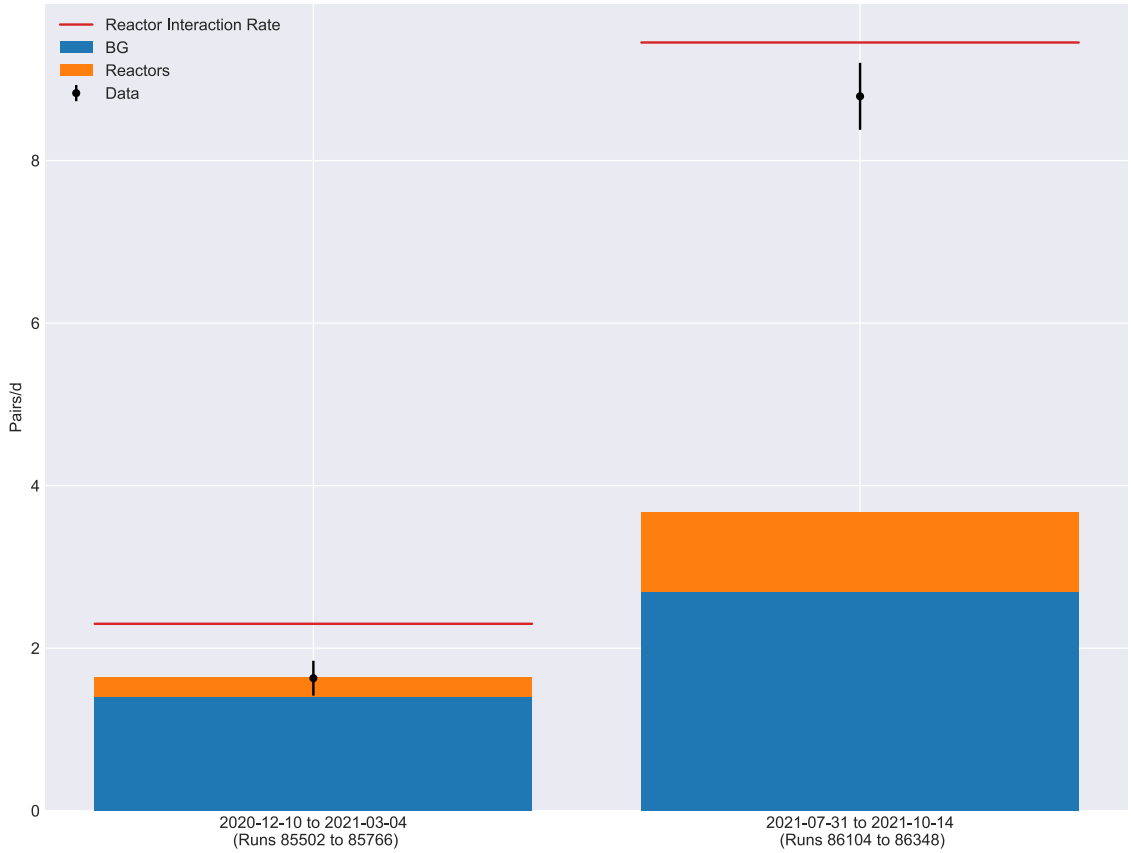


Figure 6.22: The pair rate in the pre and post periods, compared to the expected BG and reactor rates. The red line represents the reactor rate with 100% efficiency. Errors are statistical only.

Figure 6.23 shows that the weekly-binned rates match the pattern of the predicted background, albeit with a large offset. The implication of this is that the core principle of tracking the background pattern through the online pair rate is somewhat accurate, but another effect is causing the online→offline background efficiency to dramatically increase between the two periods. A potential explanation is

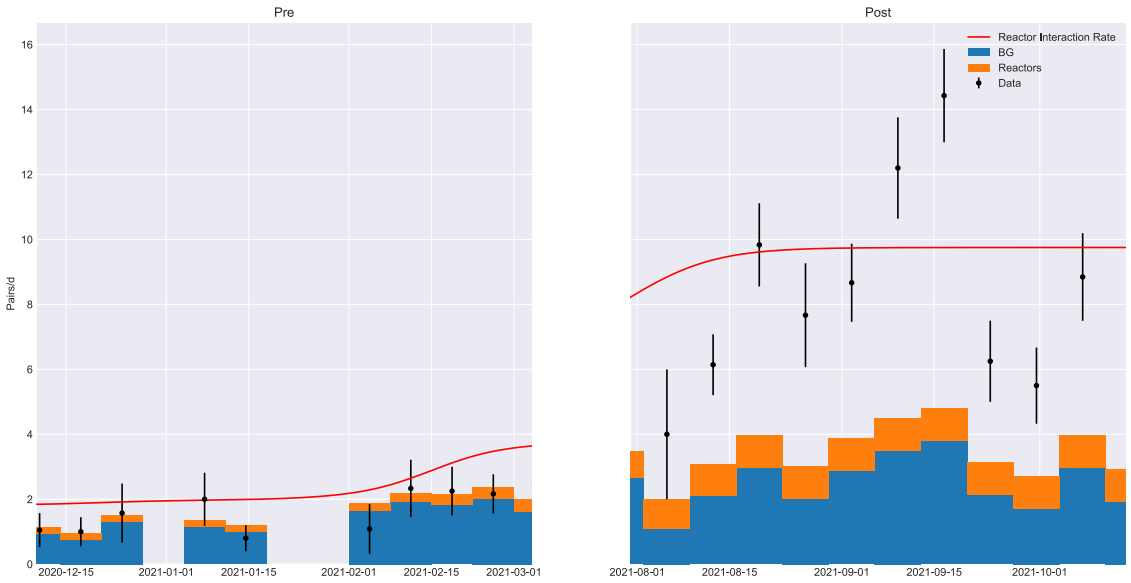


Figure 6.23: The pair rate in the pre and post periods, weekly binned, compared to the expected BG and reactor rates, with the livetime calculated using assumed 100% run lengths. The red line represents the reactor rate with 100% efficiency. Errors are statistical only.

that the signal efficiency derived from MC is underestimated, meaning what is classified as BG in fact may contain signal events and the increase in rate would be greater than predicted.

The data in figure 6.23 is plotted alongside the predicted reactor interaction rate, without any efficiency scaling. While the majority of the data points are below this line, some lie above, implying that this increase is unlikely to be solely driven by an increase in reactor rate.

Figure 6.24 shows the required signal efficiency in the pre period against the increase in BG rate for the increase to be consistent with the expected reactor neutrino rate. If the BG remained constant (1 on the x-axis), and the signal efficiency similarly did not change between the pre and post periods (orange line), then the signal efficiency would have to be an unreasonable 93% to explain the increase. The predicted efficiency is 10.3% and, given that approximately half of all events exhibit neutron capture on hydrogen with around 17% WIT efficiency, an efficiency beyond 50% is generally unreasonable.

Figure 6.24 also shows these required efficiency curves for a range of relative increases in the MC efficiency between the two periods. This increase would more

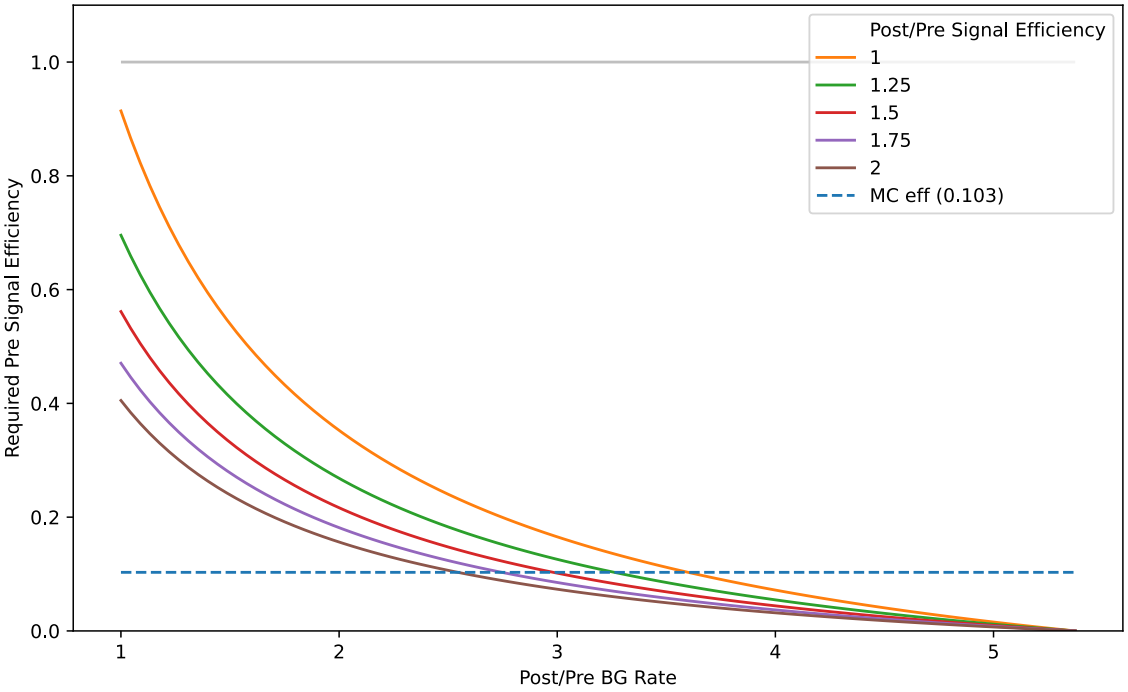


Figure 6.24: The required signal (MC) efficiency in the pre period needed to explain the observed pair rate increase for a range of BG rate increases and relative signal efficiency increases.

likely be somewhat coupled to the change in BG rate, however, and so figure 6.25 displays these curves where the efficiency of the signal increases as a factor of the BG. If this was driven by an change in dark rate, for example, the inflated number of hits in the trigger window would boost the likelihood that an IBD pair would pass the cuts, as well as accidental coincidences.

Another indication that this increase is not due to the reactors is that the week-on-week data points aren't entirely consistent with one another, and it is expected that the reactor rate would remain approximately constant once the cores had reached full power, as expected in this period. It is also expected, however, that there are no time-dependant sources of background in SK. The detector itself remained stable through each of these periods, with no major changes in its water supply or gadolinium concentration. If the number of events is expected to remain constant within each period, suspicion may then fall on the livetime used to calculate the rate.

To calculate the rates given in figures 6.22 and 6.23, the total number of events is counted and divided by the sum of the length of all runs used in the dataset. This

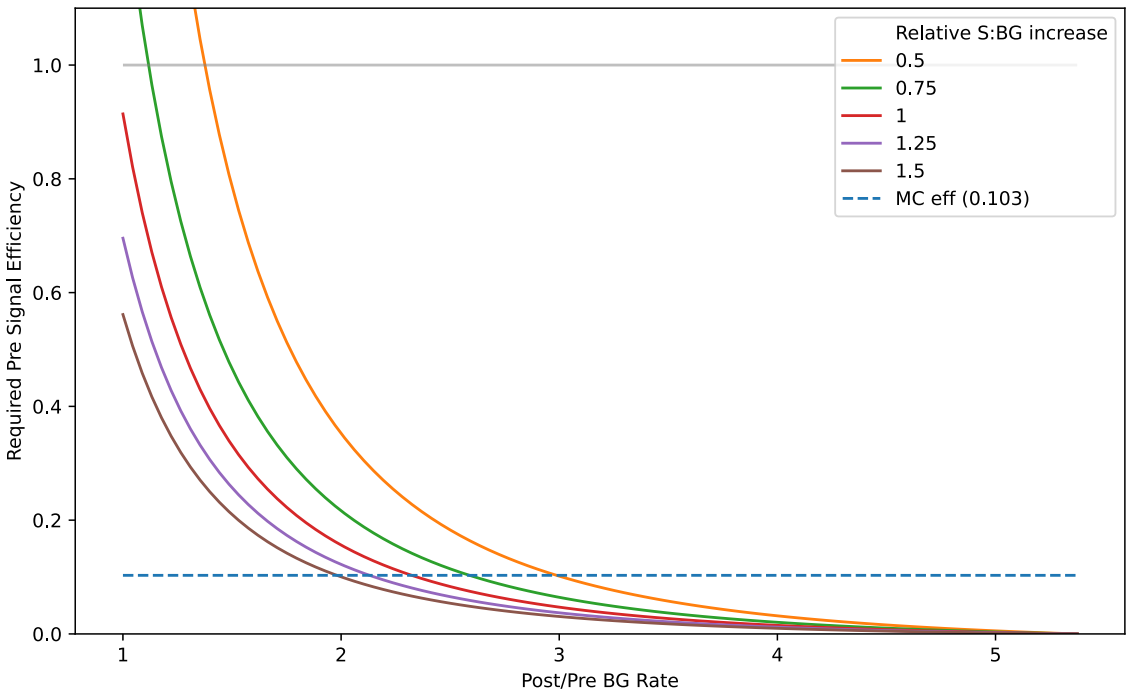


Figure 6.25: The required signal (MC) efficiency in the pre period needed to explain the observed pair rate increase for a range of BG rate increases where the MC efficiency increases as a factor of the BG increase.

assumes that each run has 100% livetime, and the number and file size differences before and after the quality cuts described in 6.5 typically reflect this. These cuts may remove more viable data than is assumed, however, resulting in a disconnect between the calculated and true livetimes.

The livetime can be alternatively calculated using the files themselves. While WIT data doesn't inherently contain a description of how long a period it covers, the length of a file can be calculated using the difference in hardware triggers in the data. This approach does not provide a fully accurate livetime, as the data in the fully reconstructed file does not necessarily contain the first and last events of the raw file produced by WIT. The file length calculated using this procedure will therefore be underestimated by an amount equal to the sum of the time between the true start of the file and the first event present, and between the last event and the true end of the file. This varies depending on the event rate; however, the difference is not insignificant. The difference between the run length and summed length of all files in a run is given in figure 6.26. This shows that, in general, the

two rarely agree, especially during the pre/post intermediary period. It is therefore difficult to quantify to what extent the file length-based calculations underestimate the length of a run, and to determine which method is more accurate.

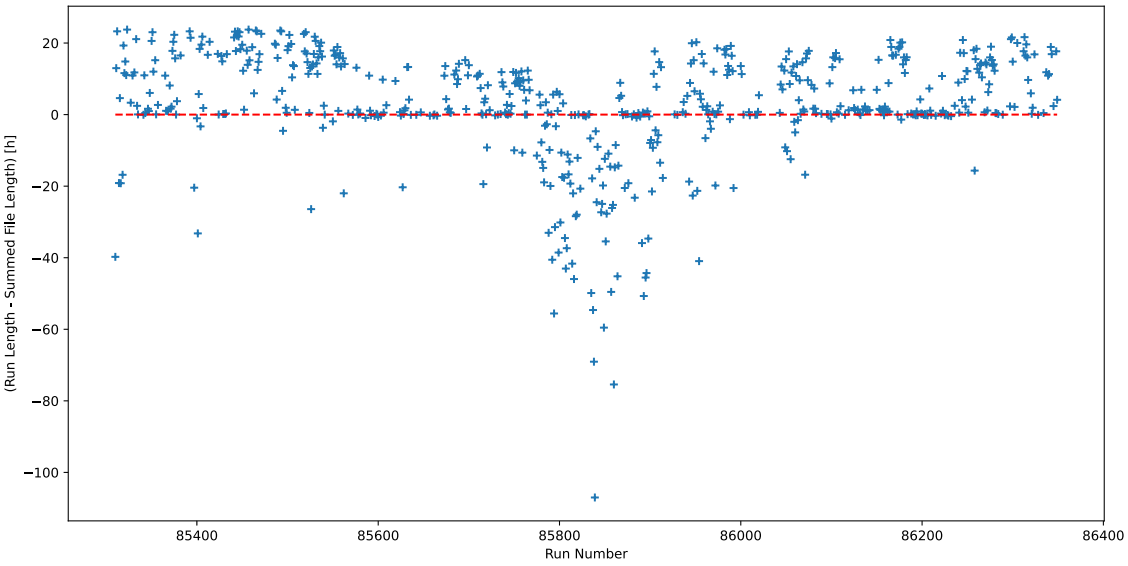


Figure 6.26: The difference between the known run length and run length derived from the sum of all calculated file lengths.

Regardless, the weekly-binned rates calculated using this alternative livetime calculation are given in figure 6.27. Using this livetime calculation, a far more stable background rate is predicted, potentially indicating that the file length calculations are more accurate than one would assume from figure 6.26. The average background rate is predicted to be far higher than when calculated using the uncut run lengths. This could confirm the theory that inaccurate run-based livetime calculations are the cause of the varying background rates, as an overprediction of livetime would lead to an underprediction of the background rates. If that were the case, figure 6.27 indicates insufficient background rejection by the BDT selection to demonstrate a reactor-driven increase, which is unsurprising given that the BDT cut was optimised using a BG rate calculated using the run-based livetime approach.

Though the more-stable background implies the alternative livetime calculation may be preferable, this approach needs more in-depth validation, and shows signs that it may not be entirely accurate.

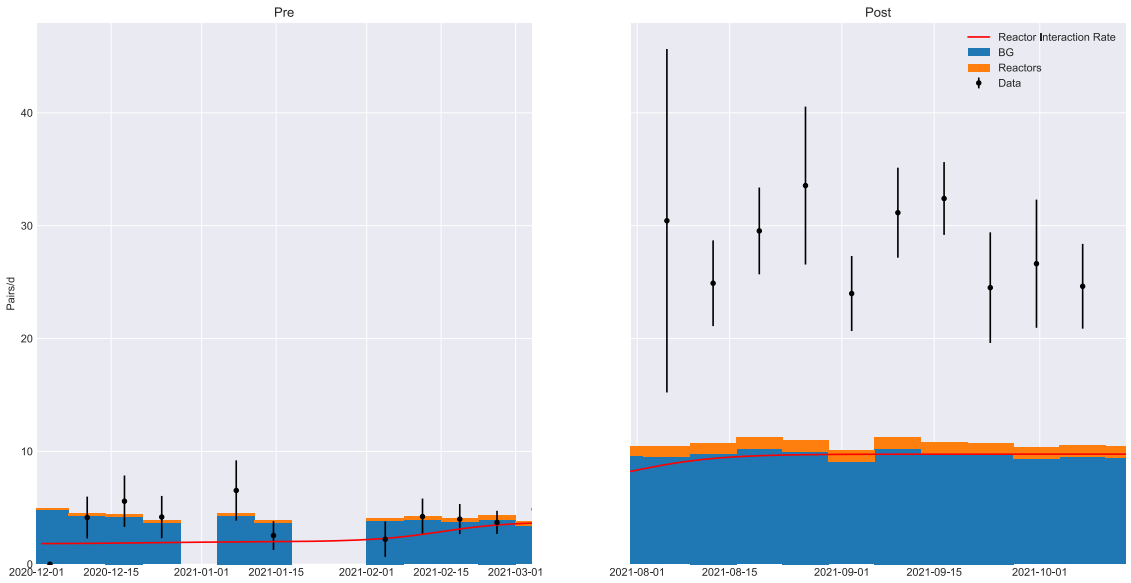


Figure 6.27: The pair rate in the pre and post periods, weekly binned, compared to the expected BG and reactor rates, with the livetime calculated using the file lengths as opposed to assumed 100% livetime run lengths. Errors are statistical only.

Firstly, the difference between the true file length and the calculated file length should be cross-checked through comparing raw WIT data with the fully reconstructed pair data to determine the true significance of the missing start and end periods. If the deficiency is large enough to explain the disagreement shown in figure 6.26, then an offset or scaling can be applied to account for this. If instead the missing time is insufficient, the file length calculations may be untrustworthy, as it is unlikely that such a large portion of WIT data would have been lost. There exists an alternative approach to calculating the length of a file which uses hardware trigger meta-information saved for each file; however, this method has not been fully explored given that some files appear to be missing this information.

Secondly, the unstable period between the pre and post periods appears to be driven at least in part by issues with the saved hardware trigger counter which would be used to make this calculation. Until this issue is diagnosed and confirmed to not affect the rest of the data, it is preferable to avoid making livetime calculations using this information. Indeed, the higher offline rate observed in the post period may be due to this issue itself causing a shift in the file length calculations, though this would most likely be reflected in the background rate predictions.

Regardless of which way the livetime is calculated, the measured rates in the post period far exceed the predicted, and the reason for this is unclear and cannot be easily identified without clear analysis of the pre-post transition period. While the issues with this data prohibit a continuous pair rate measurement linking the two datasets, some initial investigations are possible.

If the quality cut removing files of spurious file length is removed (item 4 in the quality cuts listed in section 6.5), the majority of the intermediate data remains available. While this data was specifically avoided due to issues with the hardware triggers, figure 6.28 shows that the rates calculated using the file length method suddenly increase immediately following the pre period, as expected; however, there appears to be no obvious disconnect between the intermediate and post data periods. This would imply that the issue that motivated the rejection of the intermediate period is somehow still affecting the post data and is not exclusively related to the file length calculations.

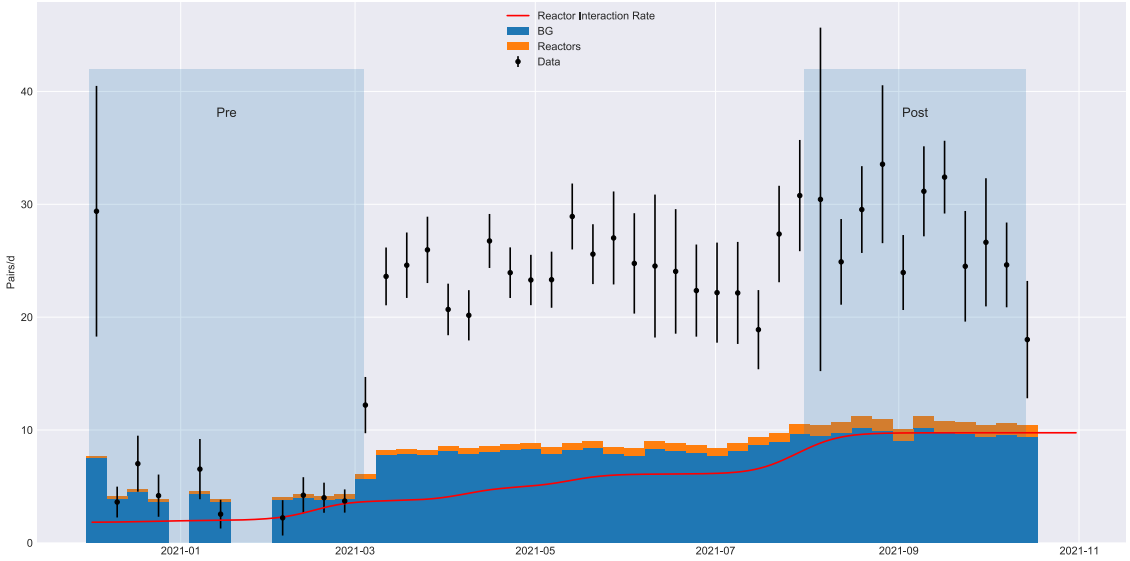


Figure 6.28: The pair rate across the whole WIT pair dataset, without quality cuts which would remove the pre→post intermediary data. Livetimes are calculated using the summed file lengths. Errors are statistical only.

Calculating livetime using the run length, shown in figure 6.29, similarly shows no obvious disconnect and arguably demonstrates more erratic week-on-week rates in

the post period compared to the intermediate, possibly indicative of the inaccuracy of the run-based livetime calculations in this period where frequent calibration occurred.

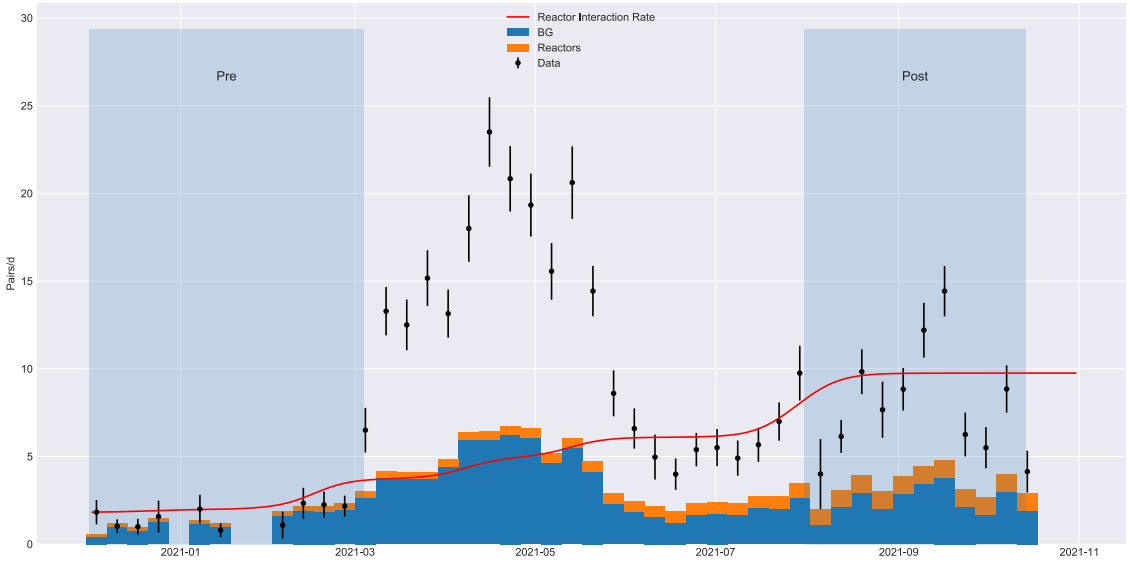


Figure 6.29: The pair rate across the whole WIT pair dataset, without quality cuts which would remove the pre→post intermediary data. Livetimes are calculated using the run lengths. Errors are statistical only.

Given that the data largely agrees with the predicted rates in the pre period when using either livetime calculations but are far in excess beyond this dataset, the WIT issues must be addressed, if possible, before any strict conclusion can be drawn.

If this data is permanently corrupted, a future analysis could, instead of detecting a new restart, monitor the monthly rate for a decrease in the pair rate coincident with the refuelling of a reactor. Such an analysis would require either the cooperation of nuclear plant operators or corroboration of identified potential restarts with those reported in the PRIS report when available. This would arguably act as a more convincing proof of the detection of reactor neutrinos, given the truly blind search for a refuelling. It would also more realistically simulate the real-world motivation for this analysis, where the activity of a single reactor core must be monitored against a reactor background for unannounced shutdowns.

7
Conclusion

Thanks to gadolinium loading, SK's sensitivity to neutron capture, and therefore inverse beta decay interactions, has increased significantly. For the first time in its long-running history, the detector is in a position where it may be able to detect a significant number of reactor neutrinos. This is compounded with the unique position the Japanese nuclear power sector is currently in, where a number of reactors, offline in some cases for over a decade, have recently seen restarts. There is a window of WIT data available from December of 2020, when the detector had stabilised following gadolinium loading, up to October of 2021, the latest data that had been fully reconstructed before an upgrade to the SK computing systems broke the processing software. In this window, a number of restarts took place which increased the expected reactor interaction rate from $1.82\,\mathrm{d}^{-1}$ to $9.75\,\mathrm{d}^{-1}$. Many steps were taken in support of this unprecedented search in SK.

A coincidence trigger has been added to WIT, the system responsible for saving the crucial ultra-low energy events of the focus of this search. This modification relaxes the trigger requirements for events in temporal and spatial coincidence, as with IBD positron/neutron pairs, and tags these coincident events with a unique identifier. This tagging simplifies the offline identification of paired events and allows pairs to skip the first offline reduction step typically applied to WIT data.

To simulate the reactor neutrino spectrum at SK, the SKReact software was created. Based on [101], this provides a tuneable, SK-specific reactor neutrino simulation package, including an estimation of the effect of reconstruction on the energy spectrum. SKReact was used to generate the spectrum used as input to SK's Monte Carlo software to produce the sample used to develop the reactor neutrino selection.

The sample selection itself was supported by the newly-developed SPLASH framework, a cut optimisation and machine learning-based analysis platform for use with WIT and general SK data. SPLASH was used to develop an optimised cut-based analysis; however, this was shown to provide insufficient background rejection to provide a significantly detectable rate increase over the course of 2021 due to reactors.

A boosted decision tree was then developed using the XGBoost algorithm, trained on 2021-spectrum MC events and real WIT data from the pre-reactor restart period which was collectively treated as background. Though trained on a limited MC dataset, the BDT selection was optimised and shown to provide a significant signal above the expected background rate based on pre-reactor restart event rates in WIT. The expected signal efficiency was calculated to be 10.3%, with a predicted background rate of less than one event per day.

Inspection of the data quality throughout this 2020/2021 dataset revealed that, between March and August of 2021, the WIT data appeared to have erratic data rates with inconsistent data file properties. As such, two periods were defined, the "pre" and "post" periods signalling the data before and after this intermediate period, and the intention of this study was to detect a significant increase in rate between the two periods.

The background rate in each period was predicted by first applying a soft "online" BDT cut to a two-week period straddling the 2021 new year. Events which pass this pre-selection are referred to as online events. The efficiency of online background events passing the full "offline" BDT cut was then measured. Finally, the number of online events in the pre and post periods were scaled by this efficiency. This revealed a significant increase in the background rate, as indicated by the increase in the

online pair rate. Inspection of a run of pre and post data indicated there were some shifts in the BDT input variables which may be indicative of an increase in dark rate, though the small sample size was insufficient to draw any clear conclusions.

Once unblinded, the data in the post period was found to be far in excess of the predicted rate which itself significantly exceeded both the measured and predicted rates in the pre period. Though not inconsistent with an increase due to reactor neutrinos, the signal efficiency would have to have been underestimated to explain the increase; what is defined as "background" would then instead contain a significant number of signal events. Further inspection of the pre and post datasets as well as the intermediary period signalled that the increase may be due to persistent, as-yet undiagnosed issues with the WIT data.

Before any conclusion can be drawn, the issues with the WIT data must be addressed and the confidence in the MC efficiency and therefore BG composition further investigated. Based on the Monte-Carlo studies in this thesis, SK has the capability to detect a significant increase in reactor neutrino rate due to a reactor restart once these issues are addressed. An alternate analysis perhaps aiming to identify a decrease in rate due to a reactor refuelling may be a sufficiently convincing signal to confirm the identification of reactor neutrinos in SK. The work in this thesis has laid all the necessary groundwork, from data acquisition to the final analysis, to support such a search.

Appendices

A

Japanese and South Korean Nuclear Reactors

Table A.1: All nuclear reactors in Japan and South Korea, ordered by $(P_{th})/(R_{SK})^2$ where P_{th} is thermal power output capacity of the reactor and R_{SK} is the distance from the core to the SK detector [75]. An asterisk by a plant's name denotes a South Korean plant. A status of pending refers to plants that have applied for a restart and are pending approval from either the local government or the NRA.

Plant	Core	Type	P_{Th} [MW]	R_{SK} [km]	P_{Th}/R_{SK}^2 [MW km^{-2}]	Status
Shika	2	BWR	3926	87.5	0.513	Pending
	1	BWR	1593	87.5	0.208	Offline
Tsuruga	2	PWR	3411	138.8	0.177	Pending
Kashiwazaki-Kariwa	7	BWR	3926	159.9	0.154	Pending
	6	BWR	3926	159.9	0.153	Pending
	5	BWR	3293	160.0	0.129	Offline
	4	BWR	3293	160.0	0.129	Offline
	3	BWR	3293	160.1	0.129	Offline
	2	BWR	3293	160.1	0.128	Offline
	1	BWR	3293	160.1	0.128	Offline
Mihama	3	PWR	2440	145.6	0.115	Online
Ohi	1	PWR	3423	178.7	0.107	Offline
	2	PWR	3423	178.7	0.107	Offline
	4	PWR	3423	178.8	0.107	Online
	3	PWR	3423	178.8	0.107	Online
Hamaoka	5	BWR	3926	214.4	0.085	Offline
Takahama	4	PWR	2660	191.4	0.073	Online
	3	PWR	2660	191.4	0.073	Online

Hamaoka	3	BWR	3293	214.3	0.072	Pending
	4	BWR	3293	214.3	0.072	Pending
Mihama	2	PWR	1456	145.7	0.069	Offline
Takahama	2	PWR	2440	191.4	0.067	Pending
	1	PWR	2440	191.5	0.067	Pending
Tsuruga	1	BWR	1070	138.9	0.055	Offline
Hamaoka	2	BWR	2436	214.6	0.053	Offline
Mihama	1	PWR	1031	145.6	0.049	Offline
Tokai	2	BWR	3293	295.8	0.038	Pending
Hamaoka	1	BWR	1593	214.6	0.035	Offline
Fukushima-Daini	1	BWR	3293	345.1	0.028	Offline
	2	BWR	3293	345.1	0.028	Offline
	3	BWR	3293	345.2	0.028	Offline
	4	BWR	3293	345.3	0.028	Offline
Fukushima-Daiichi	6	BWR	3293	349.2	0.027	Offline
	2	BWR	2381	349.1	0.02	Offline
	3	BWR	2381	349.2	0.02	Offline
	5	BWR	2381	349.2	0.02	Offline
	4	BWR	2381	349.2	0.02	Offline
Shimane	2	BWR	2436	400.6	0.015	Pending
Onagawa	2	BWR	2436	430.6	0.013	Pending
	3	BWR	2436	430.6	0.013	Offline
Fukushima-Daiichi	1	BWR	1380	349.0	0.011	Offline
Shimane	1	BWR	1380	400.7	0.009	Offline
Onagawa	1	BWR	1593	430.6	0.009	Offline
Ikata	3	PWR	2660	560.3	0.008	Online
Higashidori (Tohoku)	1	BWR	3293	636.1	0.008	Pending
Shin-Kori*	4	PWR	3983	732.0	0.007	Online
	3	PWR	3983	732.0	0.007	Online
Genkai	3	PWR	3423	753.7	0.006	Online
	4	PWR	3423	753.7	0.006	Online
Shin-Wolsong*	1	PWR	2825	708.9	0.006	Online
	2	PWR	2825	708.9	0.006	Online
Hanul*	6	PWR	2825	709.0	0.006	Online
	5	PWR	2825	709.0	0.006	Online
	4	PWR	2825	709.1	0.006	Online
	3	PWR	2825	709.2	0.006	Online
	2	PWR	2775	709.3	0.006	Online
	1	PWR	2775	709.4	0.006	Online
Kori*	4	PWR	2912	732.5	0.005	Online
	3	PWR	2912	732.6	0.005	Online
Shin-Kori*	1	PWR	2825	732.0	0.005	Online
	2	PWR	2825	732.0	0.005	Online
Ikata	2	PWR	1650	560.4	0.005	Offline
	1	PWR	1650	560.5	0.005	Offline

Tomari	3	PWR	2660	783.1	0.004	Pending
Wolsong*	4	PHWR	2061	708.6	0.004	Online
	3	PHWR	2061	708.7	0.004	Online
	2	PHWR	2061	708.8	0.004	Online
	1	PHWR	2061	708.9	0.004	Online
Sendai	2	PWR	2660	829.8	0.004	Online
	1	PWR	2660	829.8	0.004	Online
Kori*	2	PWR	1882	732.7	0.004	Online
	1	PWR	1729	732.7	0.003	Offline
Genkai	1	PWR	1650	753.6	0.003	Offline
	2	PWR	1650	753.7	0.003	Offline
Hanbit*	3	PWR	2825	987.0	0.003	Online
	4	PWR	2825	987.0	0.003	Online
	5	PWR	2825	987.0	0.003	Online
	6	PWR	2825	987.0	0.003	Online
	1	PWR	2787	987.0	0.003	Online
	2	PWR	2787	987.0	0.003	Online
Tomari	1	PWR	1650	783.3	0.003	Pending
	2	PWR	1650	783.4	0.003	Pending

B

Coefficients for Antineutrino Emission Spectra of Reactor Fuels

Table B.1: Coefficients α_{pk} of the Mueller *et al.* fifth-order polynomial parameterisation of antineutrino flux from elements $k = {}^{235}\mathrm{U}$, ${}^{238}\mathrm{U}$, ${}^{239}\mathrm{Pu}$ and ${}^{241}\mathrm{Pu}$. In the column $\delta\alpha_{pk}$ the 1σ error on α_{pk} are given [78].

$k = {}^{235}\mathrm{U}$			Correlation Matrix $\rho^k_{pp'}$					
p	α_{pk}	$\delta\alpha_{pk}$	1	2	3	4	5	6
1	3.217	4.09(-2)	1.00	-0.86	0.60	0.07	-0.17	-0.14
2	-3.111	2.34(-2)	-0.86	1.00	-0.84	0.12	0.25	0.01
3	1.395	4.88(-3)	0.60	-0.84	1.00	-0.56	-0.19	0.24
4	-3.690(-1)	6.08(-4)	0.07	0.12	-0.56	1.00	-0.42	-0.14
5	4.445(-2)	7.77(-5)	-0.17	0.25	-0.19	-0.42	1.00	-0.77
6	-2.053(-3)	6.79(-6)	-0.14	0.01	0.24	-0.14	-0.77	1.00

$k = {}^{238}\mathrm{U}$			Correlation Matrix $\rho^k_{pp'}$					
p	α_{pk}	$\delta\alpha_{pk}$	1	2	3	4	5	6
1	4.833(-1)	1.24(-1)	1.00	-0.86	0.20	0.30	0.08	-0.27
2	1.927(-1)	5.86(-2)	-0.86	1.00	-0.58	-0.21	0.04	0.23
3	-1.283(-1)	1.11(-2)	0.20	-0.58	1.00	-0.48	-0.17	0.20
4	-6.762(-3)	1.92(-3)	0.30	-0.21	-0.48	1.00	-0.36	-0.20
5	2.233(-3)	2.84(-4)	0.08	0.04	-0.17	-0.36	1.00	-0.77
6	-1.536(-4)	2.86(-5)	-0.27	0.23	0.20	-0.20	-0.77	1.00

$k = {}^{239}$Pu			Correlation Matrix $\rho^k_{pp'}$					
p	α_{pk}	$\delta\alpha_{pk}$	1	2	3	4	5	6
1	6.413	4.57(-2)	1.00	-0.86	0.60	0.10	-0.17	-0.13
2	-7.432	2.85(-2)	-0.86	1.00	-0.84	0.08	0.25	-0.01
3	3.535	6.44(-3)	0.60	-0.84	1.00	-0.54	-0.20	0.26
4	-8.820(-1)	9.11(-4)	0.10	0.08	-0.54	1.00	-0.45	-0.08
5	1.025(-1)	1.38(-4)	-0.17	0.25	-0.20	-0.45	1.00	-0.79
6	-4.550(-3)	1.29(-5)	-0.13	-0.01	0.26	-0.08	-0.79	1.00

$k = {}^{241}$Pu			Correlation Matrix $\rho^k_{pp'}$					
p	α_{pk}	$\delta\alpha_{pk}$	1	2	3	4	5	6
1	3.251	4.37(-2)	1.00	0.87	-0.60	-0.08	0.17	0.13
2	-3.204	2.60(-2)	0.87	1.00	-0.84	0.11	0.25	-0.00
3	1.428	5.66(-3)	-0.60	-0.84	1.00	-0.56	-0.19	0.26
4	-3.675(-1)	7.49(-4)	-0.08	0.11	-0.56	1.00	-0.43	-0.11
5	4.254(-2)	1.02(-4)	0.17	0.25	-0.19	-0.43	1.00	-0.78
6	-1.896(-3)	9.03(-6)	0.13	0.00	0.26	-0.11	-0.78	1.00

C

Fuel Contribution Fractions

Core Type	Frac	^{235}U	^{238}U	^{239}Pu	^{241}Pu	
PWR BWR LWGR GCR	f_i	0.538	0.078	0.328	0.056	[115]
		0.614	0.074	0.274	0.038	
		0.620	0.074	0.274	0.042	
		0.584	0.068	0.298	0.050	
		0.543	0.070	0.329	0.058	
		0.607	0.074	0.277	0.042	
		0.603	0.076	0.276	0.045	
		0.606	0.074	0.277	0.043	
		0.557	0.076	0.313	0.054	
		0.606	0.074	0.274	0.046	
		0.488	0.087	0.359	0.067	[116]
		0.580	0.074	0.292	0.054	[117]
		0.544	0.075	0.318	0.063	
		0.577	0.074	0.292	0.057	
		0.590	0.070	0.290	0.050	[118]
		0.570	0.078	0.295	0.057	[119]
		0.568	0.078	0.297	0.057	[120]
		0.563	0.079	0.301	0.057	[121]
		0.650	0.070	0.240	0.040	[122]
		0.560	0.070	0.310	0.060	
		0.480	0.070	0.370	0.080	
	p_i	0.560	0.080	0.300	0.060	[123]
MOX	p_i	0.000	0.081	0.708	0.212	[123]
PHWR	p_i	0.543	0.411	0.022	0.024	[124]

Table C.1: ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu fission (f_i) or power (p_i) fractions for a range of reactor designs, compiled by [101].

References

[1] K. Abe et al. “First gadolinium loading to Super-Kamiokande”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1027 (2022), p. 166248. URL: https://www.sciencedirect.com/science/article/pii/S0168900221010883.

[2] Tanner Kaptanoglu. “Reactor Antineutrinos In The SNO+ Water Phase And Detector R&D For Large-Scale Neutrino Detectors”. PhD thesis. University of Pennsylvania, Jan. 2020.

[3] Yasuhiro Nakajima. “Recent results and future prospects from Super-Kamiokande”. The XXIX International Conference on Neutrino Physics and Astrophysics (NEUTRINO2020). June 2020. URL: https://doi.org/10.5281/zenodo.4134680.

[4] Couchot, F. et al. “Cosmological constraints on the neutrino mass including systematic uncertainties”. In: *A&A* 606 (2017), A104. URL: https://doi.org/10.1051/0004-6361/201730927.

[5] M. Aker et al. “Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN”. In: *Phys. Rev. Lett.* 123 (22 Nov. 2019), p. 221802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.221802.

[6] W. Pauli. *Open Letter to the Group of Radioactive People at the Gauverein Meeting in Tübingen*. 1930.

[7] F. Reines and C. Cowan. “The Neutrino”. In: *Nature* 178 (1956), pp. 446–449.

[8] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. “Search for Neutrinos from the Sun”. In: *Phys. Rev. Lett.* 20 (21 May 1968), pp. 1205–1209. URL: https://link.aps.org/doi/10.1103/PhysRevLett.20.1205.

[9] B. Pontecorvo. “Mesonium and Antimesonium”. In: *Journal of Experimental and Theoretical Physics* 33 (1957), pp. 549–551.

[10] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. “Remarks on the Unified Model of Elementary Particles”. In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. eprint: https://academic.oup.com/ptp/article-pdf/28/5/870/5258750/28-5-870.pdf. URL: https://doi.org/10.1143/PTP.28.870.

[11] K. Abe et al. “Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations”. In: *Nature* 580.7803 (2020). [Erratum: Nature 583, E16 (2020)], pp. 339–344. arXiv: 1910.03887 [hep-ex].

[12] Y. Fukuda et al. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. URL: https://link.aps.org/doi/10.1103/PhysRevLett.81.1562.

[13] Q. R. Ahmad et al. "Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87 (2001), p. 071301. arXiv: nucl-ex/0106015.

[14] A Yu Smirnov. "The MSW Effect and Matter Effects in Neutrino Oscillations". In: *Physica Scripta* T121 (Jan. 2005), pp. 57–64. URL: http://dx.doi.org/10.1088/0031-8949/2005/T121/008.

[15] Gianluigi Fogli and Eligio Lisi. "Evidence for the MSW effect". In: *New Journal of Physics* 6 (Oct. 2004), pp. 139–139. URL: https://doi.org/10.1088/1367-2630/6/1/139.

[16] A Faessler et al. "Search for the Cosmic Neutrino Background". In: *Journal of Physics: Conference Series* 580 (Feb. 2015), p. 012040. URL: https://doi.org/10.1088/1742-6596/580/1/012040.

[17] M. Agostini et al. "Comprehensive measurement of *pp*-chain solar neutrinos". In: *Nature* 562.7728 (2018), pp. 505–510.

[18] J. Hosaka et al. "Solar neutrino measurements in Super-Kamiokande-I". In: *Physical Review D* 73.11 (June 2006). URL: http://dx.doi.org/10.1103/PhysRevD.73.112001.

[19] J. P. Cravens et al. "Solar neutrino measurements in Super-Kamiokande-II". In: *Physical Review D* 78.3 (Aug. 2008). URL: http://dx.doi.org/10.1103/PhysRevD.78.032002.

[20] K. Abe et al. "Solar neutrino results in Super-Kamiokande-III". In: *Physical Review D* 83.5 (Mar. 2011). URL: http://dx.doi.org/10.1103/PhysRevD.83.052010.

[21] K. Abe et al. "Solar neutrino measurements in Super-Kamiokande-IV". In: *Phys. Rev. D* 94 (5 Nov. 2016), p. 052010. URL: https://link.aps.org/doi/10.1103/PhysRevD.94.052010.

[22] M. H. Ahn et al. "Measurement of neutrino oscillation by the K2K experiment". In: *Phys. Rev. D* 74 (7 Oct. 2006), p. 072003. URL: https://link.aps.org/doi/10.1103/PhysRevD.74.072003.

[23] J-PARC. *Neutrino Experimental Facility*. 2011. URL: https://j-parc.jp/Neutrino/en/nu-facility.html (visited on 08/18/2022).

[24] K. Abe et al. "The T2K experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (Dec. 2011), pp. 106–135. URL: http://dx.doi.org/10.1016/j.nima.2011.06.067.

[25] D. S. Ayres et al. *The NOvA Technical Design Report*. Tech. rep. Fermi National Accelerator Laboratory (FNAL), Oct. 2007.

[26] Babak Abi et al. *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics*. Tech. rep. Fermi National Accelerator Laboratory (FNAL), Feb. 2020. arXiv: 2002.03005 [hep-ex].

[27] K. Abe et al. *Hyper-Kamiokande Design Report*. Tech. rep. May 2018. arXiv: 1805.04163 [physics.ins-det].

[28] L. Aliaga et al. "Design, calibration, and performance of the MINERvA detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 743 (Apr. 2014), pp. 130–159. URL: http://dx.doi.org/10.1016/j.nima.2013.12.053.

[29] Justin Evans. "The MINOS Experiment: Results and Prospects". In: *Adv. High Energy Phys.* 2013 (2013), p. 182537. arXiv: 1307.0721 [hep-ex].

[30] L. N. Machado et al. "Pre-supernova Alert System for Super-Kamiokande". In: *The Astrophysical Journal* 935.1 (Aug. 2022), p. 40. URL: https://doi.org/10.3847/1538-4357/ac7f9c.

[31] C. Simpson et al. "Sensitivity of Super-Kamiokande with Gadolinium to Low Energy Anti-neutrinos from Pre-supernova Emission". In: *The Astrophysical Journal* 885 (2019). arXiv: 1908.07551.

[32] K. Asakura et al. "KamLAND Sensitivity to Neutrinos From Pre-Supernova Stars". In: *The Astrophysical Journal* 818.1 (Feb. 2016), p. 91. URL: http://dx.doi.org/10.3847/0004-637X/818/1/91.

[33] Ben Margalit and Brian D. Metzger. "Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817". In: *The Astrophysical Journal* 850.2 (Nov. 2017), p. L19. URL: http://dx.doi.org/10.3847/2041-8213/aa991c.

[34] Stan Woosley and Thomas Janka. "The physics of core-collapse supernovae". In: *Nature Physics* 1.3 (Dec. 2005), pp. 147–154. URL: http://dx.doi.org/10.1038/nphys172.

[35] A. Burrows and D. Vartanyan. "Core-collapse supernova explosion theory". In: *Nature* 589.7840 (Jan. 2021), pp. 29–39. URL: http://dx.doi.org/10.1038/s41586-020-03059-w.

[36] Ken'ichiro Nakazato et al. "Supernova Neutrino Light Curves And Spectra For Various Progenitor Stars: From Core Collapse To Proto-neutron Star Cooling". In: *The Astrophysical Journal Supplement Series* 205.1 (Feb. 2013), p. 2. URL: http://dx.doi.org/10.1088/0067-0049/205/1/2.

[37] Karolina Rozwadowska, Francesco Vissani, and Enrico Cappellaro. "On The Rate Of Core Collapse Supernovae In The Milky Way". In: *New Astronomy* 83 (Feb. 2021), p. 101498. URL: http://dx.doi.org/10.1016/j.newast.2020.101498.

[38] Scott M. Adams et al. "Observing The Next Galactic Supernova". In: *The Astrophysical Journal* 778.2 (Nov. 2013), p. 164. URL: http://dx.doi.org/10.1088/0004-637X/778/2/164.

[39] K. S. Hirata et al. "Observation in the Kamiokande-II Detector of the Neutrino Burst from Supernova SN 1987a". In: *Phys. Rev. D* 38 (1988), pp. 448–458.

[40] Todd Haines et al. "Neutrinos from SN1987a in the IMB detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 264.1 (1988), pp. 28–31. URL: https://www.sciencedirect.com/science/article/pii/0168900288910972.

[41] E.N. Alexeyev et al. "Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope". In: *Physics Letters B* 205.2 (1988), pp. 209–214. URL: https://www.sciencedirect.com/science/article/pii/0370269388916516.

[42] K. Abe et al. "Real-time supernova neutrino burst monitor at Super-Kamiokande". In: *Astroparticle Physics* 81 (2016), pp. 39–48. URL: https://www.sciencedirect.com/science/article/pii/S0927650516300457.

[43] K. Abe et al. "Diffuse supernova neutrino background search at Super-Kamiokande". In: *Phys. Rev. D* 104 (12 Dec. 2021), p. 122002. URL: https://link.aps.org/doi/10.1103/PhysRevD.104.122002.

[44] John F. Beacom and Mark R. Vagins. "Antineutrino Spectroscopy with Large Water Čerenkov Detectors". In: *Phys. Rev. Lett.* 93 (17 Oct. 2004), p. 171101. URL: https://link.aps.org/doi/10.1103/PhysRevLett.93.171101.

[45] G. Bellini et al. "Geo-neutrinos". In: *Progress in Particle and Nuclear Physics* 73 (Nov. 2013), pp. 1–34. URL: http://dx.doi.org/10.1016/j.ppnp.2013.07.001.

[46] G. Fiorentini et al. "Mantle geoneutrinos in KamLAND and Borexino". In: *Physical Review D* 86.3 (Aug. 2012). URL: http://dx.doi.org/10.1103/PhysRevD.86.033004.

[47] Virginia Strati et al. *Expected geoneutrino signal at JUNO*. Dec. 2014. arXiv: 1412.3324 [physics.geo-ph].

[48] M. G. Aartsen et al. "Constraints on Ultrahigh-Energy Cosmic-Ray Sources from a Search for Neutrinos above 10 PeV with IceCube". In: *Phys. Rev. Lett.* 117 (24 Dec. 2016), p. 241101. URL: https://link.aps.org/doi/10.1103/PhysRevLett.117.241101.

[49] M. G. Aartsen et al. "First Observation of PeV-Energy Neutrinos with IceCube". In: *Physical Review Letters* 111.2 (July 2013). URL: http://dx.doi.org/10.1103/PhysRevLett.111.021103.

[50] M. G. Aartsen et al. "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data". In: *Phys. Rev. Lett.* 113 (10 Sept. 2014), p. 101101. URL: https://link.aps.org/doi/10.1103/PhysRevLett.113.101101.

[51] Rafael Alves Batista et al. "Open Questions in Cosmic-Ray Research at Ultrahigh Energies". In: *Frontiers in Astronomy and Space Sciences* 6 (June 2019). URL: http://dx.doi.org/10.3389/fspas.2019.00023.

[52] D. Biehl et al. "Tidally disrupted stars as a possible origin of both cosmic rays and neutrinos at the highest energies." In: *Scientific Reports* 8.10828 (2018).

[53] Jonathan Arons. "Magnetars in the Metagalaxy: An Origin for Ultra–High-Energy Cosmic Rays in the Nearby Universe". In: *The Astrophysical Journal* 589.2 (June 2003), pp. 871–892. URL: https://doi.org/10.1086/374776.

[54] Luis Alfredo Anchordoqui. "Acceleration of ultrahigh-energy cosmic rays in starburst superwinds". In: *Phys. Rev. D* 97 (6 Mar. 2018), p. 063010. URL: https://link.aps.org/doi/10.1103/PhysRevD.97.063010.

[55] Mark Aartsen et al. “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 361.6398 (2018), eaat1378. eprint: https://www.science.org/doi/pdf/10.1126/science.aat1378. URL: https://www.science.org/doi/abs/10.1126/science.aat1378.

[56] M. Apollonio et al. “Search for neutrino oscillations on a long baseline at the CHOOZ nuclear power station”. In: *Eur. Phys. J. C* 27 (2003), pp. 331–374. arXiv: hep-ex/0301017.

[57] Ya. Nikitenko. “Determination of the direction to a source of antineutrinos via inverse beta decay in Double Chooz”. In: *Phys. Part. Nucl.* 47.6 (2016), pp. 1009–1013.

[58] Alessandro Strumia and Francesco Vissani. “Precise quasielastic neutrino/nucleon cross-section”. In: *Phys. Lett. B* 564 (2003), pp. 42–54. arXiv: astro-ph/0302055.

[59] S. Fukuda et al. “The Super-Kamiokande detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 501.2 (2003), pp. 418–462. URL: https://www.sciencedirect.com/science/article/pii/S016890020300425X.

[60] K. Abe et al. “Calibration of the Super-Kamiokande detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 737 (2014), pp. 253–272. URL: https://www.sciencedirect.com/science/article/pii/S016890021301646X.

[61] Michael Smy. “Low Energy Event Reconstruction and Selection in Super-Kamiokande-III”. In: *30th International Cosmic Ray Conference*. July 2007.

[62] Giada Carminati. “The new Wide-band Solar Neutrino Trigger for Super-Kamiokande”. In: *Physics Procedia* 61 (2015). 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013, pp. 666–672. URL: https://www.sciencedirect.com/science/article/pii/S1875389214006816.

[63] A. Suzuki et al. “Improvement of 20 in. diameter photomultiplier tubes”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 329.1 (1993), pp. 299–313. URL: https://www.sciencedirect.com/science/article/pii/016890029390949I.

[64] R. Brun et al. *GEANT3*. CERN Data Handling Division, 1987.

[65] S. Agostinelli et al. “Geant4 - A simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. URL: https://www.sciencedirect.com/science/article/pii/S0168900203013688.

[66] Masayuki Harada, Yusuke Koshio, and Kaito Hagiwara. “Development of Geant4 based simulation for Super-Kamiokande”. In: vol. KMI2019. 2019, p. 049.

[67] ORGANO Corporation. *Product data sheet (AMBERJET™ 1020)*. 2021.

[68] ORGANO Corporation. *Product data sheet (AMBERJET™ 4400)*. 2021.

[69] Ll. Marti et al. “Evaluation of gadolinium’s action on water Cherenkov detector systems with EGADS”. In: *Nucl. Instrum. Meth. A* 959 (2020), p. 163549. arXiv: 1908.11532 [physics.ins-det].

[70] E. Fermi et al. “Radioattività provocata da bombardamento di neutroni III”. In: *La Ricerca Scientifica* 5.1 (1934), pp. 452–453.

[71] Lise Meitner and Otto Robert Frisch. “Disintegration of Uranium by Neutrons: a New Type of Nuclear Reaction”. In: *Nature* 143.3615 (1939), pp. 239–240.

[72] United States Department of Energy. *The First Reactor*. United States Atomic Energy Commission, 1967.

[73] World Nuclear Association. *Nuclear Power Reactors*. 2021. URL: https://world-nuclear.org/information-library/nuclear-fuel-cycle/nuclear-power-reactors/nuclear-power-reactors.aspx (visited on 09/22/2021).

[74] Yukihiro Asada Akio Tsuji Akio Endoh. “Completion of ABWR Plant: Kashiwazaki-Kariwa-6, 7 reactors”. In: *Hitachi Review* 47.47 (1998), pp. 153–158.

[75] World Nuclear Association. *Nuclear Power in Japan*. 2022. URL: http://www.world-nuclear.org/information-library/country-profiles/countries-g-n/japan-nuclear-power.aspx (visited on 01/26/2022).

[76] Dong-Liang Fang, Yu-Feng Li, and Di Zhang. “Ab initio calculations of reactor antineutrino fluxes with exact lepton wave functions”. In: *Phys. Lett. B* 813 (2021), p. 136067. arXiv: 2001.01689 [hep-ph].

[77] Yu-Feng Li and Di Zhang. “New Realization of the Conversion Calculation for Reactor Antineutrino Fluxes”. In: *Phys. Rev. D* 100.5 (2019), p. 053005. arXiv: 1904.07791 [hep-ph].

[78] Th. A. Mueller et al. “Improved predictions of reactor antineutrino spectra”. In: *Phys. Rev. C* 83 (5 May 2011), p. 054615. URL: https://link.aps.org/doi/10.1103/PhysRevC.83.054615.

[79] W.Mampe et al. “The double focusing iron-core electron-spectrometer “BILL” for high resolution (n, e^-) measurements at the high flux reactor in Grenoble”. In: *Nuclear Instruments and Methods* 154 (1978), pp. 127–149.

[80] Nuclear Power. *Fuel Burnup*. https://www.nuclear-power.net/nuclear-power/reactor-physics/reactor-operation/fuel-burnup/. Accessed: 2020-01-27.

[81] International Atomic Energy Agency. *The Power Reactor Information System*. https://pris.iaea.org/pris/.

[82] X. B. Ma et al. “Improved calculation of the energy release in neutron-induced fission”. In: *Phys. Rev. C* 88 (1 July 2013), p. 014605. URL: https://link.aps.org/doi/10.1103/PhysRevC.88.014605.

[83] F.P. An et al. “The Detector System of The Daya Bay Reactor Antineutrino Experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A Accelerators Spectrometers Detectors and Associated Equipment* 811 (Aug. 2015).

[84] F. P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Phys. Rev. Lett.* 108 (2012), p. 171803. arXiv: 1203.1669 [hep-ex].

[85] J. K. Ahn et al. *RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang*. Mar. 2010. arXiv: 1003.1391 [hep-ex].

[86] Carmen Palomares. "Double-Chooz Neutrino Experiment". In: *PoS* EPS-HEP2009 (2009), p. 275. arXiv: 0911.3227 [hep-ex].

[87] Angel Abusleme et al. *JUNO Physics and Detector*. Apr. 2021. arXiv: 2104.02565 [hep-ex].

[88] G. Mention et al. "Reactor antineutrino anomaly". In: *Phys. Rev. D* 83 (7 Apr. 2011), p. 073006. URL: https://link.aps.org/doi/10.1103/PhysRevD.83.073006.

[89] A. Aguilar et al. "Evidence for neutrino oscillations from the observation of $\overline{\nu}_e$ appearance in a $\overline{\nu}_\mu$ beam". In: *Phys. Rev. D* 64 (11 Nov. 2001), p. 112007. URL: https://link.aps.org/doi/10.1103/PhysRevD.64.112007.

[90] A. A. Aguilar-Arevalo et al. "Significant Excess of Electronlike Events in the MiniBooNE Short-Baseline Neutrino Experiment". In: *Phys. Rev. Lett.* 121 (22 Nov. 2018), p. 221801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.221801.

[91] H. Almazán et al. "Improved sterile neutrino constraints from the STEREO experiment with 179 days of reactor-on data". In: *Phys. Rev. D* 102 (5 Sept. 2020), p. 052002. URL: https://link.aps.org/doi/10.1103/PhysRevD.102.052002.

[92] Mitchell Soderberg. "MicroBooNE: A New Liquid Argon Time Projection Chamber Experiment". In: *AIP Conference Proceedings* 1189.1 (2009), pp. 83–87. eprint: https://aip.scitation.org/doi/pdf/10.1063/1.3274193. URL: https://aip.scitation.org/doi/abs/10.1063/1.3274193.

[93] Maurizio Bonesini. "The Short Baseline Neutrino Program at Fermilab". In: *PoS* NuFact2021 (2022), p. 009. arXiv: 2203.05814 [hep-ex].

[94] C. Giunti et al. "Reactor antineutrino anomaly in light of recent flux model refinements". In: *Phys. Lett. B* 829 (2022), p. 137054. arXiv: 2110.06820 [hep-ph].

[95] A. Suzuki. "Present Status of KamLAND". In: *Nucl. Phys. B Proc. Suppl.* 77 (1999). Ed. by Y. Suzuki and Y. Totsuka, pp. 171–176.

[96] Ivan Esteban et al. "The fate of hints: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2020.9 (Sept. 2020). URL: http://dx.doi.org/10.1007/JHEP09(2020)178.

[97] United Nations. *Treaty on the Non-Proliferation of Nuclear Weapons*. 1970. URL: https://www.un.org/disarmament/wmd/nuclear/npt/ (visited on 01/24/2022).

[98] The Leaders of the Five Nuclear-Weapon States. *Joint Statement of the Leaders of the Five Nuclear-Weapon States on Preventing Nuclear War and Avoiding Arms Races*. Jan. 2022. URL: https://www.whitehouse.gov/briefing-room/statements-releases/2022/01/03/p5-statement-on-preventing-nuclear-war-and-avoiding-arms-races/ (visited on 01/24/2022).

[99] M. Askins et al. *The Physics and Nuclear Nonproliferation Goals of WATCHMAN: A WAter CHerenkov Monitor for ANtineutrinos*. Feb. 2015. arXiv: 1502.01132 [physics.ins-det].

[100] Marica Baldoncini et al. *Reactor antineutrinos all over the world.* http://www.fe.infn.it/antineutrino/.

[101] Marica Baldoncini et al. "Reference worldwide model for antineutrinos from reactors". In: *Phys. Rev. D* 91.6 (2015), p. 065002. arXiv: 1411.6475 [physics.ins-det].

[102] M. Tanabashi et al. "Review of Particle Physics". In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 030001. URL: https://link.aps.org/doi/10.1103/PhysRevD.98.030001.

[103] Ivan Esteban et al. "Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23}, δ_{CP}, and the mass ordering". In: *JHEP* 01 (2019), p. 106. arXiv: 1811.05487 [hep-ph].

[104] Duke Neutrino & Cosmology Group. *NUANCE Format.* http://neutrino.phy.duke.edu/nuance-format/. Originally developed by Uni. Calif., Irvine.

[105] Yuuki Nakano. "Radon background study in Super-Kamiokande". In: *J. Phys. Conf. Ser.* 888.1 (2017), p. 012191.

[106] H. Zhang et al. "Supernova Relic Neutrino search with neutron tagging at Super-Kamiokande-IV". In: *Astroparticle Physics* 60 (2015), pp. 41–46. URL: https://www.sciencedirect.com/science/article/pii/S0927650514000681.

[107] Rene Brun and Fons Rademakers. "ROOT - An Object Oriented Data Analysis Framework". In: *AIHENP'96 Workshop, Lausane.* Vol. 389. 1997, pp. 81–86.

[108] Wes McKinney. "Data Structures for Statistical Computing in Python". In: *Proceedings of the 9th Python in Science Conference.* Ed. by Stéfan van der Walt and Jarrod Millman. 2010, pp. 56–61.

[109] Charles R. Harris et al. "Array programming with NumPy". In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. URL: https://doi.org/10.1038/s41586-020-2649-2.

[110] Jim Pivarski et al. *scikit-hep/uproot: 3.12.0.* Version 3.12.0. July 2020. URL: https://doi.org/10.5281/zenodo.3952728.

[111] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272.

[112] *Water Cherenkov Machine Learning - WatChMaL.* watchmal.org. Accessed: 2022-08-04.

[113] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR).* 2016, pp. 770–778.

[114] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* KDD '16. San Francisco, California, USA: ACM, 2016, pp. 785–794. URL: http://doi.acm.org/10.1145/2939672.2939785.

[115] G. Mention et al. "The Reactor Antineutrino Anomaly". In: *Phys. Rev. D* 83 (2011), p. 073006. arXiv: 1101.2755 [hep-ex].

[116] Y. Abe et al. "Indication of Reactor $\overline{\nu}_e$ Disappearance in the Double Chooz Experiment". In: *Phys. Rev. Lett.* 108 (13 Mar. 2012), p. 131801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.108.131801.

[117] Z. Djurcic et al. "Uncertainties in the Anti-neutrino Production at Nuclear Reactors". In: *J. Phys. G* 36 (2009), p. 045002. arXiv: 0808.0747 [hep-ex].

[118] V. Kopeikin, L. Mikaelyan, and V. Sinev. "Reactor as a source of antineutrinos: Thermal fission energy". In: *Phys. Atom. Nucl.* 67 (2004), pp. 1892–1899. arXiv: hep-ph/0410100.

[119] S. Abe et al. "Precision Measurement of Neutrino Oscillation Parameters with KamLAND". In: *Phys. Rev. Lett.* 100 (2008), p. 221803. arXiv: 0801.4589 [hep-ex].

[120] K. Eguchi et al. "First results from KamLAND: Evidence for reactor anti-neutrino disappearance". In: *Phys. Rev. Lett.* 90 (2003), p. 021802. arXiv: hep-ex/0212021.

[121] T. Araki et al. "Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion". In: *Phys. Rev. Lett.* 94 (2005), p. 081801. arXiv: hep-ex/0406035.

[122] V. I. Kopeikin. "Flux and spectrum of reactor antineutrinos". In: *Phys. Atom. Nucl.* 75 (2012). [Yad. Fiz.75N2,165(2012)], pp. 143–152.

[123] G. Bellini et al. "Observation of Geo-Neutrinos". In: *Phys. Lett. B* 687 (2010), pp. 299–304. arXiv: 1003.0284 [hep-ex].

[124] G. Bellini et al. "Measurement of geo-neutrinos from 1353 days of Borexino". In: *Phys. Lett. B* 722 (2013), pp. 295–300. arXiv: 1303.2571 [hep-ex].