Measurement of the Neutrino–Proton Neutral Current Elastic Scattering Cross Section on Carbon with the T2K Off-Axis Near Detector

AND

Development of Calibration Optics for the Hyper-Kamiokande Detector

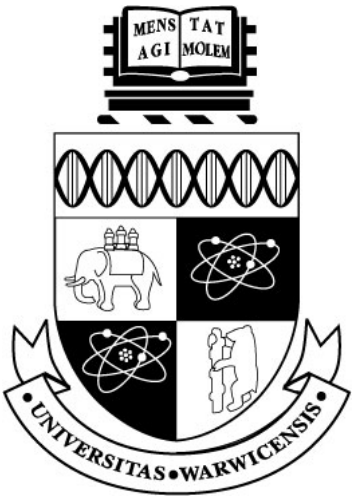

William G.S. Vinning

Department of Physics
University of Warwick

Submitted to the University of Warwick
in partial fulfilment of the requirements
for admission to the degree of
Doctor of Philosophy

May 2022

Declaration

The studentship was funded wholly by the United Kingdom's Science and Technology Facilities Council under project reference number 1917439. The material presented in this thesis has not been submitted for examination at another institute nor has it been published elsewhere. My own work is presented unless stated otherwise below.

Chapters 2 and 3 are reviews of the history and theoretical underpinnings of electroweak and neutrino physics—the work of multitudes of past and present physicists, acknowledged or otherwise.

Chapters 7 and 8 study detector data and simulation data which are the dedicated and careful work of many past and present collaborators belonging to the T2K experiment. The cross section analysis recounted in Chapter 7 was in large formulated and implemented by Dr. David Hadley. The near detector covariance matrix, used to evaluate the systematic uncertainties in the same chapter, and its interface through the HighLAND2 analysis framework was developed and validated by T2K collaborators. The near detector flux prediction was also developed by T2K collaborators and was constrained with external data from the NA61/SHINE experiment.

Chapters 12 and 13 recount the development of calibration optics and their subsequent testing within the Super-Kamiokande tank. These optics were designed and produced by the Hyper-Kamiokande research group at the University of Warwick with advice from the wider Hyper-Kamiokande calibration group. The installation of the optics and subsequent data collection from the Super-Kamiokande tank would not be possible without essential work performed by collaborators within the Hyper-Kamiokande and Super-Kamiokande calibration groups. The analysis of detector data in Chapter 13 makes use of a table of photo-sensor gain corrections and a table of photo-sensor production years produced by the Super-Kamiokande collaboration.

William G.S. Vinning
May 2022

Abstract

T2K (Tōkai–to–Kamioka) is a long-baseline (295 km), off-axis neutrino oscillation experiment that measures the appearance of electron neutrinos from a high-energy ($\sim$1 GeV) beam composed of mostly muon neutrinos. An analysis of the neutral current elastic cross section is presented using neutrino data collected at the T2K near detector facility. The signal topology is defined as neutral current interactions on the upstream fine-grained detector, FGD1, with at least one proton having $E_k > 125\,\text{MeV}$ and $\cos\theta > 0.4$ but no pions present in the final state. With a high purity (0.53) sample selected by the AdaBoost classification algorithm, the flux-averaged cross section was measured as $(3.31 \pm 0.36) \times 10^{-40}\,\text{cm}^2\,\text{nucleon}^{-1}$—in agreement with the value predicted by the nominal configuration of GENIE, which is $3.10 \times 10^{-40}\,\text{cm}^2\,\text{nucleon}^{-1}$. Compared to the same measurement performed with a cut-based signal selection, the fractional uncertainty on the cross section is reduced by 28%. Also measured are the differential cross sections with respect to the kinetic energy $d\sigma/dE_k$ and polar angle $d\sigma/d\theta$ of the primary proton within a reduced phase-space of $E_k > 181.71\,\text{MeV}$ and $\cos\theta > 0.752$. These measurements are placed in comparison with various interaction model predictions.

Hyper-Kamiokande is a next-generation, megaton-scale neutrino observatory that will succeed Super-Kamiokande as the far detector in an upgraded configuration of T2K known as T2HK. A light injection system for Hyper-Kamiokande is in development by various UK research groups to calibrate its photo-sensor responses and monitor its levels of optical scattering and absorption in-situ. A set of calibration optics developed for this system are presented. Prototypes of these devices were successfully installed at five vertical positions in the Super-Kamiokande tank in mid-2018 and have since been integrated into the detector's automated calibration routine. An analysis of detector data collected during signal tests of these optics is presented.

To my mother, father, sister,
grandmothers and grandfathers.

To Prof. Gary Barker, Dr. Steve Boyd,
Keith Jewkes, Dr. Ankush Mitra, and Dr.
David Hadley most of all, I am indebted
for their unwavering academic support
and patience throughout my time at
Warwick.

Whence nature doth forme, growth, and food impart
To all, whither their dissolv'd frames revert,
What we in reasoning the first matter call
Generative bodies, and the seeds of all,
How these first bodies include every thing,
How out of them all other bodies spring.

—TITUS LUCRETIUS CARUS, *De Rerum Natura*

Table of Contents

List of Figures

List of Tables

INTRODUCTION

Matter appears to greatly outnumber anti-matter not just locally, but at the cosmological scale also. The exact nature of the physical mechanism which birthed this asymmetry during the early universe—namely *baryogenesis*—remains to be seen, however. Per Sakharov [1], this mystery process necessarily possesses three characteristics, these are: *i)* non-conservation of baryon number, *ii)* charge conjugation (C) asymmetry and charge conjugation parity (CP) asymmetry, and *iii)* departure from thermal equilibrium. Only the first two conditions are satisfied by the Standard Model (SM) of particle physics as it currently stands—the answer likely lies beyond.

One particle of interest to the search for baryogenesis candidates is the neutrino: an electrically neutral, weakly interacting lepton with an extremely small mass. It is only during the past thirty years or so that our knowledge of neutrinos has quickly accelerated, though certainly many aspects of their nature remain unknown. A peculiar phenomenon of neutrinos is that they can shift between their three possible flavours—e, μ and τ—as they travel. Evidence for flavour oscillations in neutrinos was a landmark discovery not just for the field but for particle physics as a whole, as it confirmed that neutrinos are indeed massive rather than massless as was then assumed by the SM. Since then, piecing together the nature of these oscillations has been the collective focus of many kinds of different experiments across the globe. With the coming generation of long-baseline experiments, however, the amount of CP-violation present in neutrino oscillations, and hence in the leptonic sector, will become discoverable—a harbinger that the field has recently arrived at its *precision-era*. In 2020, the T2K collaboration announced results drawn from a decade of observations indicating some preference for maximal CP asymmetry [2]. A resolute measurement, however, will be attainable to the giant detectors of the near future—namely, Hyper-Kamiokande in Japan and DUNE in the USA. The data these detectors will accumulate over their lifetimes, due to their target masses, will not suffer from limited statistics as typical of previous oscillation

analyses; for the first time, systematic uncertainties will dominate instead. A particularly large source of systematic uncertainty faced by oscillation analyses is the modelling of neutrino interactions. Hence, the vital preparation work for these future experiments is not limited to the development of new detector hardware and techniques; validating and tuning these interaction models with increasingly precise cross section measurements is also an essential task. In one way or another, this thesis documents work undertaken on both of these fronts.

This thesis presents two independent analyses performed with two nonetheless interrelated neutrino experiments: in Part II a measurement of the neutral current elastic scattering cross section at T2K's near detector is delivered; then, in Part III, the prototyping of optical calibration hardware for the future Hyper-Kamiokande detector is documented. In their respective parts, the hardware of both detectors will be recounted for the reader. Lastly, in Part IV, the thesis will be concluded; but, first, in Part I, the history and theory behind the field of neutrino physics will be introduced.

PART I

THE NEUTRINO

CHAPTER 1

INTRODUCTION TO PART I

In preparation for the analyses delivered in Parts II and III of the thesis, the reader should first be made familiar with the neutrino and its phenomena. To this end, a brief history—spanning from the mid-20^{th} century up until the present date—of the various experimental activities surrounding the neutrino will first be offered in Chapter 2. Then, in Chapter 3, contemporary theoretical perspectives on the interactions and oscillations of neutrinos will be reported.

Introduction to Part I

CHAPTER 2

STORY OF THE NEUTRINO

2.1 Origins and Discovery

In 1930, Pauli[1] proposed a new particle [3] to resolve the apparent non-conservation of energy, momentum and spin observed in β-decay (${}^{A}_{Z}N \rightarrow {}^{A}_{Z+1}N' + e^{-} + \nu_e$) when considered merely as a two-body process—i.e. with solely the daughter nucleus and electron as the final state particles. Unlike the emissions of α and γ decay which are discrete in energy, the spectrum of β-particle energies is continuous up until a maximum value where the process ceases. Pauli's remedy particle—which would be ultra-light; have a spin of 1/2, to balance the unity spin of the nuclei and the electron's half-spin; and would be weakly interacting, having evaded detection hitherto—would share the total energy of the process with the e^{-}, thus preserving energy conservation on an event-by-event basis. Amaldi christened the new particle *neutrino*, a name which was adopted by the field under the influence of Fermi[2] and his theory of β-decay ($p \rightarrow ne^{+}\nu_e$).

The (anti-)neutrino's existence would not be confirmed until 1956, when inverse β-decay interactions ($\bar{\nu}_e + p \rightarrow e^{+} + n$) were detected around the Savannah River nuclear reactor[3] by Reines[4] and Cowan [4, 5]. The rate of this interaction was calculated prior to be tiny—even with a source of (anti-)neutrinos as abundant as a reactor. There, the field's tradition of massive detection systems was kick-started. The detector consisted of two water tanks—200 l in total and doped with 40 kg of $CdCl_2$—in alternation with three liquid scintillator tanks instrumented with photo-multiplier tubes. This set-up enabled signals from both final state

[1] Nobel Prize in Physics, 1945
[2] Nobel Prize in Physics, 1938
[3] South Carolina, USA
[4] Nobel Prize in Physics, 1995

particles to be detected in coincidence: the 511 keV photons induced by the annihilation of the e^+ and light from the neutron's capture on Cd a few μs later. In the meantime, however, Davis[5] had set out to prove with his detector next to the Brookhaven reactor[6] whether ν and $\bar{\nu}$ were different particles [6]. By finding no signs of the reaction $\bar{\nu}_e + {}^{37}\mathrm{Cl} \rightarrow e^- + {}^{37}\mathrm{Ar}$, it can be inferred that the interaction $\bar{\nu}_e + n \rightarrow e^- + p$ (that is otherwise possible for ν_e) is not physically viable. Davis, dissatisfied with the neutrino flux of the Brookhaven reactor, would later relocate to Savannah River to continue his research with a larger detector, and concluding in 1959 that ν and $\bar{\nu}$ were indeed different particles [7].

A second flavour of neutrino, that produces μ rather than e on interaction, ν_μ, was identified in 1962 by Lederman[7], Schwartz[7] and Steinberger[7] [8]. High-energy protons from the Alternating Gradient Synchrotron at Brookhaven National Laboratory were aimed at a fixed Be target to produce a source of mesons. The decay of those mesons in-flight ($\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$, $\pi^+ \rightarrow \mu^+ + \bar{\nu}_\mu$) created a beam dominated by ν_μ, that were detected downstream by an array of spark chambers. The third flavour, ν_τ, would go undetected until 2001 by the DONUT[8] experiment [9]; nonetheless, its existence was implied by the discovery of τ in 1975 [10] and the value of the invisible width of the Z^0 in 1998 [11].

2.2 Solar Neutrinos

Retiring his research at Savannah River, Davis turned his interest toward the Sun. Stars of relatively low mass, such as the Sun, create He from H mostly from a set of fusion reactions known as the proton–proton chain, as listed in Table 2.1. And four of those reactions, which occur at different rates, produce ν_e of some characteristic energy at the MeV scale as a by-product. However, only the ν_e from the reactions involving ^{7}Be and ^{8}B would be frequent and energetic enough (>0.814 MeV) to be observable from the Sun with the radio-chemical detection principles that Davis previously employed. The first detection attempt was made with a 3800 l tank of C_2Cl_4 at the Barberton Limestone Mine[9], and failed to produce any signs of neutrino-induced ^{37}Ar [12].

On the second attempt, a 378 000 l tank was built further underground within the, now famed, Homestake Mine[10]. Here, in 1968, the *solar neutrino problem* was birthed—as Davis saw just around a third of the solar flux predicted by the *standard solar model* [13]. That same

[5] Nobel Prize in Physics, 2002
[6] New York, USA
[7] Nobel Prize in Physics, 1988
[8] Fermi National Accelerator Laboratory (Fermilab), Illinois, USA
[9] Ohio, USA
[10] South Dakota, USA

Reaction	Frequency	Energy [MeV]	Name
p–p I			
$p + p \rightarrow {}^{2}\mathrm{H} + e^{+} + \nu_e$	99.75%	0.0–0.42	*pp*
$p + e^{-} + p \rightarrow {}^{2}\mathrm{H} + \nu_e$	0.25%	1.44	*pep*
${}^{2}\mathrm{H} + p \rightarrow {}^{3}\mathrm{H} + \gamma$	100%	—	
${}^{3}\mathrm{He} + {}^{3}\mathrm{He} \rightarrow {}^{4}\mathrm{He} + 2p$	85%	—	
p–p II			
${}^{3}\mathrm{He} + {}^{4}\mathrm{He} \rightarrow {}^{7}\mathrm{Be} + \gamma$	15%	—	
$e^{-} + {}^{7}\mathrm{Be} \rightarrow {}^{7}\mathrm{Li} + \nu_e$	99.99%	0.86, 0.38	${}^{7}\mathrm{Be}$
$p + {}^{7}\mathrm{Li} \rightarrow {}^{4}\mathrm{He} + {}^{4}\mathrm{He}$	100%	—	
p–p III			
$p + {}^{7}\mathrm{Be} \rightarrow {}^{8}\mathrm{B} + \gamma$	0.01%	—	
${}^{8}\mathrm{B} \rightarrow {}^{4}\mathrm{He} + {}^{4}\mathrm{He} + e^{+} + \nu_e$	100%	0–14.1	${}^{8}\mathrm{B}$

Table 2.1 Processes within the proton–proton chain and their frequency. For reactions that produce electron-neutrinos as a by-product, the common name of the process and neutrino energy are given.

deficit was seen again and again through the subsequent decades—at the radio-chemical detectors GALLEX[11] [14] and SAGE[12][15] in 1992 and 1999, and the KamiokaNDE[13] (Kamioka Nucleon Decay Experiment; Kamiokande henceforth) water Cherenkov detector in 1989 [16]—and would go unsolved until the turn of the millennium when neutrino oscillations were discovered, which will be covered in Section 2.6.

2.3 Atmospheric Neutrinos

Molecules in the Earth's upper atmosphere like O_2 or N_2 shield us from *primary cosmic rays*—high-energy hadrons (98%) and electrons (2%) that originate either from outside the Solar System or, to a lesser extent, from the Sun [17]. Most of these primaries are protons ($\sim$85%), and some are heavier nuclei like α-particles ($\sim$11%), but all produce a variety of particles called *secondary cosmic rays* when absorbed by the atmosphere. Charged mesons ($\pi^{\pm}$ and, to a lesser extent, $K^{\pm}$) form a component of those secondaries, and subsequently decay in flight ($\pi^{+} \rightarrow \mu^{+}\nu_\mu$, $\pi^{-} \rightarrow \mu^{-}\bar{\nu}_\mu$, $\mu^{+} \rightarrow e^{+}\nu_e\bar{\nu}_\mu$, $\mu^{-} \rightarrow e^{-}\bar{\nu}_e\nu_\mu$) to create the atmospheric neutrino flux. The accompanying flux of cosmic $\mu^{\pm}$ presents as a background

[11] Gran Sasso, Italy

[12] North Caucasus, Russia

[13] Kamioka, Toyama, Japan

to the measurement of the neutrinos; it is for this reason that neutrino observatories are built deep underground where they cannot penetrate.

The first observations of the atmospheric ν_μ were announced in 1965 by experiments at the Kolar Gold Field[14] and the East Rand[15] mines [18, 19], and the latter measured the ratio of fluxes $\phi\left(\nu_\mu + \bar{\nu}_\mu\right)/\phi\left(\nu_e + \bar{\nu}_e\right)$ in 1978 [20]. According to the decay processes of the secondary cosmic rays, for every ν_e ($\bar{\nu}_e$) there should be two ν_μ ($\bar{\nu}_\mu$)—however, it would later prove to be naive to assume that the same proportion would be measured at the Earth's surface. Despite that, as the measurement was fairly imprecise—owing to large cross section and flux uncertainties—the result was consistent with expectations. It was not until 1988, when Kamiokande announced a significant deficit of μ-like events [21], that the *atmospheric neutrino anomaly* transpired. Like the solar neutrino problem, the atmospheric anomaly would be solved by considering flavour transitions between the neutrino production and detection points—in this case, namely, $\nu_\mu \rightleftarrows \nu_\tau$. In the meantime however, this anomaly would be echoed by the Irvine–Michigan–Brookhaven[9] (IMB) detector in 1991 [22] and Soudan-2[16] in 1997 [23, 24], but not by the Fe calorimeters NUSEX[17] [25] and Frejus[18] [26]. As mentioned, the history of neutrino oscillations will be discussed in Section 2.6—but not before other important developments in the field prior to its discovery.

2.4 Supernova Neutrinos

If the core of a massive star ($M \gtrsim 8M_\odot$) collapses during its evolution, there is a sudden explosion of material before the star degenerates into a neutron star or a black hole; such events are called *core-collapse supernovae*, and are the most profuse source of MeV scale neutrinos of astrophysical origin. In the early stages of collapse, ν_e and neutron-rich nuclei are produced by electron–capture on the Fe nuclei that constitute the star's core—the star then passes into the *neutronisation phase* as those unstable nuclei undergo β-decay, producing $\bar{\nu}_e$ in the process. These two processes chip away at the energy of the progenitor's core and thus cause it to eventually give way under gravity. The neutrinos, however, become trapped as the core density increases. But, as the core density reaches the level of nuclear density, the repulsion from nuclear pressure decelerates the collapse and induces a shock wave that sparks the explosion. The trapped neutrinos are released, before any other kinds of radiation, in an initial intense *neutrino burst* that lasts a few tens of ms—this, in turn, damps the shock

[14] Karnataka, India
[15] Boksburg, South Africa
[16] Minnesota, USA
[17] Mont Blanc, Italy
[18] Modane Underground Laboratory, France

wave, provoking an in-fall of mass, or *accretion phase*, that lasts $\sim$100 ms and culminates in the formation of a neutron star. In the *cooling phase*, there is similar emission of all neutrino flavours as thermal processes, such as pair-production ($e^+e^- \to \nu_\ell \bar{\nu}_\ell$) and Bremsstrahlung ($e^\pm + \mathrm{N} \to e^\pm + \mathrm{N} + \nu_\ell + \bar{\nu}_\ell$, $\mathrm{N} + \mathrm{N} \to \mathrm{N} + \mathrm{N} + \nu_\ell + \bar{\nu}_\ell$), prevail. As $\sim$10 s elapse since the shock, the neutrino emissions alone will have carried away $\sim$99% of the supernova's energy, after which ν emission drops considerably.

On the 24$^{\text{th}}$ February 1987, multiple telescopes identified a type-II supernova some 50 kpc away from Earth in the Large Magellanic Cloud. The four neutrino experiments in operation at the time rushed to analyse their recently collected data in the hope to find signs of core collapse—and three of them would. Across those detectors, there was a burst of two dozen events within a 12 s time window the day before: Kamiokande-II, most of all, measured 11 neutrinos; IMB, 8; and Baksan[12], 5. The other detector in operation at the time, Mont Blanc, drew a blank. The event was labelled *SN1978A*, and is so far the only supernova observed by neutrino telescopes, and thus also, in combination with observations from optical telescopes, the only multi-messenger observation of a supernova.

As neutrinos from a supernova arrive some time before any light, detection of the former provides an early warning for the arrival of the latter. And so, since SN1987A, several neutrino telescopes around the world cooperate in the SuperNova Early Warning System project to which many optical telescopes around the world subscribe.

2.5 Historical Experiments in Electroweak Physics

A transformation of parity flips a spatial coordinate system into its mirror image. Accordingly, if a process respects parity symmetry, then the outcome of that process is the same in a mirrored universe; if not, the outcome differs, and consequently that process can be used to establish universal definitions of *left* and *right*. Around 1950, it was known that K^+ seemed to decay to both two and three pions in different modes—this was dubbed the $\tau - \theta$ *problem*. At that time, parity violation was the last thing on many physicists' minds. Though it was evident that electromagnetism, the strong interaction and gravity were parity symmetric, Wu *et. al* set out in 1957 to observe this in the weak interaction [27] as proposed by Lee and Yang[19] [28]. They would observe the angular distribution of β-particles from the natural β-decay of a sample of ^{60}Co ($^{60}\mathrm{Co} \to {}^{60}\mathrm{Co}^* + e^- + \bar{\nu}_e$) when polarised with a magnetic field. The mirror-image experiment would be repeated with the magnetic field flipped. If β-decay, and hence the weak interaction, respects parity, then the angular distribution will be simply transformed $\theta \to \theta - \pi$ under mirroring. Wu found that the e^- preferred to be emitted in the

[19] Nobel Prize in Physics, 1957

opposite spin-direction of the ^{60}Co nucleus; the violation parity was in fact maximal. And the $\tau - \theta$ *problem* was thus solved as this was a weak process.

The neutrino was discovered through its charged-current (CC) interaction; the weak neutral-current (NC) was predicted by theorists when the weak and electromagnetic interactions were unified as the *electroweak* interaction [29]. In 1973, the Gargamelle bubble chamber[20] identified one event out of 360,000 scans in which a single ∼400 MeV electron, accompanied neither by a charged lepton nor hadrons nor photons, appeared to originate in the detector material [30]. That same year, Gargamelle reported a sizeable sample of hadronic events which also appeared in the detector material without a charged lepton [31]. These two findings respectively were evidence for both leptonic ($\bar{\nu}_\mu + e^- \to \bar{\nu}_\mu + e^-$) and hadronic ($\nu_\mu + N \to \nu_\mu + X$) neutral currents.

The bosons themselves that mediate the CC and NC processes, $W^\pm$ and Z^0 respectively, were not directly observed until the UA1 p–$\bar{p}$ experiment, which operated from 1981–1993. The UA1 detector operated at CERN's proton-antiproton collider, observing the unique processes $Z^0 \to e^+e^-$ ($Z^0 \to \mu^-\mu^+$) and $W^+ \to e^+\nu_e$ ($W^+ \to \mu^+\bar{\nu}_\mu$) $W^- \to e^-\bar{\nu}_e$ ($W^- \to \mu^-\bar{\nu}_\mu$) [32]. The $W^\pm$ and Z^0 bosons were found to be very heavy ($M_{Z^0} = (91.1876 \pm 0.0021)\,\text{GeV}$, $M_{W^\pm} = (80.379 \pm 0.012)\,\text{GeV}$ [33])—a reflection of the short range of the weak interaction.

2.6 Flavour Mixing

2.6.1 Origins and discovery

Neutral particles tend to turn back and forth between their particle and anti-particle states if they are able to interact in a mode that proceeds by the non-conservation of a particular quantum number. As nature is exact in conserving charge and baryon number, this phenomenon seems to only occur in the neutral meson systems $K^0 \rightleftarrows \bar{K}^0$, $B^0 \rightleftarrows \bar{B}^0$, $B_s^0 \rightleftarrows \bar{B}_s^0$ and $D^0 \rightleftarrows \bar{D}^0$ which proceed via the weak interaction (which, in turn, fails to conserve the quark flavour quantum numbers). Gell-Mann[21] and Pais foresaw this phenomenon for the $\bar{K}^0 \rightleftarrows \bar{K}^0$ system in 1955—it was Pontecorvo who suggested the same for $\nu \rightleftarrows \bar{\nu}$ if neutrinos should possess non-zero masses [34]. There exists another kind of particle oscillation that involves, rather, transitions between particle flavours. Such flavour oscillations can be generated by distinguishing, and mixing, the flavour and mass eigenstates of a particle. This was prescribed to the neutrino sector by Maki *et al.* [35] in 1962, to which Pontecorvo speculated, in 1967, a

[20] European Organization for Nuclear Research (CERN), Geneva, Switzerland

[21] Nobel Prize in Physics, 1969

subsequent tendency to undergo $\nu_\mu \rightleftarrows \nu_e$ oscillations if leptonic charge is not conserved [36]. The formalism of such phenomena will later be covered in Section 3.3; the rest of this section will be devoted to the experimental history of neutrino oscillations.

The discovery of neutrino oscillations was two-pronged: through atmospheric neutrinos with Super-Kamiokande, thus resolving the atmospheric neutrino anomaly, and through solar neutrinos with SNO[22], thus resolving the solar neutrino problem. First, we turn to the former. In an updated analysis in 1994, its predecessor Kamiokande saw that the deficit of μ-like events was weakly dependent on the zenith angle [37]. The higher the angle subtending the ring axis of a neutrino event and the zenith axis, the further away is the production point of that neutrino and, hence, the longer the distance that neutrino has traversed before interacting. The data suggested that the up-going flux (i.e. from neutrinos that cross some distance through the Earth before interacting) was outweighed by the down-going flux, though the sample was too statistically limited for this to be convincing. The much larger successor to Kamiokande—Super-Kamiokande—verified this with greater statistics in 1998, and the up-down flux asymmetry was consistent with $\nu_\mu \rightleftarrows \nu_\tau$ transitions [38], as no corresponding excess of ν_e was seen. SNO found flavour transitions in a model-independent manner using ^{8}B neutrinos from the Sun in 2001 [39]. At ^{8}B energies ($<$14.1 MeV), the only charged current interactions are contributed by ν_e; hence this channel is sensitive to potential oscillations. The neutral current channel, on the other hand, is inclusive of all flavours and therefore the flux is insensitive to oscillations between the active neutrinos, and may be used to constrain the total flux. With these two complementary measurements, the flux due to ν_μ and ν_τ that have appeared and the ν_e that has accordingly disappeared over the solar distance can be inferred. SNO found that the neutral current flux was in agreement with that of the standard solar model and that roughly two-thirds of the solar ν_e flux had transitioned to the other flavours. These discoveries, which were awarded the 2015 Nobel Prize in Physics, opened the door to a rich new area of neutrino physics that has perhaps been its principal topic in the last twenty years: measurement of the neutrino mixing parameters.

[22] Sudbury, Ontario, Canada

2.6.2 Measurements of the mixing parameters

The flavour oscillations that affect solar and atmospheric neutrinos do not occur with the same amplitude or periodicity; they are different modes, each driven by a different neutrino mass splitting $\Delta m^2_{ij}\left(\equiv m^2_j - m^2_i\right)$ and substantiated by particular mixing angles θ_{ij}. The atmospheric mass splitting is orders of magnitude larger than that of the solar, therefore by convention that $\Delta m^2_{21} < |\Delta m^2_{32}|$, Δm^2_{atm} is identified as Δm^2_{32} and Δm^2_{sol} as Δm^2_{21}. In general, oscillations experiments are characterised by their L/E (see Section 3.3)—the ratio of traversal length L between the neutrino production point and detector, and energy E of the source neutrinos—which sets the oscillation modes which the detector is exposed to, and whether they seek to measure the relative *appearance* ($\nu_\alpha \to \nu_\beta$) or *disappearance* ($\nu_\alpha \to \nu_\alpha$) of a particular flavour supplied by the neutrino source. It follows that the oscillations observed through accelerator, reactor, atmospheric and solar neutrinos with various flavours, and at different baselines when possible, are complementary in global fits, and necessary to cover the full phase space of the mixing parameters. The current values of those parameters are given in Table 2.2 [33].

Parameter	Value of Best-fit ($\pm 1\sigma$)
Δm^2_{21} $[10^{-5}\,\mathrm{eV}^2]$	7.53 ± 0.18
Δm^2_{32} $[10^{-3}\,\mathrm{eV}^2]$	2.453 ± 0.033 (-2.536 ± 0.034)
$\sin^2\theta_{12}$	0.307 ± 0.013
$\sin^2\theta_{23}$	0.546 ± 0.021 (0.539 ± 0.022)
$\sin^2\theta_{13}$	0.0220 ± 0.0007
δ $[\pi]$	$1.36^{+0.20}_{-0.16}$

Table 2.2 Values and 1σ uncertainties of a global fit of the 3ν neutrino mixing parameters. Entries containing an un-bracketed and bracketed value correspond to the values assuming normal and inverted orderings of the neutrino masses, otherwise the parameter is independent of mass orderings. Values reproduced from [33].

Mass splitting Δm^2_{21} and mixing angle θ_{12}

Prior to SNO, the combined measurement of $\nu_e \to \nu_e$ in solar neutrinos with Super-Kamiokande, GALLEX and SAGE left ambiguity in their solutions in terms of Δm^2 and $\sin^2\theta$. The four resolutions to this puzzle included the vacuum solution ($\Delta m^2 \simeq 10^{-11}\,\mathrm{eV}^2$, $\sin^2 2\theta \simeq 0.8$) and also the three solutions inclusive of the matter effect: the small mixing angle solution (SMA; $\Delta m^2 \simeq 10^{-6}\,\mathrm{eV}^2$, $\sin^2 2\theta \simeq 10^{-4}$), the large mixing angle (LMA; $\Delta m^2 \simeq 10^{-5}\,\mathrm{eV}^2$, $\sin^2 2\theta \simeq 0.8$) and the low probability, low mass solution (LOW; $\Delta m^2 \simeq 10^{-9}\,\mathrm{eV}^2$, $\sin^2 2\theta \simeq 1.0$) [40]. SNO [39] and the long baseline reactor experiment KamLAND[13] [41] ($L/E \simeq \frac{100\,\mathrm{km}}{1\,\mathrm{MeV}}$) would later provide decisive evidence that the LMA solution was the correct one. Today, the most precise measurements of Δm^2_{21} are contributed by KamLAND for antineutrinos and Super-Kamiokande for neutrinos.

Both solar observatories and KamLAND are subject to the survival probability function

$$P_{\nu_e\to\nu_e} = \cos^4\theta_{13} P^{2\nu}_{\nu_e\to\nu_e} + \sin^4\theta_{13} \tag{2.1}$$

where $P^{2\nu}$ is the two neutrino survival probability function

$$P^{2\nu}_{\nu_e\to\nu_e} = 1 - \sin^2 2\theta_{12M} \sin^2 \frac{\Delta m^2_{21M} L}{4E} \tag{2.2}$$

where Δm^2_{21M} and θ_{12M} are the matter-modified solar mixing parameters.

Mass splitting Δm^2_{32}, mixing angle θ_{23} and the mass hierarchy

The atmospheric parameters, Δm^2_{32} and θ_{23}, are extracted at atmospheric or long baseline accelerator experiments either from a combined study of the $\nu_\mu \to \nu_\mu$, $\bar\nu_\mu \to \bar\nu_\mu$, $\nu_\mu \to \nu_e$ and $\bar\nu_\mu \to \bar\nu_e$ channels or solely from $\nu_\mu \to \nu_\tau$ when neutrino high energies are available—though examples of the latter method are limited to the OPERA[11] experiment [42].

The ν_μ survival probability function in vacuum takes the form [43]

$$P_{\nu_\mu\to\nu_\mu} \approx 1 - 4\cos^2\theta_{13}\sin^2\theta_{23}\left(1 - \cos^2\theta_{13}\sin^2\theta_{23}\right)\sin^2\frac{\Delta m^2_{32} L}{4E} \tag{2.3}$$

and the ν_μ to ν_e transition probability function takes the form [44]

$$P_{\nu_\mu \to \nu_e} \approx \sin^2\theta_{23} \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2_{31} L}{4E} \mp \frac{\sin 2\theta_{12} \sin 2\theta_{23}}{2 \sin\theta_{13}} \sin \frac{\Delta m^2_{21} L}{4E} \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2_{31} L}{4E} \sin\delta \tag{2.4}$$

where $\mp$ evaluates to $-$ for neutrinos and $+$ for antineutrinos.

The most precise constraints of $|\Delta m^2_{32}|$ and θ_{23} are contributed by the accelerator experiments MINOS[8] [45], T2K[23] [2] and NOvA[8] [46] and the atmospheric observatories Super-Kamiokande[13] [47] and IceCube[24] [48]. At present however, there are two ambiguities in this sector: the sign of Δm^2_{32}, i.e. whether the neutrino mass ordering is normal ($\Delta m^2_{32} > 0$) or inverted ($\Delta m^2_{32} < 0$), and the octant of θ_{23}, i.e. whether $\sin^2\theta_{23} < 0.5$ (lower octant) or $\sin^2\theta_{23} > 0.5$ (upper octant). The two mass orderings can be discriminated in very long baseline ($\sim$1000 km) experiments by the behaviour of the matter effect in the $\nu_\mu \to \nu_e$ and $\bar{\nu}_\mu \to \bar{\nu}_e$ channels. Given normal ordering, the matter effect enhances the appearance of ν_e but suppresses the appearance $\bar{\nu}_e$, whereas the inverse is true for inverted ordering. At present, however, there is some preference towards normal ordering and θ_{23} belonging to the upper octant based on global data [49].

Mixing angle θ_{13}

Sensitivity to the mixing angle θ_{13} is available at the first maximum of oscillations driven by Δm^2_{32} ($L/E \sim 1\,\mathrm{km\,GeV^{-1}}$). Therefore, θ_{13} is possible to be measured with either reactor neutrino experiments at medium baselines ($L/E \simeq \frac{1\,\mathrm{km}}{1\,\mathrm{GeV}}$), in searches for $\bar{\nu}_e \to \bar{\nu}_e$ with the survival probability function [17]

$$P_{\bar{\nu}_e \to \bar{\nu}_e} \approx 1 - \sin^2\theta_{13} \frac{\Delta m^2_{13} L}{4E} - \sin^2\theta_{12} \cos^4\theta_{13} \sin^2 \frac{\Delta m^2_{12} L}{4E} \tag{2.5}$$

or with accelerator experiments at long baselines $L/E \simeq \frac{10^4\,\mathrm{km}}{1\,\mathrm{GeV}} \simeq 10^4\,\mathrm{km\,GeV^{-1}}$, in searches for $\nu_\mu \to \nu_e$ with the transition probability function given in Eq. (2.4).

Prior to a ν_e appearance study by T2K[23] in 2011 [50], the searches for these channels with CHOOZ[25] [51] and MINOS[8] [52] were inconclusive. As demonstrated in Eq. (3.51), charge–parity asymmetry is allowed only when all three mixing angles are non-zero; thus,

[23] Japan Proton Accelerator Research Complex (J-PARC), Ibaraki, Japan
[24] Amundsen–Scott South Pole Station, South Pole
[25] Ardennes, France

given that θ_{12} and θ_{23} were already known at that time to be substantial, strict charge–parity conservation in the neutrino sector was implied by these null results. In the years following, however, it would be uncovered by precision measurements undertaken by the reactor experiments Daya Bay[26] [53], RENO[27] [54] and Double Chooz[25] [55] that, though smaller than the other angles by an order of magnitude, θ_{13} would be larger than suspected.

CP violating phase δ

The *CP*-violating phase δ is accessible through the $\nu_\mu \to \nu_e$, $\bar{\nu}_\mu \to \bar{\nu}_e$ channels driven by Δm^2_{32}, as seen in Eq. (2.4). A precision measurement of δ, however, requires sufficient constraints on all other mixing parameters; therefore, the most feasible method of extracting δ is the simultaneous study of all the oscillation channels available at long baseline accelerator experiments ($\nu_\mu \to \nu_e$, $\bar{\nu}_\mu \to \bar{\nu}_e$, $\nu_\mu \to \nu_\mu$, $\bar{\nu}_\mu \to \bar{\nu}_\mu$) with external constraints placed on the solar mixing parameters and θ_{13}. The best constraints of δ come from NOvA[8] [46] in 2019 and T2K[23] [2] in 2020, which both indicate a preference for maximal *CP* violation ($\delta \sim -\frac{\pi}{2}$) in either mass ordering scenario. Both of these measurements are statistically limited however and fall short of excluding *CP* conserving values of δ to 3σ.

2.6.3 Searches for sterile flavours

In nature, unlike all other leptons, it seems neutrinos are found only in left-handed chiral states; Pontecorvo, again, was the first to speculate upon their right-handed counterparts—otherwise known as *sterile* neutrinos [36]. Sterile flavours are stripped of the coupling to the $W^\pm$ and Z^0 bosons that the *active*, left-handed flavours possess, and thus make their presence known strictly either through gravity or flavour oscillations. The minimal extension to standard 3ν mixing is the inclusion of a single sterile flavour—known as the *3+1 model* [56]. In terms of flavour oscillations, a fourth mass eigenstate m_4 introduces three additional squared splittings to the mass hierarchy (Δm^2_{41}, Δm^2_{42}, Δm^2_{43})—which may or may not be virtually degenerate, depending on the mass of m_4—and three additional mixing angles to the mixing matrix ($\theta_{14}, \theta_{24}, \theta_{34}$). Hence such an extension adds new appearance and disappearance modes for the active flavours which, uniquely, would distort spectra in both the charged current and neutral current interaction channels. Sterile neutrinos of varying mass scales hold numerous implications for cosmology, but this aspect will be overlooked for the sake of brevity.

[26] Shenzhen, China
[27] Jeollanam-do, South Korea

Several anomalies spotted at short baselines ($L/E \sim 1\,\mathrm{eV}^{-2}$) have been interpreted as evidence for a light ($m_4 \sim 1$ eV) sterile neutrino—namely, *i)* the accelerator anomalies, which comprises the apparent $\bar{\nu}_\mu \to \bar{\nu}_e$ transitions at LSND[28] (at 3.8σ) [57, 58] and in both neutrino and anti-neutrino mode (at 4.8σ) at MiniBooNE[8] [59]; *ii)* the *reactor anomaly*, which comprises apparent $\bar{\nu}_e \to \bar{\nu}_e$ transitions (at 3.0σ globally) at reactor experiments with baselines around 10–100 m [60]; and *iii)* the *gallium anomaly*, which comprises apparent $\nu_e \to \nu_e$ transitions (at 2.3σ) at GALLEX [61] and SAGE [62] when exposed to ^{37}Ar and ^{51}Cr calibration sources at a baseline of $\sim$1 m [63]. The accelerator excesses, confusingly, are not strictly compatible with each other, as they occur at different values of L/E; to put this all to rest, the Short Baseline Neutrino (SBN) program at Fermilab is an experiment to begin in the near future that will exclude 99% of the oscillation phase space indicated by LSND to 5σ [64].

2.7 Other Neutrino Sources

Radioactive material in the Earth—mostly ^{238}U, ^{232}Th and ^{40}K—decay to produce about half of all geothermal heat—the heat that drives plate tectonics, mantle convection and other geophysical processes. *Geo-neutrinos* are MeV scale neutrinos that are a by-product of these decays. Thus the geo-neutrino flux offers an indirect means to observe radiogenic heating and carries crucial insights into the Earth's thermal history and radioactive composition. KamLAND[13] [65] and Borexino[11] [66] are detectors that observe geo-neutrinos.

One second after the Big Bang, the universe had cooled enough such that neutrinos decoupled from matter. These neutrinos—known as *relic neutrinos*—persist today and form the yet-to-be-discovered cosmic neutrino background. Although extremely abundant, the energies of relic neutrinos are so low (about 1 μeV to 100 μeV) to be rendered practically invisible to current detection systems. PTOLEMY [67] is a future high-resolution detector proposed to directly detect the cosmic neutrino background.

[28] Los Alamos National Laboratory, New Mexico, USA

2.8 Future Prospects

At present, a few key questions about the basic nature of neutrinos remain unanswered. First of all, it is undetermined whether the neutrino is its own anti-particle—-that is, whether neutrinos are Majorana or Dirac type fermions. Experiments that seek this answer search for neutrinoless double-beta decay ($(A,Z) \rightarrow (A,Z+2)+2e^{-}$), a reaction unique to Majorana neutrinos that violates lepton number conservation by two units. Some detectors placing ever stringent limits on the half-life of this process include KamLAND-Zen[13] [68], SNO+[22] [69] and SuperNEMO[18] [70]. As well, not only are the absolute neutrino masses unknown but neither their ordering nor their mass scale. While their ordering is exposed by the matter effect in long-baseline oscillation experiments, the mass scale can be established by high-precision measurements of the beta-decay spectrum. KATRIN[29] [71] is an ongoing experiment studying the decay of tritium that can establish a bound of 0.2 eV on the electron neutrino mass within three years of exposure.

2.8.1 Flavour Mixing

Oscillation physics has entered an era of precision physics and the study of perhaps two of its most crucial topics are accessible within the next generation of experiments, these being the observation of *CP* symmetry, or asymmetry, and determination of the neutrino mass ordering. These feats will be performed by the long baseline accelerator experiments Hyper-Kamiokande[13] and the Deep Underground Neutrino Experiment (DUNE)[8] [72]; the technical design and physics prospects belonging to the former will be covered in detail later in Chapter 11.

At a very long baseline 1300 km—stretching from Chicago to South Dakota—DUNE will exploit the matter effect to boost sensitivity to the mass hierarchy. Its four far detector modules, which will together amass 70 kt, will make use of the excellent tracking capabilities of state-of-the-art liquid argon technology. With a nominal beam power of 1.2 MW, DUNE is projected to discover the mass hierarchy and exclude 50% of the phase space of δ to 3σ within five years of exposure. Initial construction efforts began in 2019, and data-taking is expected to begin in the early 2030s.

[29] Karlsruhe, Germany

CHAPTER 3

THEORY OF THE NEUTRINO

3.1 The Standard Model and the Electroweak Interaction

The SM unites the standard quantum field descriptions of the electroweak force and the strong force—namely, Glashow–Weinberg-Salam (GWS) theory and Quantum Chromodynamics (QCD)—in a singular framework. Accordingly, the gauge symmetry group associated with the SM takes the form $\mathrm{SU}(3)\otimes\mathrm{SU}(2)\otimes\mathrm{U}(1)$ where SU(3) is the colour-charge group, SU(2) is the weak isospin group and U(1) is the hypercharge group. This section will henceforth focus on the electroweak force—corresponding to the $\mathrm{SU}(2)\otimes\mathrm{U}(1)$ sector of the SM group—which is, in turn, a unified description of the weak and electromagnetic forces.

The weak interaction is set apart from the other interactions in several ways: it is flavour changing, parity-violating and couples only to the left-handed fermions. The electroweak theory of Glashow[1], Weinberg[1] and Salam[1] is hence a chiral theory in which the left-handed fermions transform as doublets

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix}_L, \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}_L, \quad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}_L, \quad \begin{pmatrix} u \\ d' \end{pmatrix}_L, \quad \begin{pmatrix} c \\ s' \end{pmatrix}_L, \quad \begin{pmatrix} t \\ b' \end{pmatrix}_L \tag{3.1}$$

and the right-handed fermions transform as singlets

$$e_R, \quad \mu_R, \quad \tau_R, \quad u_R, \quad d_R, \quad c_R, \quad s_R, \quad t_R, \quad b_R \tag{3.2}$$

[1] Nobel Prize in Physics, 1979

under arbitrary SU(2) rotations. The up-type quarks u,c,t are collected with the flavour eigenstate d',s',b' of their corresponding down-type quark, which are distinct from their mass eigenstates d,s,b. The mass and flavour eigenstates of quarks are related by the Cabibbo-Kobayashi-Maskawa (CKM) mixing matrix V according to

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{3.3}$$

The weak isospin group grants a triplet of vector bosons W^i_μ which couple to the weak isospin currents J^i_μ with strength g, where $i \in \{1,2,3\}$, and the weak hypercharge group grants a singlet vector boson B_μ which couples to the weak hypercharge current j^Y_μ with strength g'. The electroweak interaction thus takes the form

$$\mathscr{L}_{int} = -ig\left(J^i\right)^\mu W^i_\mu - ig'\left(j^Y\right)^\mu B_\mu. \tag{3.4}$$

Spontaneous symmetry breaking via the Higgs mechanism generates the boson masses and brings about the four following physical gauge fields. The weak isospin fields W^1_μ and W^2_μ combine to form the two massive charged bosons W^+ and W^-

$$W^\pm_\mu = \frac{1}{\sqrt{2}}\left(W^1_\mu \mp iW^2_\mu\right) \tag{3.5}$$

and B_μ and the third isopin field W^3_μ mix according to the Weinberg angle θ_W to form the massive Z^0

$$Z_\mu = \cos\theta_W W^3_\mu - \sin\theta_W B_\mu \tag{3.6}$$

and the massless A_μ

$$A_\mu = \cos\theta_W W^3_\mu + \sin\theta_W B_\mu \tag{3.7}$$

which is otherwise identified as the photon field. The $W^\pm$ and Z^0 masses are very large; it is this fact that restricts the range of the weak interaction and suppresses its strength at low energy scales.

The electroweak interaction thus includes the leptonic charged current interactions

$$J^{\mu}_{CC} = \bar{\nu}_{eL}\gamma^{\mu}e_L$$
$$= \bar{\nu}_e\gamma^{\mu}[(1-\gamma^5)/2]e$$

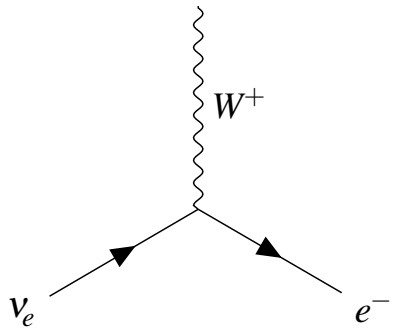

and

$$(J^{\mu}_{CC})^{\dagger} = \bar{e}_L\gamma^{\mu}\nu_{eL}$$
$$= \bar{e}\gamma^{\mu}[(1-\gamma^5)/2]\nu_e$$

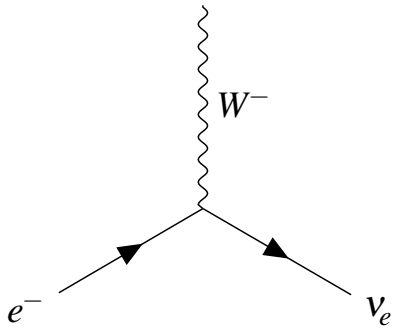

and the weak neutral current interaction

$$J^{\mu}_{NC} = \frac{1}{2}\left(\bar{\nu}_{eL}\gamma^{\mu}\nu_{eL} - \bar{e}\gamma^{\mu}e_L\right) - \sin^2\theta_W J^{\mu}_{EM}$$

where

$$J^{\mu}_{EM} = -\bar{e}_L\gamma^{\mu}e_L - \bar{e}_R\gamma^{\mu}e_R = -\bar{e}\gamma^{\mu}e.$$

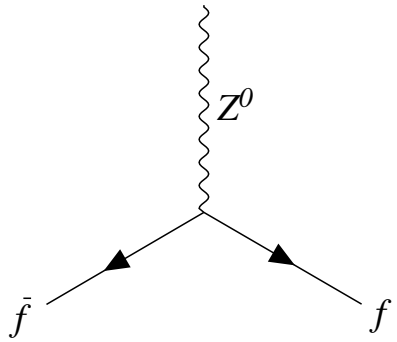

Apparent in the CC interaction is the *vector minus axial-vector*, or $V-A$, structure which is parity-violating by nature. All in all, GWS theory has three free parameters: the fine structure constant α, the Weinberg θ_W and the Fermi constant G_F.

3.2 Neutrino Interactions

The cross sections of neutrino–matter interactions are of particular interest to the field as they are often a significant source of uncertainty for other physics analyses, though they are interesting to study in their own right. The interactions of neutrinos with leptons and nucleons will be briefly recounted.

3.2.1 Neutrino–lepton interactions

Neutrino–electron scattering is a fundamental reaction to electroweak theory; both being point particles, they are also rather simple processes to consider at the tree level. The elastic scattering reactions $\nu_\ell e^- \to \nu_\ell e^-$ and $\bar{\nu}_\ell e^- \to \bar{\nu}_\ell e^-$ depicted in Fig. 3.1, which proceed exclusively through the NC if $\ell \in \{\mu, \tau\}$ or through either the CC or NC if $\ell \equiv e$.

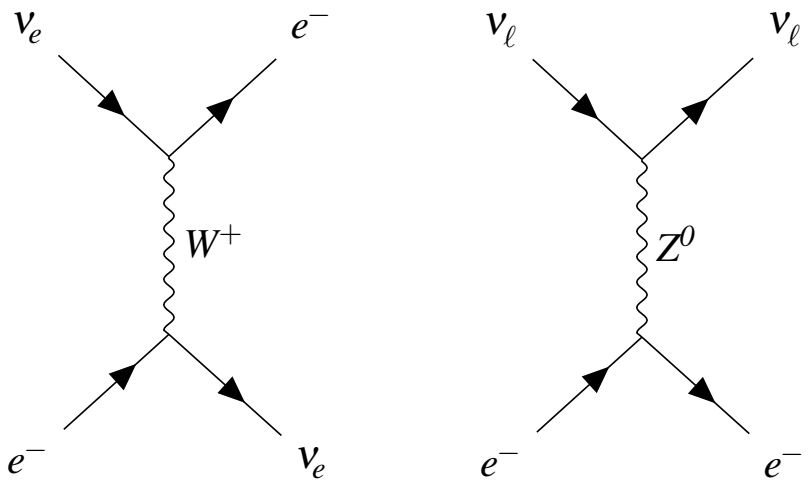

(a) Charged current mode. (b) Neutral current mode.

Fig. 3.1 Feynman diagrams for the neutrino–electron elastic scattering reactions in charged current (Fig. 3.1a) and neutral current modes (Fig. 3.1b).

Following Feynman rules, the invariant matrix element of the CC process takes the following form in the case of free electrons [73]

$$\mathscr{M}_{CC} = \frac{G_F}{\sqrt{2}} \left[\bar{e}\gamma_\mu(1-\gamma^5)\nu_\ell\right] \left[\bar{\nu}_\ell\gamma^\mu(1-\gamma^5)e\right] \tag{3.8}$$

where G_F is the Fermi constant (equal to $(1.166\,37 \pm 0.000\,01) \times 10^{-5}\,\mathrm{GeV}^{-2}$ [33]) and e, ν_ℓ, $\bar{e}$, and $\bar{\nu}_\ell$ are the fields of the electron, the neutrino and their anti-particle forms. Using this, the differential cross section with respect to the inelasticity y—or, the fractional energy transferred to the electron—of the same process is given by

$$\frac{d\sigma(\nu_\ell e^- \to \nu_\ell e^-)_{CC}}{dy} = \frac{2m_e G_F^2 E_\nu}{\pi} \tag{3.9}$$

where m_e is the mass of the electron, E_ν is the incident neutrino energy and y is defined by

$$y = 1 - \frac{E_e'}{E_e} \tag{3.10}$$

where E_e and E_e' are the initial and final state energies of the electron [74].

The invariant matrix element corresponding to the NC process is given by

$$\mathscr{M}_{NC} = \frac{G_F}{\sqrt{2}} \left[\nu_\ell \gamma^\mu (1-\gamma^5)\nu_\ell\right] \left[\bar{e}\gamma_\mu (g_V - g_A\gamma^5)e\right] \tag{3.11}$$

where g_V and g_A are the vector and axial-vector coupling constants. The differential cross section for the same process is given by

$$\frac{d\sigma(\nu_\ell e^- \to \nu_\ell e^-)_{NC}}{dy} = \frac{m_e G_F^2 E_\nu}{2\pi} \left((g_V \pm g_A)^2 + (g_V \mp g_A)^2 (1-y)^2 - (g_V^2 - g_A^2)\frac{m_e y}{E_\nu} \right) \tag{3.12}$$

where $\pm$ evaluates to $(+1)$ for neutrinos and (-1) for antineutrinos.

Observing said neutrino–lepton scatterings prove to be difficult in practice, owing to their tiny cross sections ($\sim 10^{-42}\,\mathrm{cm}^2$) at typical beam energies ($\sim$ GeV)—which are about an order of a magnitude less than that of neutrino–nucleon reactions.

3.2.2 Neutrino–nucleon interactions

The physics of interactions with extended objects, such as nucleons, is immeasurably more complicated than that with point-like particles; the internal structure of nuclei is modelled by the presence of form factors within the hadronic current.

Neutrino–nucleon scatterings are readily described by GWS theory; however, nuclear medium effects—significant in the heavy targets typical of neutrino experiments—are poorly understood and notoriously difficult to model. Hadrons in the final-state of some primary reactions more often than not also undergo interactions within the nucleus: for example, a π produced from some primary neutrino–nucleon interaction can scatter, be absorbed, undergo charge exchange or interact to emit further particles with the other nucleons in the nuclear medium. Such events are known collectively as Final State Interactions (FSI).

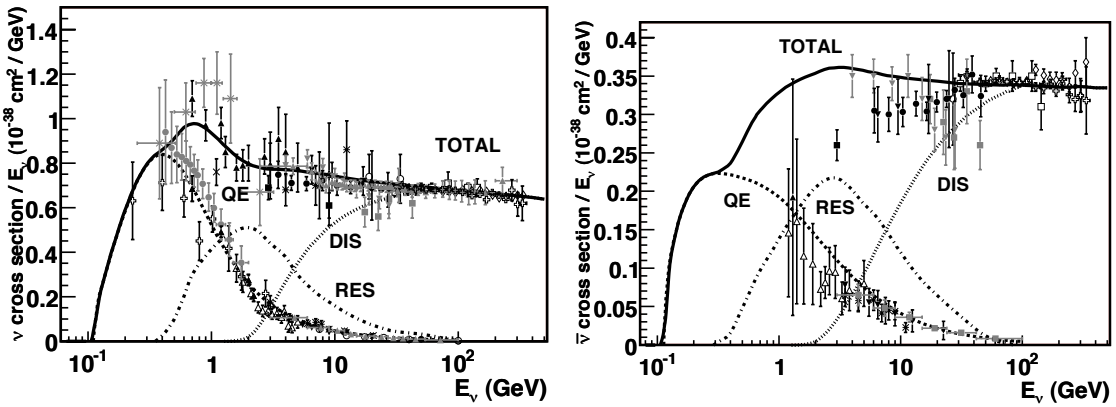

(a) For neutrinos. (b) For anti-neutrinos.

Fig. 3.2 Total neutrino–nucleon and anti-neutrino–nucleon charged current cross sections per nucleon as a function of reconstructed neutrino energy. Reproduced from Ref. [73].

Charged current quasi-elastic scattering

Charged current quasi-elastic (CCQE) scattering

$$\nu_\ell + n \to \ell^- + p, \qquad \bar{\nu}_\ell + p \to \ell^+ + n \tag{3.13}$$

is the most probable neutrino–nucleon interaction mechanism at neutrino energies around and below 1 GeV. The term *quasi-elastic* is used to describe these processes as they neither proceed entirely elastically, as there is a mass transfer (albeit small compared to the momentum transfer q^2 of the process) when the nucleon type changes, or inelastically, as the nucleon remains intact in the final state. These processes in neutrino and anti-neutrino mode are depicted in Fig. 3.3. As a reaction that dominates at many neutrino beam energies, in which both final-state particles are, in principle, reconstructable and the outgoing lepton gives the flavour of the incident neutrino, this channel is naturally the signal for most oscillation experiments.

The CCQE differential cross section, per Llewellyn-Smith [75], assumes the following form in the case of free nucleons

$$\frac{d\sigma}{d|q^2|} = \frac{G_F^2 M^2}{8\pi E_\nu^2}\left(A(q^2) \pm \frac{s-u}{M^2}B(q^2) + \frac{(s-u)^2}{M^4}C(q^2)\right) \tag{3.14}$$

where $\pm$ evaluates to -1 for neutrinos and $+1$ for antineutrinos, M is the nucleon mass, m is the lepton mass, E_ν is the neutrino energy, and s and u are Mandelstam variables. The factors

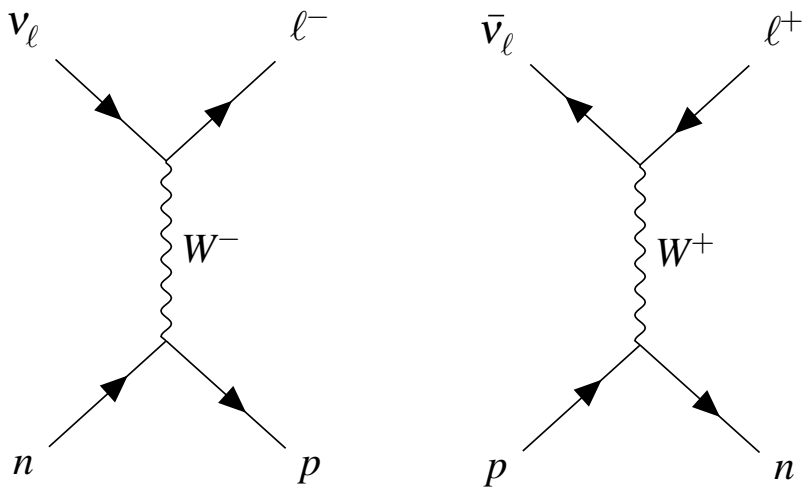

(a) Neutrino mode. (b) Antineutrino mode.

Fig. 3.3 Feynman diagrams of the quasi-elastic charge current in neutrino mode (Fig. 3.3a) and antineutrino mode (Fig. 3.3b).

A, B and C are functions of the four form factors of the nucleon: these being, the two vector form factors, F_1 and F_2, which are found from electron–nucleon scattering; the axial-vector form factor, F_A; and the pseudo-scalar form factor, F_P, which may be expressed in terms of F_A. The factors A, B and C are defined as [73]

$$A = \frac{m^2 - q^2}{M^2}\left[(1+\eta)F_A^2 - (1-\eta)F_1^2 + \eta(1-\eta)F_2^2 + 4\eta F_1 F_2 - \frac{m^2}{4M^2}\left((F_1+F_2)^2 + (F_A+2F_P)^2 + \left(\frac{q^2}{M^2}+4\right)F_P^2\right)\right], \tag{3.15}$$

$$B = -\frac{q^2}{M^2}F_A(F_1+F_2), \tag{3.16}$$

$$C = \frac{1}{4}\left(F_A^2 + F_1^2 + \eta F_2^2\right) \tag{3.17}$$

where

$$\eta = -q^2/4M^2. \tag{3.18}$$

The axial form factor F_A is commonly assumed to take the following *dipole* form at low q^2

$$F_A(q^2) = \frac{g_A}{\left(1 - q^2/M_A^2\right)^2} \tag{3.19}$$

where g_A is the axial coupling constant—determined from β decay at $q^2 \simeq 0$—and M_A is the axial mass—determined from either CCQE neutrino–nucleon scattering or pion electro-production.

The value of the axial mass remains topical. The combined CCQE measurements performed by deuterium target experiments achieve an average M_A of (1.016 ± 0.026) GeV, to which the average of pion electro-production measurements lies in good agreement at (1.014 ± 0.016) GeV [76]. However, more recent measurements of CCQE on heavy targets favour higher values of M_A. For example, MiniBooNE has reported an effective value of (1.35 ± 0.17) GeV [77] with which subsequent measurements made at K2K [78], MINOS [79] and T2K [80] find agreement. These results, which may be biased by nuclear effects, have introduced into the field a particular focus on the tuning and selection of nuclear interaction models. For example, models of the *2p-2h* (two-particle, two-hole) mechanism [81, 82], which account for interactions of neutrinos with multiple correlated nucleons, are of particular interest. Such reactions ultimately result in two final state nucleons—unlike tree-level processes which can otherwise be labelled as *1p-1h*. In the event that only one of these nucleons can be reconstructed—owing either to final state interactions or to insufficient data—these interactions can easily fake a CCQE signature. Factoring in 2p-2h, along with Random Phase Approximation [83] corrections, the anomalous MiniBooNE M_A results can ultimately be brought into agreement with that of the deuterium experiments.

Neutral current elastic scattering

Neutral current elastic (NCE) scattering

$$\nu_\ell + n \to \nu_\ell + n, \qquad \bar{\nu}_\ell + n \to \bar{\nu}_\ell + n, \tag{3.20}$$

$$\nu_\ell + p \to \nu_\ell + p, \qquad \bar{\nu}_\ell + p \to \bar{\nu}_\ell + p \tag{3.21}$$

is about the simplest way a neutrino can interact with a nucleon. These processes are depicted in Fig. 3.4. Given that only the recoil nucleon is, in principle, detectable, whether the incident particle undergoing NCE scattering is a neutrino or anti-neutrino is experimentally indistinguishable.

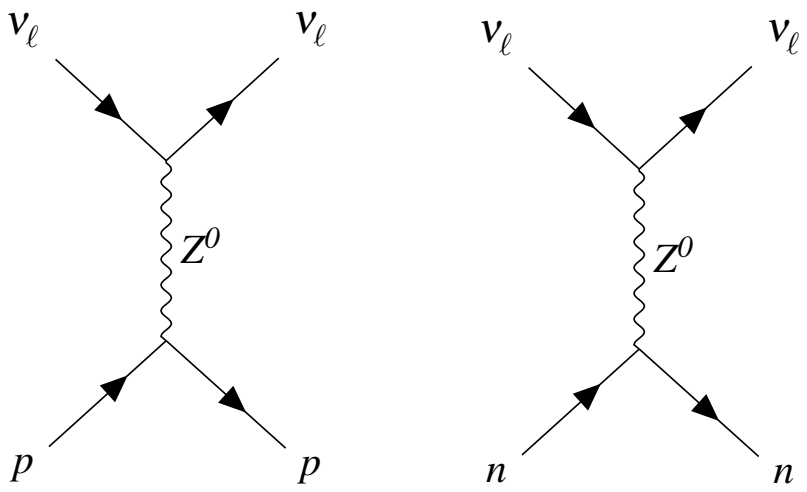

(a) On proton. (b) On neutron.

Fig. 3.4 Feynman diagrams of the neutral current interaction of neutrino–protons (Fig. 3.4a) and neutrino–neutron (Fig. 3.4b).

This channel has previously been studied at the following experiments: firstly in 1976 by the Columbia-Illinois-Rockefeller [84] (CIR) and Harvard-Pennsylvania-Wisconsin [85] (HPW) collaborations at BNL, at CERN's Gargamelle [86], at BNL-E734 [87, 88], at MiniBoone [89, 90] and at T2K [91, 92]. More details on these measurements can be found in Table 3.1.

The differential cross section for NCE [93] upon free nucleons may be written as

$$\frac{d\sigma}{d|q^2|} = \frac{G_F^2}{2\pi}\frac{|q^2|}{E_\nu^2}\left(A(q^2) \pm By + C(q^2)y^2\right) \tag{3.22}$$

where

$$y = \frac{4E_\nu}{M} + \frac{q^2}{M^2} \tag{3.23}$$

and where A, B and C are functions of the nucleon form factors, where

$$A(q^2) \equiv \frac{1}{4}\left[F_A^2\left(1-\frac{q^2}{4M^2}\right) - \left(F_1^2 + \frac{q^2}{4M^2}F_2^2\right)\left(1+\frac{q^2}{4M^2}\right) - F_1F_2\frac{q^2}{M^2}\right], \tag{3.24}$$

$$B(q^2) \equiv -\frac{1}{4}F_A\left(F_1+F_2\right), \tag{3.25}$$

$$C(q^2) \equiv \frac{M^2}{16q^2}\left[F_2^2\frac{q^2}{4M^2} - F_A^2 - F_1^2\right]. \tag{3.26}$$

In addition to that associated with CC, the neutral weak current axial form factor F_A contains a contribution from the strange quark as follows

$$F_A(q^2) = \pm\frac{1}{2}\frac{F_A\left(q^2=0\right)}{\left(1-q^2/M_A\right)^2} - \frac{1}{2}F_A^s(q^2) \tag{3.27}$$

where $\pm$ evaluates to $+1$ for protons and -1 for neutrons; $F_A^s(q^2)$ is the strange axial form factor, which is commonly assumed to take the dipole form as follows

$$F_A^s(q^2) = \frac{\Delta s}{\left(1-q^2/M_A^2\right)^2} \tag{3.28}$$

where Δs is the strange quark contribution to the nucleon spin—a valuable quantity to be known in the age of the *proton spin crisis*. The proton spin crisis is an outstanding problem in physics, whereby the spins given by the up and down valence quarks were found by the EMC experiment [94] to provide subdominant contributions to the overall proton spin, implying that contributions by sea quarks or exchange gluons are significant. Results from BNL-E734 have suggested that $\Delta s < 0$ at 67% CL [87, 88], whilst MiniBooNE have reported $\Delta s = 0.08 \pm 0.26$ from studying the ratio $\sigma(\nu_\mu p \to \nu_\mu p)/\sigma(\nu_\mu n \to \mu^- p)$ [89]. One global analysis suggests that $\Delta s \approx -0.15$ [95] albeit with significant extrapolation uncertainties. Needless to say, the measurements at present are ambiguous and therefore any further efforts to measure ν–N NCE scatterings would be welcomed by the field.

Experiment	Target	Signal	Neutrino Energy	Measurement	Reference(s)
CIR	CH_2	p track	1 GeV to 2 GeV	$\sigma(\nu_\mu p \to \nu_\mu p)/\sigma(\nu_\mu n \to \mu^- p)$	[84]
HPW	C, Al	—	—	—	[85]
Gargamelle	CF_3Br	—	—	—	[31, 30, 86]
BNL E-0734	C	—	1 GeV to 3 GeV	$d\sigma/dQ^2$ with ν_μ, $\bar{\nu}_\mu$, $\nu_\mu/\bar{\nu}_\mu$	[87, 88]
MiniBooNE	CH_2	$p + n$ scintillation light	0.8 GeV	$d\sigma/dQ^2$ with ν_μ, $\bar{\nu}_\mu$, $\nu_\mu/\bar{\nu}_\mu$	[89, 90]
T2K	H_2O	γ de-excitation	0.6 GeV	$\langle\sigma\rangle_\Phi$ with ν	[91, 92]

Table 3.1 Summary of previous neutrino–nucleon neutral current elastic cross sections measurements.

Resonant pion production

Neutrino–nucleon scatterings may also proceed inelastically and induce the target nucleon into an excited state, its subsequent decay most often leads to a π and nucleon as the final-state particles. As seen in Fig. 3.2, this channel dominates at $\sim$2 GeV.

Deep inelastic scattering

Deep inelastic scattering (DIS) is the dominant neutrino–nucleon interaction mode for neutrinos with energies around and above 10 GeV. In this channel, an incident neutrino interacts with one of the constituent quarks of a target hadron to produce a lepton and a hadronic system, which ultimately showers, in the final state.

Coherent pion scattering

Coherent scattering of neutrinos from nuclei is a π producing interaction mode facilitated at low Q^2, in which a single forward-going π is created. This can proceed either in CC mode ($\nu_\ell + N \to \ell^- + N + \pi^+$) or in NC mode ($\nu_\ell + N \to \nu_\ell + N + \pi^0$).

3.3 Flavour Mixing

It is because neutrinos are massive particles that we may consider distinct their mass eigenstates, that have definite mass, and flavour eigenstates, that have definite flavour. The flavour eigenstates shall be denoted by $|\nu_\alpha\rangle$, where $\alpha \in \{e, \mu, \tau\}$, and the mass eigenstates shall be denoted by $|\nu_i\rangle$, where $i \in \{1,2,3\}$. Given that these are eigenbases of the same system, it follows that any particular flavour state can be expressed as a superposition of the mass eigenstates, and vice versa. These two eigenbases may be transformed between using the relations

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle\,, \qquad\qquad |\nu_i\rangle = \sum_\alpha U^*_{\alpha i} |\nu_\alpha\rangle \tag{3.29}$$

where U is a 3×3 complex unitary matrix called the *Pontecorvo–Maki–Nakagawa–Sakata* (PMNS) matrix which may be represented as

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}. \tag{3.30}$$

The magnitude of a particular element of the PMNS matrix may be physically interpreted as the extent of the overlap between a particular flavour eigenstate and a particular mass eigenstate; for example, $|\langle \nu_1 | \nu_e \rangle|^2 = |U_{e1}|^2$. The condition of unitarity imposed on the PMNS matrix, that ensures particle number is conserved, allows its nine elements to be parametrised by just four real and independent parameters[2]. The standard parametrisation of neutrino mixing is an angular parametrisation that consists of the three Euler angles, θ_{12}, θ_{23} and θ_{13}, and a phase δ, given by

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \tag{3.31}$$

where $s_{ij} \equiv \sin\theta_{ij}$ and $c_{ij} \equiv \cos\theta_{ij}$.

The rotation angles can be said to transform the orthonormal $\mathbb{C}^3$ basis consisting of the mass eigenstates to a basis expressed in the flavour eigenstates, where θ_{ij} is the angle of rotation to be enacted in the plane of $|\nu_j\rangle$–$|\nu_i\rangle$. The rotation angle θ_{13} is simple to interpret physically, it parametrises the extent of the overlap between flavour state $|\nu_e\rangle$ and mass state $|\nu_3\rangle$, as $|\langle \nu_3 | \nu_e \rangle|^2 = \sin^2\theta_{13}$. The other two angles can be said to express the following proportions between mass–flavour overlaps, $|\langle \nu_2 | \nu_e \rangle|^2 = \tan^2\theta_{12} |\langle \nu_1 | \nu_e \rangle|^2$ and $\langle \nu_3 | \nu_\mu \rangle|^2 = \tan^2\theta_{23} |\langle \nu_3 | \nu_\tau \rangle|^2$. The phase δ introduces an imaginary component to some of the PMNS elements when non-zero—its role as the generator of neutrino sector *CP* symmetry, or asymmetry, becomes evident when a conjugation of charge is applied to Eq. (3.29) to transform the neutrino to its antiparticle counterpart.

Combining the two equations in Eq. (3.29), it follows that a neutrino of definite flavour can be expressed as a superposition of other flavours, according to

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} \sum_\beta U^*_{\beta i} |\nu_\beta\rangle . \tag{3.32}$$

It is exercising to consider the case where the mixing angles are vanishing. Eq. (3.32) may be expressed in terms of the diagonal and off-diagonal elements of the PMNS matrix according to

[2]The two Majorana phases, α_1 and α_2, have been neglected for simplicity.

$$|\nu_\alpha\rangle = \sum_{i\neq\beta} U_{\alpha i}U^*_{\beta i}|\nu_\beta\rangle + \sum_{i=\beta} U_{\alpha i}U^*_{\beta i}|\nu_\beta\rangle\,. \tag{3.33}$$

If there is no mixing, the PMNS matrix is necessarily diagonal, therefore

$$|\nu_\alpha\rangle = (0) + |U_{\alpha\alpha}|^2\,|\nu_\alpha\rangle \tag{3.34}$$

$$= |\nu_\alpha\rangle \tag{3.35}$$

and the flavour and mass eigenbases exactly coincide.

3.3.1 Oscillations in vacuum

The phenomena of flavour transitions as a natural consequence of flavour mixing will be demonstrated. Consider first the time evolution t of a neutrino with well-defined mass j propagating in vacuum with momentum p, the time dependence is revealed by the Schrödinger equation

$$i\frac{\partial\,|\nu_j(t)\rangle}{\partial t} = H\,|\nu_j(t)\rangle \tag{3.36}$$

where H is the free Hamiltonian in the case of vacuum. The mass eigenstates are stationary states, so the solution is the standing wave

$$|\nu_j(t)\rangle = e^{-iEt}\,|\nu_j(t=0)\rangle \tag{3.37}$$

where E is the neutrino energy. Consider now the time-dependence of a neutrino of definite flavour ν_α emitted by a source at $t=0$ in terms of the propagation of its constituent mass eigenstates, given by

$$|\nu_\alpha(t)\rangle = \sum_j U_{\alpha j}|\nu_j(t)\rangle = \sum_j e^{iE_jt}U_{\alpha j}\,|\nu_j(t=0)\rangle\,. \tag{3.38}$$

At this stage, it is important to note that if the energies of the mass eigenstates are non-degenerate, then these states necessarily fall into decoherence with time. Consider now the propagation of ν_α in terms of the stationary flavour states, given by

$$|\nu_\alpha(t)\rangle = \sum_j e^{iE_jt}U_{\alpha j}\sum_\beta U^*_{\beta j}\,|\nu_\beta(t=0)\rangle\,. \tag{3.39}$$

With knowledge that the flavour eigenstates form an orthonormal basis according to

$$\langle \nu_\alpha(t=0) | \nu_\beta(t=0) \rangle = \delta_{\alpha\beta} = \begin{cases} 1, & \text{if } \alpha = \beta, \\ 0, & \text{if } \alpha \neq \beta \end{cases}, \tag{3.40}$$

the amplitude that ν_α may be found later at time t as ν_β, or the transition amplitude, is given by

$$\langle \nu_\beta(t) | \nu_\alpha \rangle = \sum_j e^{iE_j t} U_{\alpha j} U^*_{\beta j}. \tag{3.41}$$

The probability that ν_α may be found later at time t as ν_β, or the transition probability, is given by

$$P_{\nu_\alpha \to \nu_\beta}(t) = \left| \langle \nu_\beta(t) | \nu_\alpha \rangle \right|^2 = \left| \sum_j e^{iE_j t} U_{\alpha j} U^*_{\beta j} \right|^2 \tag{3.42}$$

$$= \sum_j \sum_k U^*_{\beta k} U_{\beta k} U_{\alpha j} U^*_{\alpha j} e^{-i(E_k - E_j)t} \tag{3.43}$$

which assumes the form of a harmonic oscillator. At this stage, it is illustrative to consider the following two scenarios: *i)* the case of vanishing mixing angles, or a diagonal PMNS matrix, once again, and *ii)* the case of degenerate neutrino masses. In the first case, a transition of flavour is impossible as the transition probability evaluates to

$$P_{\nu_\alpha \to \nu_\beta}(t) = \begin{cases} 1, & \text{if } \alpha = \beta, \\ 0, & \text{if } \alpha \neq \beta \end{cases}. \tag{3.44}$$

The outcome of the second case is demonstrated by considering that neutrinos are always found to be ultra-relativistic in nature ($p \gg m_i$). Therefore, the following approximation is reasonable:

$$E_i = \sqrt{m_i^2 + p_i^2} \simeq p_i + \frac{m_i^2}{2p_i} \simeq E + \frac{m_i^2}{2E} \tag{3.45}$$

where m_i and p_i are the mass and momentum of the mass eigenstate i. Hence, Eq. (3.43) becomes

$$P_{\nu_\alpha \to \nu_\beta}(t) = \sum_j \sum_k U^*_{\beta k} U_{\beta k} U_{\alpha j} U^*_{\alpha j} e^{\frac{-i\Delta m^2_{ji} t}{2E}} \tag{3.46}$$

where $\Delta m^2_{ji} \equiv m^2_j - m^2_i$. The propagation length L, i.e. the distance between source and detector, is more practical to measure than propagation time; therefore, given that for $L \approx t$ for ultra-relativistic particles, the following substitution is convenient:

$$P_{\nu_\alpha \to \nu_\beta}(t) = \sum_j \sum_k U^*_{\beta k} U_{\beta k} U_{\alpha j} U^*_{\alpha j} e^{\frac{-i\Delta m^2_{ji} L}{2E}}. \tag{3.47}$$

If neutrinos are degenerate, i.e. $\Delta m^2_{ji} = 0$ for all i, j, the neutrino mass eigenstates propagate coherently in Eq. (3.38), and there are no interference terms in the transition probability function as a consequence; the outcome is as if there were no mixing at all. Thus both non-degenerate neutrino masses and non-zero mixing angles form necessary conditions for flavour transitions of neutrinos in flight. And from studying the form of Eq. (3.47), it is further uncovered that the differences between the neutrino masses and the mixing angles control the rate and amplitude of the oscillation pattern.

Eq. (3.47) may be further manipulated and generalised for neutrinos and antineutrinos to take the following form:

$$P_{\nu_\alpha \to \nu_\beta} = \delta_{\alpha\beta} - 4 \sum_{j>i} \Re \left(U^*_{\alpha j} U_{\beta j} U_{\alpha i} U^*_{\beta i} \right) \sin^2 \frac{\Delta m^2_{ji} L}{4E} \tag{3.48}$$

$$\pm 2 \sum_{j>i} \Im \left(U^*_{\alpha j} U_{\beta j} U_{\alpha i} U^*_{\beta i} \right) \sin \frac{\Delta m^2_{ji} L}{2E} \tag{3.49}$$

where $(\pm)$ evaluates to $(+)$ for neutrinos and $(-)$ for antineutrinos. Remembering that the phase δ introduces an imaginary component to the PMNS matrix when δ is not a factor of π, it is clear that the oscillation patterns of neutrinos and antineutrinos are identical if δ is *CP*-conserving. Hence, the difference in oscillation pattern observed of neutrinos and antineutrinos is a residue of the extent of leptonic *CP* symmetry.

3.3.2 *CP* violation

The difference in oscillation pattern between neutrino and antineutrinos is given by

$$\Delta P^{CP}_{\nu_\alpha \to \nu_\beta} = -16 J_{\alpha\beta} \sin \left(\frac{\Delta m^2_{12} L}{4E} \right) \sin \left(\frac{\Delta m^2_{23} L}{4E} \right) \sin \left(\frac{\Delta m^2_{13} L}{4E} \right) \tag{3.50}$$

where $J_{\alpha\beta}$ is the Jarlskog invariant, which provides a definition of the magnitude of CP-asymmetry that is invariant of the chosen parametrisation of the neutrino mixing matrix. The Jarlskog invariant [74] is defined as

$$J_{\alpha\beta} \equiv \Im(U_{\alpha 1}U^*_{\alpha 2}U^*U_{\beta 1}U_{\beta 2}) = \pm c_{12}s_{12}c_{23}s_{23}c^2_{13}s_{13}\sin\delta \tag{3.51}$$

so, all mixing angles first must be non-zero for *CP* violating effects to take hold.

3.3.3 Mass hierarchy

Given three non-degenerate neutrino masses, there are two non-negative independent mass squared differences. These will be denoted by Δm^2_{big} and Δm^2_{small} for now, where $\Delta m^2_{big} > \Delta m^2_{small}$. The smaller mass splitting Δm^2_{small} is always labelled Δm^2_{21} by convention. There are two ways of arranging the neutrino masses in order of their squared masses: either Δm^2_{small} is the mass squared difference between the eigenstates lowest in mass, such that $m_1 < m_2 < m_3$; or Δm^2_{small} is the mass squared difference between the eigenstates highest in mass, such that $m_3 < m_1 < m_2$. The former hierarchy is identified as normal ordering, where $\Delta m^2_{big} \equiv \Delta m^2_{32}$, whereas the latter is identified as inverted ordering, where $\Delta m^2_{big} \equiv \Delta m^2_{31}$.

From experiment, it has been found that $\Delta m^2_{sol} \ll \Delta m^2_{atm}$. Hence $\Delta m^2_{sol} \equiv |\Delta m^2_{21}|$, but $\Delta m^2_{atm} \equiv \Delta m^2_{32}$ under normal ordering and $\Delta m^2_{atm} \equiv -\Delta m^2_{31}$ under inverted ordering. However since $|\Delta m^2_{31}| \cong |\Delta m^2_{32}|$, Δm^2_{atm} can safely be identified as either Δm^2_{31} or Δm^2_{32}.

3.3.4 Oscillations in matter

The interaction modes that a flavour eigenstate can undergo in a particular medium each introduce a corresponding potential to the flavour Hamiltonian. In the vacuum case, as already demonstrated, the energy of the system takes the form of the free Hamiltonian; in matter, however, neutrinos of all flavours are free to undergo CC interactions with nuclei and NC interactions with both nuclei and electrons, and also ν_e may undergo CC interactions on electrons. These introduce the following potentials to the Hamiltonian:

$$V_{CC} = \sqrt{2}G_F N_e, \qquad V_{NC} = -\frac{\sqrt{2}G_F N_n}{2} \tag{3.52}$$

where N_e and N_n are the electron and neutron densities of the medium. The NC potential V_{NC} is felt by all neutrinos equally and hence ultimately cancel; the CC potential, on the other hand, alters strictly the ν_e mass in matter such that

$$m_e^2 \to m_e^2 + 2EV_{CC} \tag{3.53}$$

where E is the neutrino energy. This perturbation is known interchangeably as the *Mikheyev–Smirnov–Wolfenstein* (MSW) effect or *matter* effect, which subsequently modifies the vacuum oscillation probability pattern for neutrinos that propagate through matter.

PART II

MEASUREMENT OF THE NEUTRINO–PROTON NEUTRAL CURRENT ELASTIC SCATTERING CROSS SECTION ON CARBON WITH THE T2K OFF-AXIS NEAR DETECTOR

CHAPTER 4

INTRODUCTION TO PART II

This part of the thesis details a measurement of the neutrino–proton NC0π scattering cross section on carbon at the near detector of the T2K experiment. First of all, the experimental configuration and hardware of T2K will be summarised in Chapter 5. Then, the analysis strategy, physics dataset and analysis samples used in the cross section extraction will be introduced in Chapter 6. The development of a multivariate signal selection will be reported in Chapter 7 and its subsequent deployment in the cross section measurement will be given in Chapter 8. Finally, in Chapter 9, conclusions will be drawn and suggestions for further work will be offered.

INTRODUCTION TO PART II

CHAPTER 5

THE T2K EXPERIMENT

In 2001, a successor experiment to K2K was proposed by KEK to look for ν_μ changing to ν_e [96]. At that time, and at a limit of $\sin^2 2\theta_{13} < 0.12$ (at 90% CL), CHOOZ had found no evidence of $\bar{\nu}_e$ appearance generated by Δm^2_{23} [51]. Given θ_{13} was small (or perhaps even zero), this result proved ν_μ–ν_e transitions are sub-leading and ν_μ–ν_τ transitions dominate in flavour oscillations driven by Δm^2_{32}. The search for ν_e appearance thus became of great importance as the success of future searches for CP-violation in the neutrino sector hinged on a non-zero θ_{13}.

Given the limit set by CHOOZ, measurement of ν_e appearance was beyond the K2K's sensitivity owing to insufficient power of the KEK synchrotron; a high-intensity beam facility and new near detector were needed. The upstream site of K2K's successor became one of a few physics programs at Japan's Proton Accelerator Research Complex (J-PARC), which was jointly constructed by KEK and the Japan Atomic Energy Agency (JAEA) in the coastal village of Tōkai, Ibaraki prefecture between 2001 and 2008. With Super-Kamiokande kept as the far detector, this new long-baseline configuration was aptly named T2K (Tōkai-to-Kamioka) and began operation in 2010.

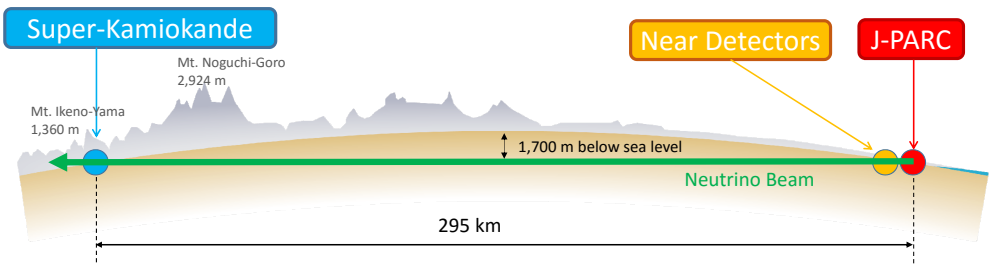

Fig. 5.1 A diagram of the configuration of the T2K experiment (reproduced from Ref. [97]).

T2K [98] measures neutrino oscillations over a baseline stretching 295 km from coast-to-coast across Honshū, the largest island of Japan. As seen in Fig. 5.2, in the east, an intense narrow-band beam of mostly ν_μ is sent from J-PARC through the earth towards the near detector facilities just 280 m downstream and ultimately Super-Kamiokande in the west. Using its world-class neutrino source, not only is T2K intended to improve the limit set by CHOOZ on θ_{13} by a factor of 20 but also to improve the precision of the ν_μ disappearance parameters: namely, θ_{23} to within 1% precision and Δm^2_{23} to within 3% precision. T2K's original physics goals were not limited to the topic of neutrino mixing, however, and also included neutrino interaction studies, sterile neutrino searches, among other exotic topics.

5.1 Beam Facility

5.1.1 J-PARC

J-PARC hosts a suite of proton accelerators for the activities of the three experimental groups on-site: the Materials and Life Sciences facility, the hadron facility and the neutrino facility. As seen in Fig. 5.2, the proton beamline consists of a linear accelerator (LINAC), a 3 GeV rapid cycle synchrotron (RCS) and a 50 GeV main ring (MR) synchrotron [99].

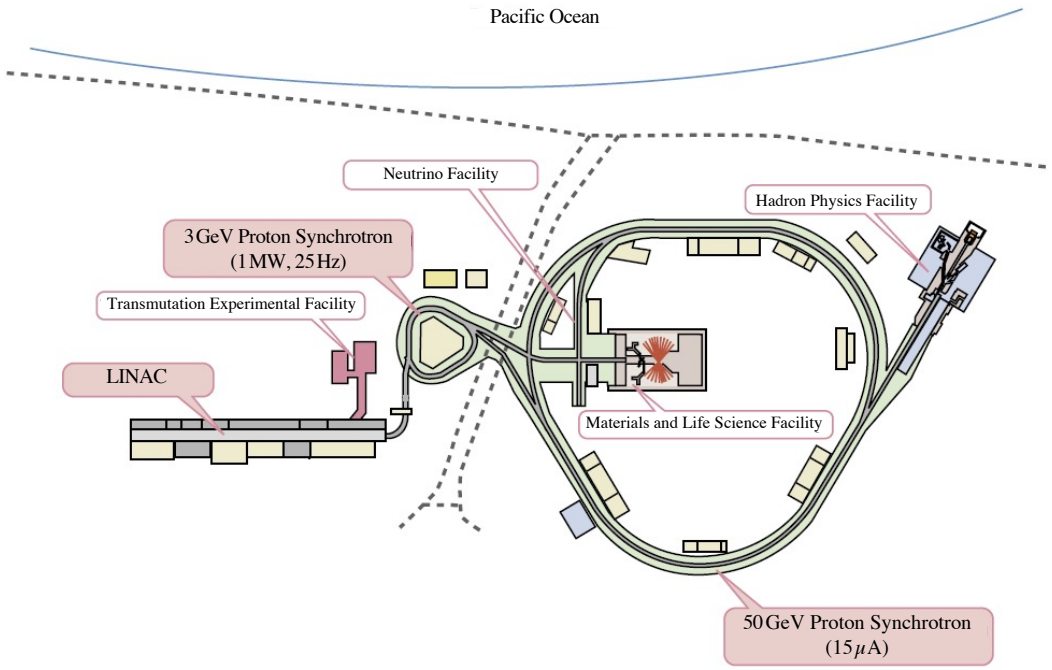

Fig. 5.2 A view of J-PARC's accelerator facility (adapted from Ref. [100]).

At the first stage of the beamline, H^- are generated and accelerated to 400 MeV by the LINAC. Before entering the RCS, the H^- are rid of their electrons by charge stripping foils to become protons. The RCS—348 m in circumference—accelerates the protons in bunches of two to 3 GeV on a 25 Hz cycle. From the RCS, ∼95% of the protons are sent to the Materials and Life Sciences Facility and the remainder are injected into the MR. The MR accelerates the protons to 30 GeV in bunches of eight over its 1568 m circumference. The protons may be extracted from the MR either to the Hadron facility, when in slow extraction mode, or the neutrino facility, when in fast extraction mode. In the latter case, the eight bunches are extracted to the neutrino beamline in a single turn by five kicker magnets. The timing structure of the proton delivery from J-PARC is well-understood from the perspective of T2K, so that it may be used to exclude various backgrounds from the detectors.

5.1.2 Neutrino beamline

The production of the neutrino beam for T2K follows the conventional fixed-target strategy: *i)* generate and extract a stream of high-energy protons, *ii)* focus the protons onto a fixed target, *iii)* let the primary particles (mostly π, K) produced from proton-target interactions decay into (mostly) μ and ν_μ (via $M^+ \to \mu^+ \nu_\mu$, M = π, K). The first step is handled by the J-PARC accelerator, the latter steps by the neutrino beamline.

The neutrino beamline is divided into two stages: a primary beamline, which focuses and steers the protons onto the target; and a secondary beamline, where the neutrinos are produced. As seen in Fig. 5.3, a 54 m straight section, a 147 m arced section and a further 37 m straight section constitute the primary beamline. Along the first straight section, a series of 11 normal conducting magnets prepare the protons extracted from the MR so that they may be accepted by the arced section. Then, a series of 14 pairs of super-conducting magnets along the arced section steer the protons 80.7° around to face Kamioka, and a further 10 normal conducting magnets along the final straight section focuses and tilts the beam 3.6° downwards onto the target and towards the detectors.

The secondary beamline is home to the target station, the focusing horns, the decay pipe and, ultimately, the beam dump. The target for the delivered protons is a graphite rod, 91.4 cm in length (1.9 interaction length) and 2.6 cm in diameter, which is encased within a 0.3 mm thick titanium shell. If a material much denser than graphite (1.8 g cm^{-3}) were to be used for the target, it would be melted by the heat generated by the pulsed beam. Protons interacting with the target nuclei cause a large number of mesons to enter the decay pipe downstream, where they decay primarily to μ and ν_μ pairs along its 96 m length. The toroidal magnetic horns, of which there are three along the secondary beamline, focus charged mesons of a preferred sign in the forward direction to increase the neutrino intensity. When each

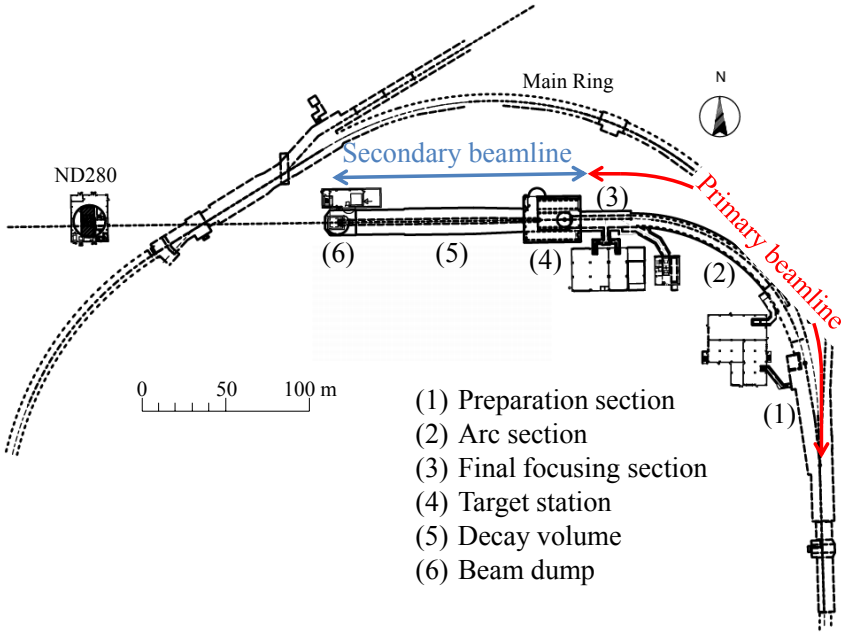

Fig. 5.3 A schematic diagram of the neutrino beamline at J-PARC (reproduced from Ref. [101]).

horn operates at a current of 320 kA, the flux at the peak of Super-Kamiokande's neutrino energy spectrum is increased by a factor of $\sim$16 compared to when the horns are switched off. Whether T2K operates on neutrinos or anti-neutrinos is decided by the polarity of the horns: a forward horn current will accept π^+, K^+ (thus ν_μ, ν_e) and reject π^-, K^- (thus $\bar{\nu}_\mu$, $\bar{\nu}_e$), a reverse horn current will perform the inverse. At the end of the decay pipe is the beam dump—a large mass of graphite and iron where muons with energies below $\sim$5 GeV and hadrons are absorbed. From hereon, the neutrinos pass through the earth towards the detectors. The high-energy muons that pass the beam dump are sampled by a monitoring system known as MUMON—from a measurement of their spatial distributions, MUMON can infer the beam intensity to better than 3% and its direction to within 0.25 mrad.

In actuality, the neutrino beamline was deliberately made to point at a small angle away from the far detector. By doing so, the peak of the neutrino energy may be tuned to a desired value by exploiting the kinematics of pion decay—the dominant mechanism by which beam neutrinos are made. The energy of the daughter neutrino E_ν produced from $\pi^\pm$ decay is given by

$$E_\nu = \frac{m_\pi^2 - m_\mu^2}{2\left(E_\pi - p_\pi \cos\theta\right)} \tag{5.1}$$

where m_π is the mass of the $\pi^\pm$, m_μ is the mass of the $\mu^\pm$, E_π is the energy of the parent $\pi^\pm$, p_π is the momentum of the $\pi^\pm$ and θ is the angle subtending the directions of the $\pi^\pm$ and daughter ν.

At a baseline of 295 km, the first maximum of the ν_e appearance probability function appears at around 600 MeV. In Fig. 5.4 it is seen that by introducing a 2.5° offset in angle between the beam direction and the position of the detector, the peak neutrino energy at the far detector is tuned such that an increased flux appears at the oscillation maximum—thus enhancing sensitivity to the ν_e appearance channel. The expected fluxes at 295 km by neutrino type are shown in Fig. 5.5. An intrinsic ν_e ($\bar{\nu}_e$) component to the beam is caused by the decay of secondary K^+ (K^-), constituting 0.4% of the total flux in neutrino (anti-neutrino) mode at 600 MeV.

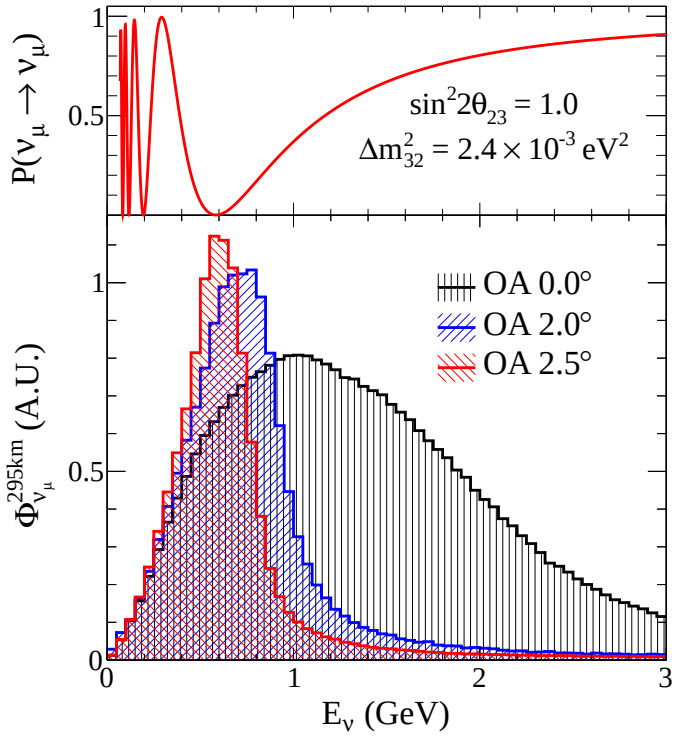

Fig. 5.4 Muon neutrino fluxes for different off-axis angles and survival probability function at the T2K far detector sent from J-PARC (reproduced from Ref. [101]).

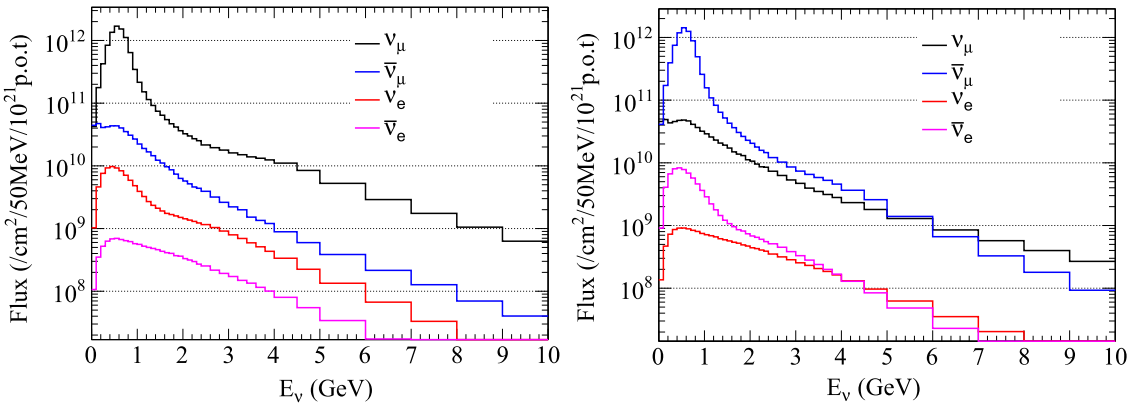

(a) Neutrino mode (forward horn current). (b) Anti-neutrino mode (reverse horn current).

Fig. 5.5 The predicted neutrino fluxes at the ND280 detector when running in neutrino mode (left) and anti-neutrino mode (right) (both reproduced from Ref. [102]).

5.2 Near Detector Facility

In the near detector hall lie three detectors: ND280, INGRID and WAGASCI-BabyMIND—all positioned ∼280 m from the target station to sample the un-oscillated neutrinos. ND280 is the primary near detector of T2K and sits at the same off-axis angle to the beam as the far detector. INGRID, however, is stationed directly on-axis to monitor the properties of the delivered beam. WAGASCI-BabyMIND—a recent addition to the near detector hall—is an auxiliary detector at an off-axis angle of ∼1.5° that exists to reduce systematic errors at the far detector.

5.2.1 ND280

ND280 (Near-Detector-At-280 m), as seen in Fig. 5.6, is a tracking detector formed of several subdetectors arranged within a 0.2 T dipole magnet[1]. At the core of ND280 is the π^0 detector (P0D), which lies most upstream, and a further adjoining system known as the tracker, formed of two fine-grained detectors (FGDs) interleaved with three time-projection chambers (TPCs). Both the P0D and tracker are surrounded by electromagnetic calorimeters (ECals), which are enclosed by the yokes of the magnet in turn. The magnet's return yokes are instrumented with layers of scintillator which constitute the Side Muon Range Detector (SMRD).

[1]Having served in CERN's UA1 experiment from 1981 to 1990 and NOMAD from 1995 to 2000, the magnet was granted a third lease of life with ND280.

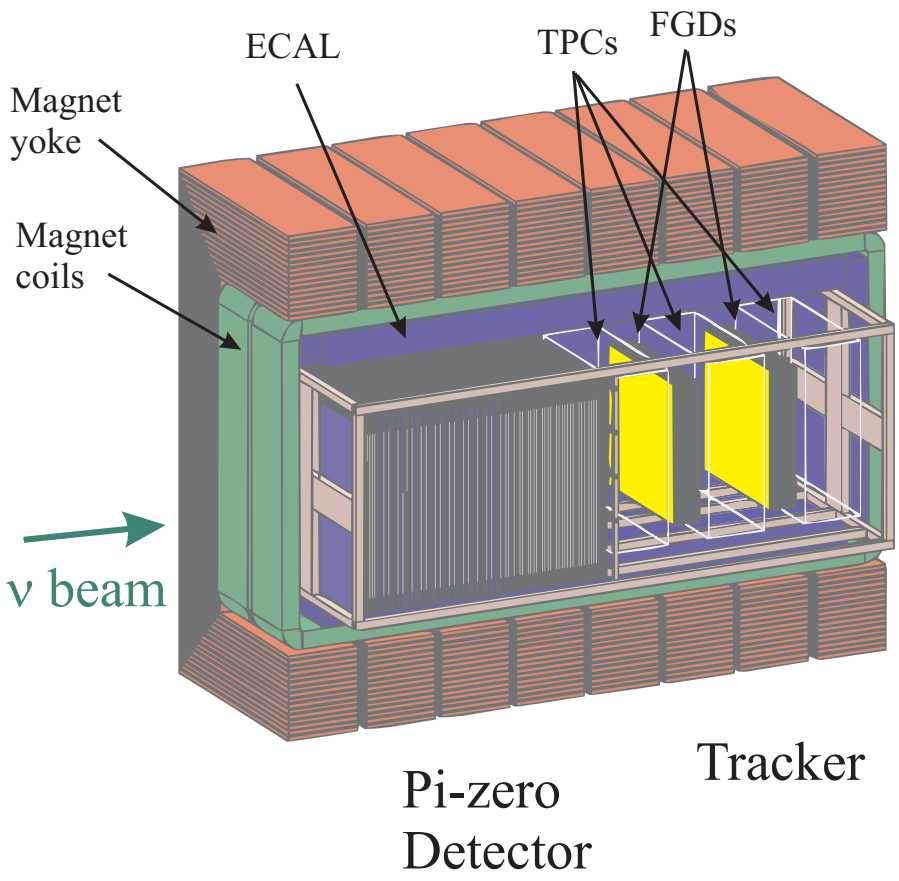

Fig. 5.6 A cross-sectional schematic diagram of T2K's ND280 detector (reproduced from Ref. [103]).

Fine-grained detectors (FGDs)

The highest purpose of the tracker is to measure the Charged Current Quasi-Elastic (CCQE) interaction—the most common interaction at T2K's beam energy. The CCQE interaction ($\nu_\ell + n \to \ell^- + p$) leaves a proton and lepton in the final state. With this in mind, the two FGDs—FGD1 upstream and FGD2 downstream—are thin enough (both are 2300 mm $\times$ 2400 mm $\times$ 365 mm) that most interaction products, like leptons, will cross to the next TPC module, but short-range particles, like protons, may stop in the FGD they were produced. CCQE events can be identified well from the kinematics of the recoil proton. It is therefore paramount that the FGDs were given high granularity: that the direction of protons produced via CCQE, given their short ranges travelled within the FGDs, may be measured with good precision.

A reoccurring unit of active material within the near detectors is the XY module: a panel of plastic scintillator bars layered with another in an orthogonal rotation, with each bar instrumented with a Wavelength-Shifting (WLS) fibre to transport light to a Multi-Pixel Photon Counter (MPPC) for read-out. A construction as such provides target material with three-dimensional tracking in-situ. The XY modules of the FGDs consist of 192 bars in each panel, and each bar is a 9.61 mm × 9.61 mm × 1864.3 mm piece of extruded polystyrene. FGD1 consists of 15 of such modules; FGD2, however, alternates 7 XY modules with 6 water modules—the reasoning follows. An oscillation measurement demands a prediction of the neutrino interaction rates in the absence of oscillations at the far detector. The cross-sections of all processes relevant to T2K are subject to various nuclear effects (such as Pauli blocking, pion re-scattering, nuclear absorption) which cannot be reliably be accounted for from a purely theoretical approach [103]. Therefore the nuclear cross-sections on water, the target material of the far detector, are directly measured by interactions on FGD2.

Time projection chambers (TPCs)

Any charged particle travelling in any of the three TPCs in the tracker will have its momentum measured and flavour content identified. In general terms, a TPC is a detector filled with a sensitive liquid or gas which traces the paths of incident charged particles in three dimensions when ionisation electrons, induced along the path the incident charged particle takes through the medium, are carried to a position-sensitive charge collector (most commonly wire-chambers) by a uniform electric field. A magnetic field acting parallel to the electric field is usually also maintained throughout the chamber to mitigate the diffusion of the ionisation electrons.

In the rectangular TPCs of ND280, the electric field is established in the chamber by a cathode plate at −25 kV, which splits the chamber in the middle, and arrays of 12 Micromegas (Micro-MEsh Gaseous Structure) detectors at −350 V, which face the cathode plate on both sides of the chamber [104]. A maximum drift distance of 897 mm within each chamber sets a field strength of 275 V cm^{-1}, which in turn sets a maximum drift velocity approaching the saturation point of the chamber medium. The chamber is filled with a gas mixture of Ar:CF_4:iC_4H_{10} (95%:3%:2%) which gives a low charge diffusion and is known to work well with Micromegas detectors. A diagram of the TPC design is shown in Fig. 5.7.

Spatial information read from the Micromegas pads necessarily provides a projection of a given particle track in the plane transverse to the field direction. The track's longitudinal coordinate within the chamber is inferred from the arrival time of the signal. Thus, when spatial and timing information are combined, tracking in three-dimensions is achieved. An important function for the TPCs is to tell between e and μ tracks—as this also means to

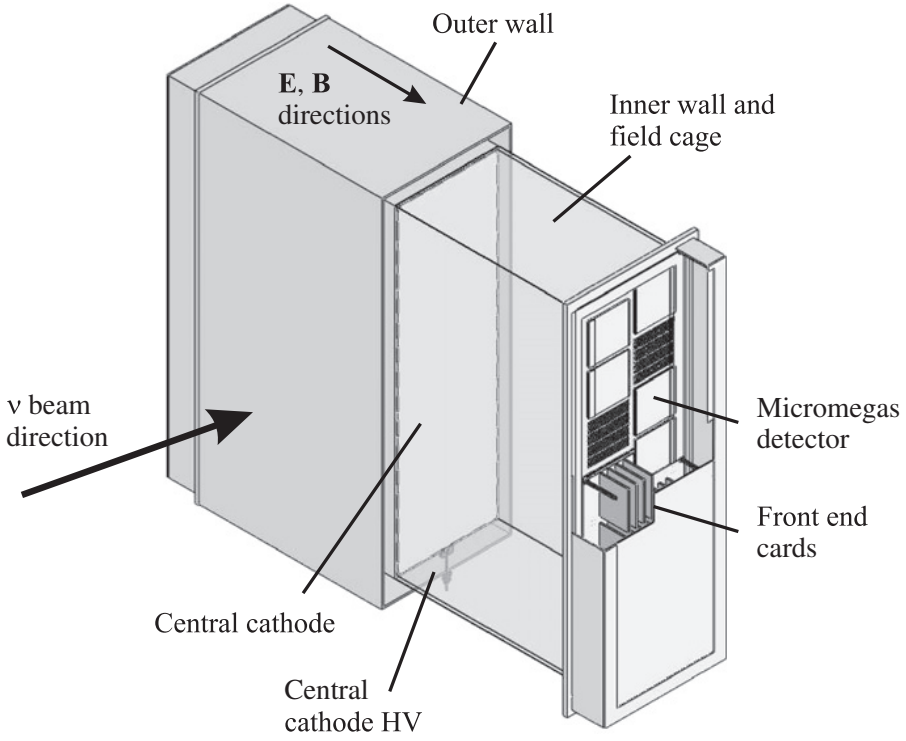

Fig. 5.7 A schematic diagram of the TPC subdetector belonging to T2K's ND280 detector (reproduced from Ref. [103]).

distinguish ν_e and ν_μ in the case of the CCQE interaction. Particles may be identified by their characteristic energy loss as they permeate the TPC. In the relevant momentum range, the energy loss of electrons is almost twice that of muons, therefore the resolution of energy loss for the TPCs must reside around 10% to achieve adequate separation. An energy resolution of $(7.8 \pm 0.2)\%$ was met for minimum ionising particles and a misclassification rate of 0.2% for muons with momenta under 1 GeV.

Pi-zero detector (P0D)

Interactions leaving a π^0 in the final state are the largest background associated with the ν_e appearance signal at the far detector; another significant contribution is given by ν_e intrinsic to the beam, or, in other words, ν_e which did not change their flavour. The P0D combats both of these backgrounds by measuring them at the near detector station—the rate of π^0 undergoing neutral current interactions ($\nu_\mu + N \to \nu_\mu + N + \pi^0 + X$) and the ν_e rates on water.

The P0D, throughout, has a sandwich structure of XY modules in alternation with target layers, though the materials of the target layer differ by region. As seen in Fig. 5.8, there

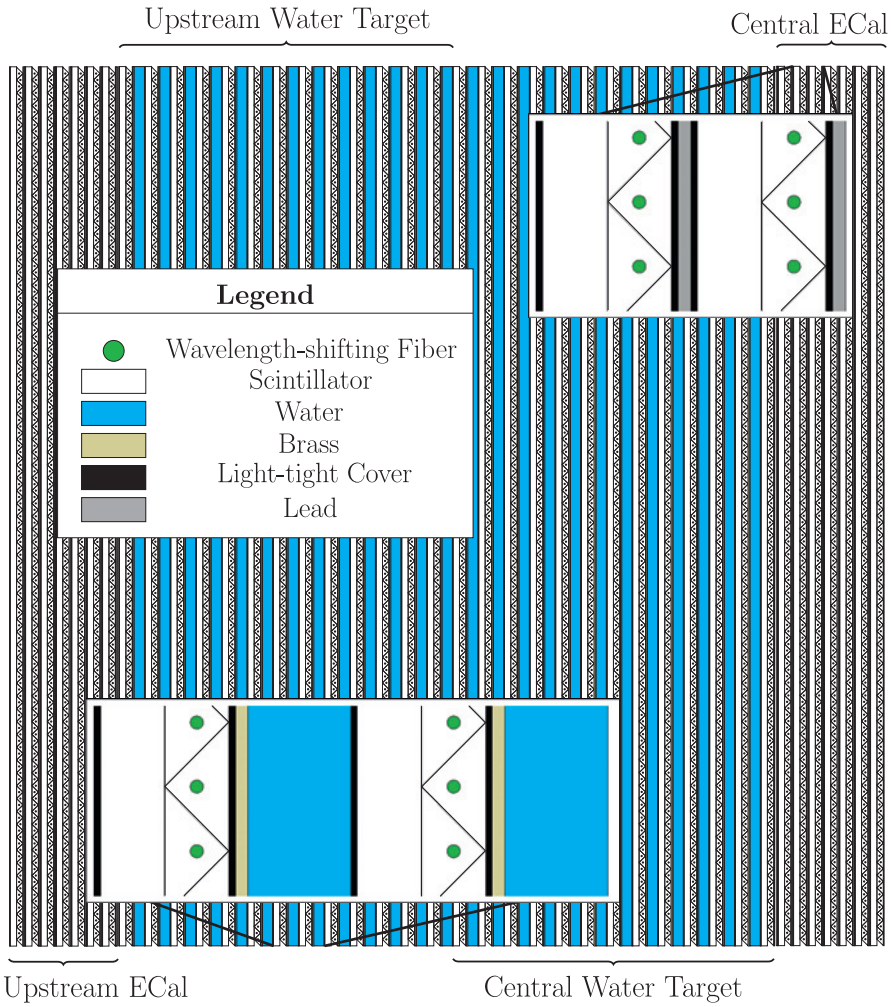

Fig. 5.8 A cross-sectional diagram of the π^0 subdetector (P0D) belonging to T2K's ND280 detector (reproduced from Ref. [103]).

are two such kinds of regions in the P0D—a central region, with target layers made of brass and either water or air, sits between two ECals[2], which have lead targets. With this layout, the ECals establish a veto region around the central region to reject particles coming in from outside the P0D and also measure electromagnetic showers. Seven XY modules instrument each ECal and 26 instrument the central region. The P0D's XY modules are relatively coarser than that of the FGDs, each consisting of 134 horizontal bars (each 2133 mm in length) and 126 vertical bars (each 2272 mm in length) triangular in profile rather than square. The water modules within the central region are sometimes filled and sometimes left empty: as the P0D,

[2]The P0D's ECals are not to be confused with the P0D-ECals, which are considered part of ND280's ECal system rather than ND280's P0D system.

by mass, is not mostly made of water, the interaction rates on water must be derived through subtraction of the rates found with and without the modules filled.

Electromagnetic calorimeter (ECal)

The ECal encases the subdetectors hitherto described (bar the front face of the P0D) to catch the energies and directions of escaping particles [105]. Thirteen modules form the ECal altogether: the six P0D-ECals, which surround the P0D (two top, two bottom, two side[3]); the six barrel ECals (two top, two bottom, two side), which surround the tracker; and the downstream ECal, which lies beyond the last TPC. The barrel and downstream ECals are often collectively referred to as the tracker-ECal for convenience.

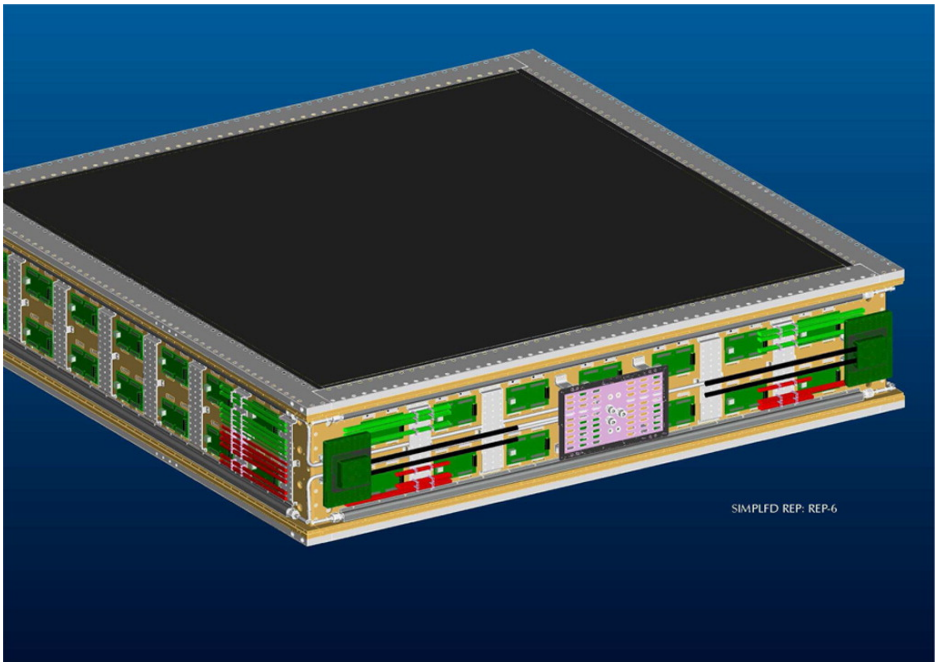

Fig. 5.9 An engineers drawing of an ECal module belonging to T2K's ND280 detector (reproduced from Ref. [103]).

As their physics goals differ, so do the designs of the tracker-ECals and P0D-ECals. Both, however, consist of alternating layers of XY modules and lead panels. To complement the tracking capabilities of the TPCs, the tracker-ECals are specially built to reconstruct electromagnetic showers in fine detail. With this in mind, there are 31 (10 X_0) scintillator-lead layers in the barrel-ECal and 34 (11 X_0) in the downstream ECal. The P0D, on the other hand, is perfectly capable of reconstructing electromagnetic showers with its own ECals. The

[3]For access to the subdetectors, the UA1 magnet is split along the beam direction; the top and bottom ECal modules are also split along the beam direction for access to the P0D and tracker.

composition of the P0D-ECal is therefore relatively coarse in comparison, with six thicker (4.3 X_0) lead layers and arranged such that the scintillator bars within its XY modules all run parallel to the beam. The particle energies sampled by the P0D-ECal assists the P0D in distinguishing μ from γ to improve identification of π^0. The hit efficiency of the downstream ECal is 99% and 95% for the barrel ECal [105].

Side muon range detector (SMRD)

The outermost subdetector of ND280 is the SMRD; by and large, it is used to measure the momenta of muons that escape the inner sub-detectors at high-angles (relative to the beam axis), and to identify cosmic rays and interactions on the magnet [106].

The SMRD comprises 440 scintillator panels that instrument the magnet. The magnet is made of two separate yokes, and each yoke is made of 8 adjacent C-shaped return yokes. Within each return yoke are 16 layers of 4.8 cm steel plates at 1.7 cm spacings; the scintillator panels which constitute the SMRD are inserted between these air-gaps.

5.2.2 INGRID

As the energy of a neutrino depends on the decay angle of its parent meson, monitoring and maintaining the direction of the neutrino beam with precision, as well as its intensity, is of great importance to T2K as an off-axis experiment. While many systems exist along the neutrino beamline to monitor the intensity, position and timing of the protons delivered onto the target, they can only give indirect information about the neutrinos delivered to T2K. Even the information offered by MUMON, as mentioned in Section 5.1.2, is limited, as only muons with high momenta are sampled. Below ND280 in the near detector hall, and centred on the beam-axis, the INGRID (Interactive Neutrino GRID) detector directly samples the neutrino beam and gives daily monitoring of its intensity, shape and direction.

As seen in Fig. 5.10a, INGRID is made of 16 identical iron scintillator modules: 14 of which are arranged in a cross formation (with seven vertical modules in front and seven horizontal modules behind) and the remaining two sit off-axis and outside of the cross. INGRID samples a 10 m $\times$ 10 m area which corresponds to approximately twice the width (1σ) of the neutrino beam at the near detector hall [107]. The core of the INGRID modules consists of nine iron plates in alternation with 11 XY modules[4], bearing 7.1 t of target mass each. The XY modules consist of 24 bars in each panel, and each bar is 1.0 cm $\times$ 5.0 cm $\times$ 120.3 cm piece of polystyrene doped with 1% PPO and 0.03% POPOP. Additional XY modules with

[4]Due to weight restrictions, there are no iron plates at the 10^{th} and 11^{th} layers of each module—however, this does not sacrifice the overall tracking performance.

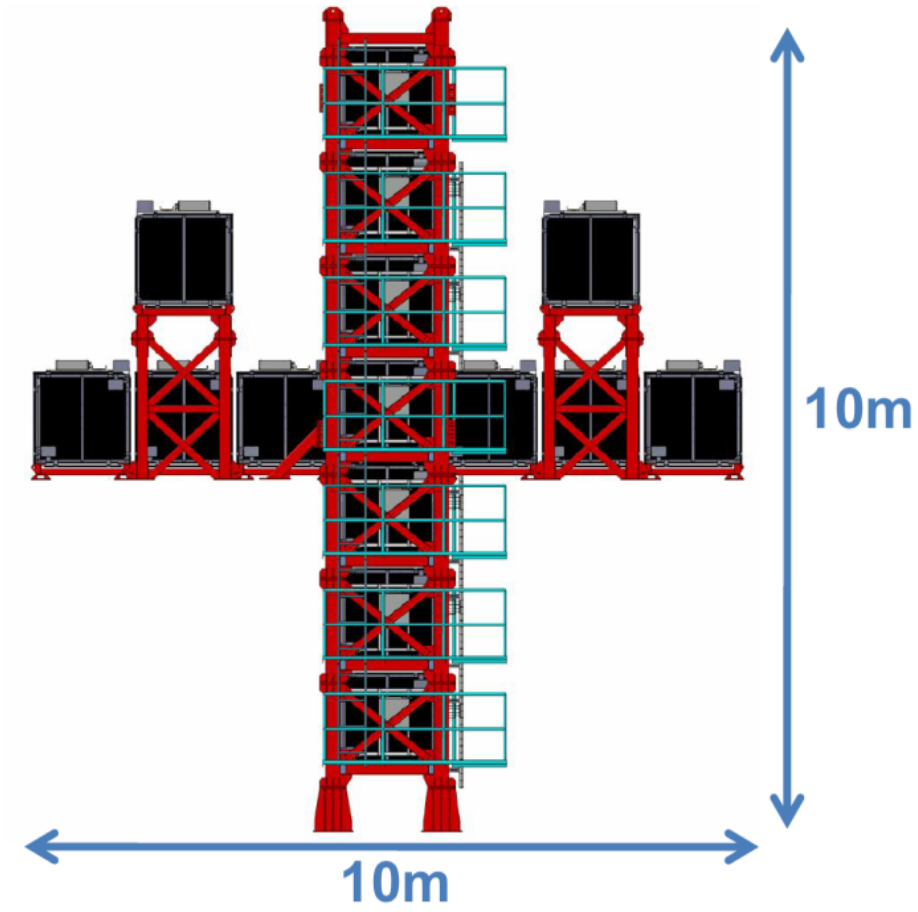

(a) A schematic diagram of INGRID's arrangement within the near detector hall (adapted from Ref. [107]).

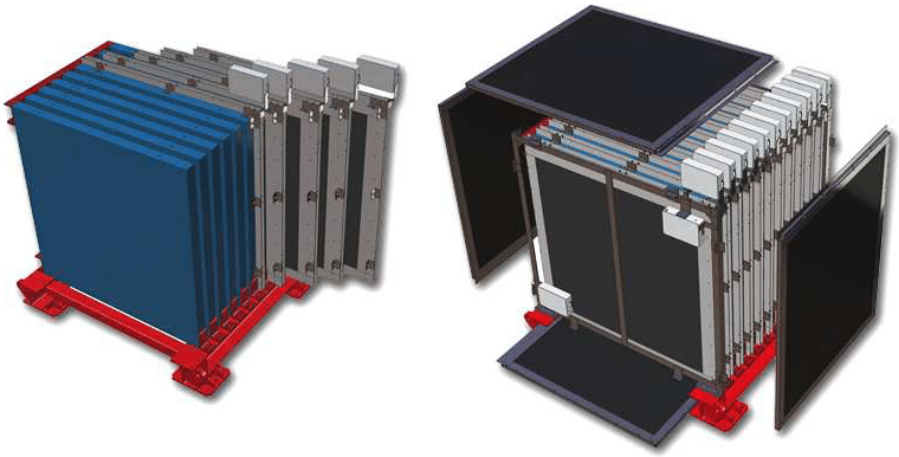

(b) Renderings of the INGRID modules without (left) and with (right) their veto planes (reproduced from Ref. [98]).

Fig. 5.10 Diagrams of T2K's INGRID detector system.

differing specifications encase the INGRID modules to reject interactions from the outside. When T2K is running at nominal beam intensity, the on-cross modules of INGRID can determine the neutrino beam centre to a precision of 10 cm (or 0.4 mrad at the near detector hall from the target) every day from the event rates found in each module. The two off-cross modules are used to probe the symmetry of the beam.

Charged particles other than μ—such as p or π—will be stopped by the iron plates in the normal INGRID modules. A special additional module lied between the centre of the horizontal and vertical modules called the Proton Module, which is used in conjunction with the rest of INGRID to measure CCQE interactions for comparison with beamline and neutrino interaction simulations, until it was replaced by the Water Module in June 2016 [108]. Unlike the INGRID modules, the Proton Module consisted only of XY modules, which are comparatively fine-grained to improve tracking, to facilitate the measurement of its namesake. The Water Module is a prototype module of the WAGASCI detector, described later in Section 5.2.3.

5.2.3 WAGASCI-BabyMIND

For oscillations to become apparent, an accurate prediction of the unoscillated spectra at the far detector is first needed. Data collected by ND280 and INGRID constrain many of the flux and cross-section parameters which form this prediction. The largest source of systematic error associated with this prediction, however, is contributed by the difference in target material (mostly CH for ND280, H_2O for SK) and difference in detector acceptance (ND280 can only see forward-going particles, SK has 4π acceptance) between the near and far detectors [109]. WAGASCI-BabyMIND (WAter Grid And SCIntillator, Baby Magnetised Iron Neutrino Detector) is a detector recently added to the near detector hall with the same acceptance and target material as SK to reduce these sources of uncertainty. WAGASCI-BabyMIND is positioned 1.5° off-axis with respect to the neutrino beam, exposed to a neutrinos with an average energy of $\sim$0.86 GeV—slightly higher than that of ND280 [110].

5.3 Far Detector Facility

After travelling 295 km through the earth, the neutrinos—now with an oscillated componen—meet with Super-Kamiokande (SK) under Mount Ikeno. As mentioned in Chapter 2, SK has enjoyed a rich history as an independent long-term experiment.

SK is a large target volume of water instrumented with photo-multiplier tubes (PMTs) and uses Cherenkov radiation to reconstruct neutrino interactions. When a charged particle

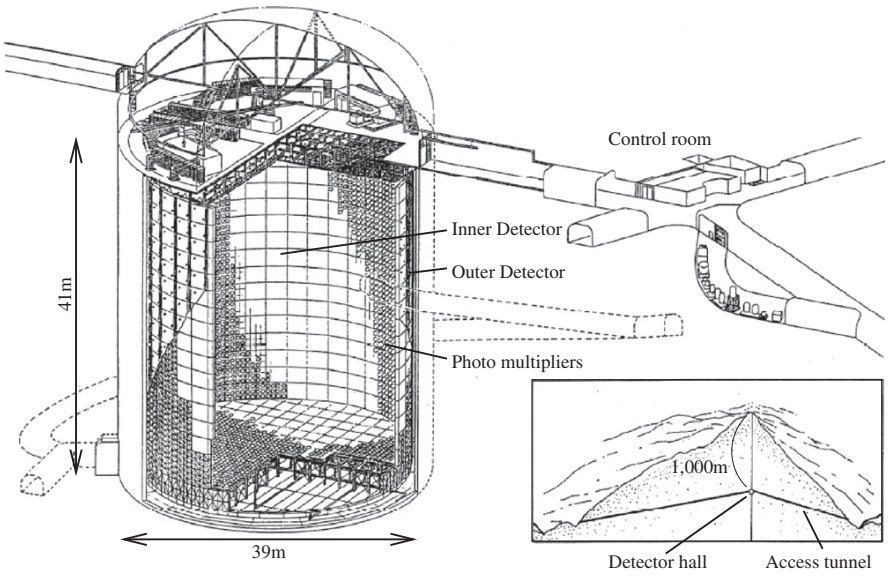

Fig. 5.11 A schematic drawing of the Super-Kamiokande detector facility (reproduced from Ref. [98]).

enters a dielectric medium where the phase velocity of light thereof is surpassed by the speed of that particle, a characteristic cone of light, known as Cherenkov radiation, is emitted. In SK, the leptons of the final state of CCQE neutrino interactions show up a ring as their Cherenkov cones are projected onto the tank's walls and collected by the PMTs. Electrons and muons are differentiated by the relative sharpness of their rings: when travelling through the tank, electrons, due to their small mass, scatter and shower and hence show up as a fuzzy ring; muons, however, as highly relativistic particles, leave rings with sharper edges. SK can differentiate electrons and muons at T2K's energy scale ($\sim$1 GeV) very well: the rate of electrons misclassified as muons is 0.7%, and 0.8% for muons misclassified as electrons [111]. The interaction position in the tank is inferred from the timing information gathered by the PMTs—the PMT timings minus the time-of-flight of the Cherenkov light should be tightly collected. From the sum of PMT charges over a ring, the momentum of the lepton may be found.

SK holds 50 kt of ultra-pure water within a cylindrical tank 41.4 m in height and 39.3 m in diameter. A stainless steel superstructure, onto which the PMTs are mounted, divides the tank into an inner detector (ID) region and outer detector (OD) region. The ID, which is the physics region, is well-instrumented and holds 11,146 20" PMTs (40% coverage) which face inwards; the OD, used to reject particles entering the detector from the outside, is comparatively sparsely instrumented and holds 1,885 8" PMTs which face outwards. These

detector regions are optically separated by a black sheet facing the ID and white Tyvek facing the OD side.

A charged particle traversing a particular medium induces a Cherenkov cone with an opening-angle θ given by

$$\cos\theta = \frac{1}{\beta n} \tag{5.2}$$

where β is the particle's velocity as a fraction of the speed of light in vacuum and n is the refractive index of the medium. In SK—where the leptons in the final state of CCQE neutrino interactions are ultra-relativistic—the opening-angle of a typical Cherenkov cone is 42° ($n = 1.34$, $\beta \approx 1$).

5.4 T2K-II and Near Detector Upgrades

T2K will increase its accumulated dataset from 7.8×10^{21} POT to 20×10^{21} POT in an extended running period known as T2K-II. During this period, running from 2022 to 2026 [112, 113], T2K will operate with an upgraded ND280 and the beam power of J-PARC's MR will be gradually increased to meet the requirements of Hyper-Kamiokande. T2K-II will be superseded by the T2HK (Tōkai-to-Hyper-Kamiokande) configuration when Hyper-Kamiokande comes online.

5.4.1 Beamline

As seen in Fig. 5.12, the power of J-PARC's MR will be first ramped up from the current operating power of $\sim$500 kW to the original design power of 750 kW following a power supply upgrade in 2022, then to 1.3 MW—and eventually beyond—following an RF upgrade in 2024.

The neutrino beamline in its present state is essentially capable of accepting a beam power of 1.3 MW—though some upgrades, mostly limited to its cooling and monitoring systems, are planned. The toroidal horns along the secondary beamline, however, will be upgraded from 250 kA to 320 kA by increasing the number of horn power supplies from two to three—granting a 10% increase in neutrino flux and a reduction in wrong sign contamination.

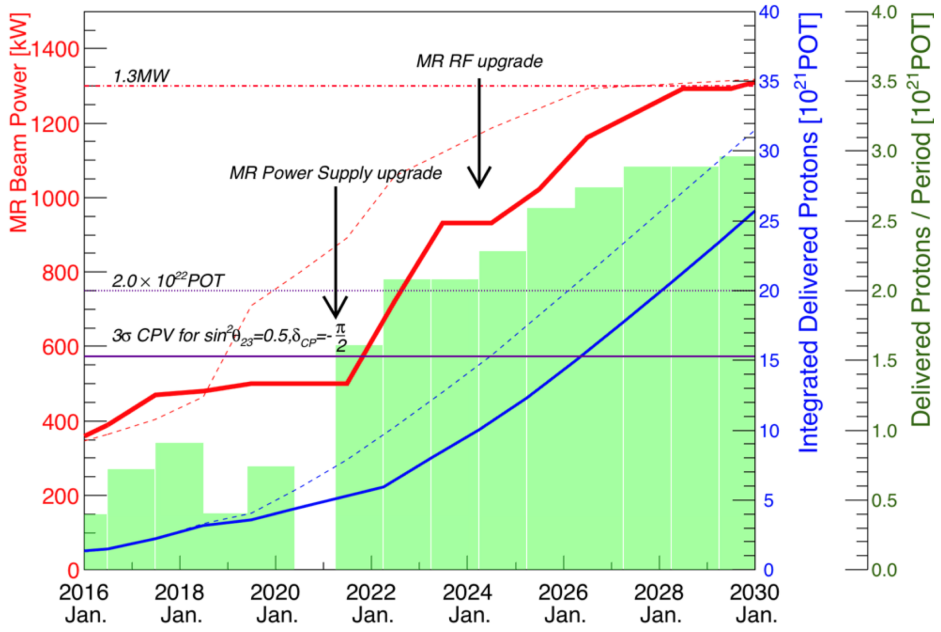

Fig. 5.12 Target beam power and accumulated POT of J-PARC's MR upgrade (reproduced from Ref. [112]).

5.4.2 Near detector

ND280 will be upgraded in the latter half of 2022 [113] by replacing most of the P0D with a new system of subdetectors to achieve 4π acceptance; the rest of the detector will be untouched. As seen in Fig. 5.13, in the place of the P0D's central module and downstream ECal will be a Super-FGD module sandwiched vertically between two High-Angle TPCs (HA-TPCs)—the P0D's upstream ECal will remain. This system, occupying roughly a $2\,\mathrm{m}^3$ cubic volume, will be surrounded on all faces by Time-Of-Flight (TOF) detectors. The upgraded ND280 will be able to better constrain the systematic uncertainties relevant to the oscillation analysis during T2K-II and T2KHK. The increased efficiency for reconstructing high-angle and backwards-going muons in particular increases ND280's sensitivity to important nuclear effects than the current configuration, thus better constraining the interaction model uncertainties—which is the dominant source of uncertainty—affecting the predicted number of signal events at the far detector. It is anticipated will be able to predict the signal event rate at SK to better than 4% [114] in the upgraded configuration.

XY modules, formed of tracking planes, possess a limited acceptance and evidently prefer particles that travel along the axis of which the tracking planes are arranged. The Super-FGD, however, is formed of many $1\,\mathrm{cm} \times 1\,\mathrm{cm} \times 1\,\mathrm{cm}$ cubes of plastic scintillator that are instrumented along all three orthogonal axes with WSL fibres, and hence possesses no such directional preference. The cubes will arranged $192 \times 192 \times 56$—totalling 2,064,384

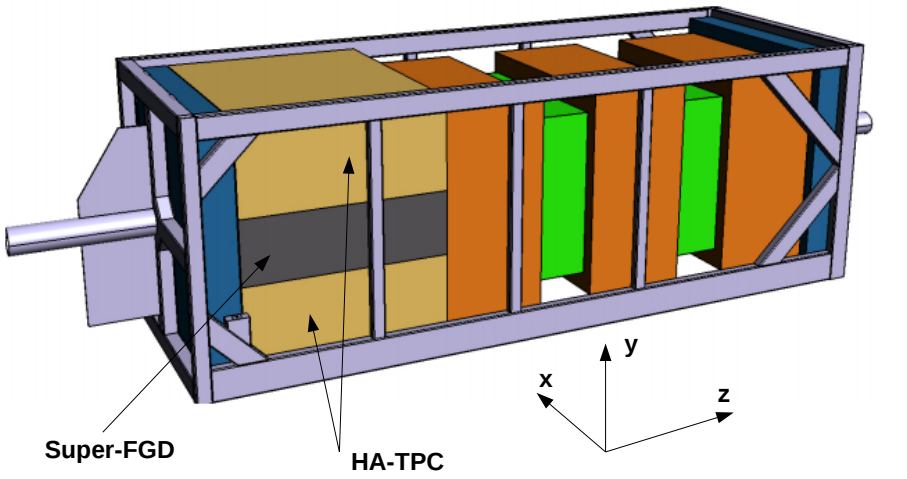

Fig. 5.13 A CAD drawing of the inner detectors of the upgraded ND280 detector for T2K (adapted from Ref. [114]).

cubes with 58,368 read-out channels. Simulation studies indicate that the reconstruction efficiency for μ will approach 90% for all particle directions. For protons, the efficiency also approaches $\sim$90% (for protons with momenta above $\sim$400 MeV) and the detection threshold is as low as $\sim$300 MeV.

As the performances of the TPCs currently found in ND280 are satisfactory, the designs of the HA-TPCs do not differ substantially. The HA-TPCs will make use of enhanced Micromegas detectors inside a thinner field cage design that maximises the tracking volume within the chamber.

The TOF modules have an excellent timing resolution ($\sim$600 ps) to measure the direction of particles. Each TOF module is a single panel of scintillator bars read out by arrays of Si photo-multipliers adhered directly to the bar.

CHAPTER 6

ANALYSIS STRATEGY

The goal of this analysis is to extract the neutrino–proton NC0π differential cross section on carbon nuclei as a function of the kinematics—the kinetic energy and polar angle—of the highest momentum outgoing proton using the ND280 tracker. The theoretical significance and experimental history of this particular channel was discussed earlier in Section 3.2.2. The interaction topology of interest is a neutrino of any flavour that interacts with a proton via the neutral current to give at least one proton in the final state and no other particles (other than a neutrino of identical flavour as the original neutrino). In this analysis, FGD1—which is made of more than 85% carbon [115]—is purposed as the target material.

The stages of the cross section analysis, later delivered in Chapter 8, are as follows. Firstly, the backgrounds contaminating the sample of signal-like events will be extracted and subtracted from said sample. As it will later be seen in Section 8.2, the levels of background contaminating the signal sample are determined by fitting the nominal Monte-Carlo to the data in the signal sample and various control samples simultaneously. The background levels determined by said fit will then be validated in Section 8.3 with an independent calculation of the background contributed by interactions of out-of-fiducial-volume neutrons by studying the arrival times of tracks within the FGD. Next, in Section 8.4, in a treatment otherwise known as *unfolding*, the newly background-subtracted spectra of signal events will be corrected for detector effects—the detector's resolution, efficiencies and thresholds—to render the true number of expected signal events. Lastly, in Section 8.5, the differential cross sections with respect to the outgoing proton kinematics will be calculated from the unfolded spectra.

The measurement was first performed *blind* to avoid experimenter's bias: data events in the signal sample were substituted for the nominal MC; data in the control samples, however, were retained. The analysis was performed again with data events reintroduced into the

signal sample once the blinded measurement was performed successfully. The analysis shows results from both the blinded and unblinded analyses simultaneously.

As well, both analyses were performed with a signal sample constructed from a series of cuts and a range of signal samples designated by a multivariate signal selection. The development of this multivariate selection is documented in Chapter 7. Unlike the cut-based selection, which produces a boolean value according to whether a given event appears to be signal or not, the multivariate selection calculates a continuous probability score. Thus, the analysis performed with the multivariate selection has an additional free parameter: the threshold applied to the signal probability scores that ultimately renders the signal sample. One objective of this analysis is to demonstrate how much the multivariate selection improves the precision of the cross section measurement over the cut based selection, which entails minimising the uncertainty of the cross section measurement based on the threshold applied. However, how the measurement responds to the makeup of the signal sample, and hence the optimum value of the applied threshold, was not known a priori. Rather, the optimal threshold value was found empirically by performing the analysis many times over a range of applied threshold values. Where possible, the results of the multivariate analysis will be shown in comparison to those of the cut-based analysis.

6.1 Signal Definition

The signal is identified by one or more proton tracks that are accompanied by no other types of reconstructed particles, that appear in a fiducial volume defined within FGD1 and cross over into the adjacent TPC, TPC1. The fiducial volume is defined as the region within the ND280 global coordinate system[1] that satisfies $|x| < 874.51\,\text{mm}$, $|y - 55\,\text{mm}| < 874.51\,\text{mm}$ and $136.875\,\text{mm} < |z| < 447.375\,\text{mm}$. There are a few classes of interactions that mimic this definition of signal and thus contaminate the sample of signal events, these backgrounds are

- **Out-Of-Fiducial-Volume (OOFV) Neutrons**: The largest background is contributed by interactions of neutrons, that are produced by neutrino interactions outside of the fiducial volume, which then cross into the fiducial volume to interact to give one or more protons in the final state. It is because neutrons are invisible to ND280 that this type of event is difficult to reject.

[1] The ND280 coordinate system is orientated such that $\hat{z}$ is aligned with the beam direction, $\hat{y}$ is vertical and $\hat{x}$ is horizontal.

- **Charged Current (CC) Interactions**: The second-largest background is contributed by charged-current neutrino interactions within FGD1 that give in the final state at least one proton reconstructed by ND280 and a lepton that evades detection.
- **Neutral Current (NC) Backgrounds**: A further source of background are neutrino–proton interactions via the neutral-current that occur within the fiducial volume to produce pions or other particles.

The overwhelming majority of the background are contributed by these three interaction types. The small remainder of background events are made up of mainly OOFV protons and OOFV sand interactions.

6.2 Dataset

The analysis uses the complete neutrino mode ND280 dataset[2]. These data include both configurations of the P0D (i.e. filled with either water or air). The accumulated number of protons on target per run are given in Table 6.1.

		Protons On Target	
Run No.	Run Date	Data	MC
2	Nov. 2010–Mar. 2011	7.93×10^{19}	2.53×10^{20}
3	Mar. 2012–Jun. 2012	1.58×10^{20}	9.02×10^{20}
4	Oct. 2012–May 2013	3.43×10^{20}	1.66×10^{21}
8	Oct. 2016–Apr. 2017	5.73×10^{20}	3.10×10^{21}
Total		1.15×10^{21}	5.92×10^{21}

Table 6.1 Number of protons on target accumulated per T2K run in the data and MC datasets.

Various processings of detector data and MC are readily available to T2K analysers; this analysis uses the processings generally recommended for physics analyses and that contain the most up-to-date calibrations[3]. Various corrections are applied to these datasets—some

[2]Specifically, these are runs 2, 3, 4 and 8. Run 1 is excluded as the near-detector operated without the barrel-ECal and with a different beam configuration compared to the remainder of the dataset.

[3]The real data for runs 2–4 are drawn from release 6M, which was produced with version 11.31.5 of the ND280 software; the real data for run 8 is drawn from release 6P, which was produced with version 11.31.27 of the ND280 software. The simulation data for runs 2–4 are drawn from release 6B, which was produced with version 11.31.0 of the ND280 software; the simulation data for run 8 is drawn from release 6L, which was produced with version 11.31.27 of the ND280 software. The MC in all runs was generated with version 2.8.0 of the GENIE event generator.

of which are considered standard [116], some of which are particular to this analysis. The standard corrections include:

- **Data quality**: includes only runs that were taken with good beam quality and with all near-detector subdetectors fully operational.
- **Energy loss of the TPC**: as electron interactions in the TPC are poorly modelled in GEANT4, some discrepancies between data and MC are seen in the distributions of the TPC particle identification variables. The differences in dE/dx response from TPC to TPC as measured in data are calibrated out and corrections are applied to the MC to achieve better data–MC agreement [117].
- **Momentum resolution of the TPC**: some discrepancy between data and MC is seen in the momentum resolution of the TPC. The momentum is smeared in MC to achieve better agreement [118].

The non-standard corrections include:

- **Integration of sand interactions**: neutrino interactions in the near-detector and its surroundings are simulated independently in the standard processings and are typically merged later at the histogramming level. In the data, however, these interactions occur simultaneously; hence—since the selection logic of the analysis samples consider the track multiplicity measured in some of the surrounding sub-detectors—this typical treatment prompts some events to migrate between the analysis samples at the histogramming stage. The sand and detector interactions are instead combined at the analysis level.
- **Pile-up**: the efficiencies of the subdetector activity veto cuts mentioned later are subject to the rate of pile-up interactions in the detector; hence some discrepancies between data and MC are seen in these efficiencies as the rate of pile-up interactions are not sufficiently modelled. The probability distribution describing the number of near-detector interactions per beam spill was remodelled based on the total number of delivered protons on target in the real detector data used in this analysis. The same was performed for the number of sand muons that enter the near-detector. For each of these, the correction is performed by re-weighting MC by the ratio between the nominal and data-driven distribution.

	Total No. Events	
Sample	Data	MC
Signal (cut-based)	1691	1768
Signal (ML, efficiency enhanced)	3292	3166
Signal (ML, purity enhanced)	1352	1278
P0D (cut-based)	1009	979
SMRD (cut-based)	741	726
Michel Electron (cut-based)	746	757

Table 6.2 Total number of data and pre-fit Monte-Carlo events in each analysis sample.

6.3 Event Selections

The four analysis samples—the signal sample and the three control samples—are drawn from the corrected datasets. As mentioned, there are two implementations of the signal selection: a cut-based one and a multivariate one. The control sample cuts are identical in the analyses performed with both signal selections; however, it will be shown later in this section that since the control samples are restricted from overlapping with the signal sample, the control samples can be different depending on the signal probability threshold applied to render the signal sample. The three control samples are made background contaminated by selecting proton-like tracks reconstructed in the tracker with additional reconstructed objects in the same event. An object is any type of reconstructed activity in any subdetector: for example, a shower, a track, a cluster of hits and so on. The P0D control sample requires that at least one object is reconstructed in the P0D, the SMRD control sample requires at least one object reconstructed in the SMRD and the Michel electron sample requires at least one Michel electron to be reconstructed in the FGD1 fiducial volume.

In Table 6.2 the number of events per analysis sample are given. Since all of the cut-based analysis samples derive from a sample of proton candidates, the construction of that sample will be described first.

6.3.1 Proton sample

The candidate primary proton track is the highest momentum track starting in the FGD1 fiducial volume with at least one matched track in the TPC with > 18 hits. The matched TPC tracks are also required to have momenta of less than 30 GeV; TPC tracks with momenta in excess of this limit are unphysical and indicative of a reconstruction failure. The following conditions are placed on the candidate proton to construct the proton sample:

- **Has good data quality**: all subdetectors across ND280 must be operating as standard at the time of the event.
- **Has positive charge**: the candidate proton track must be reconstructed with positive charge.
- **Is proton-like in TPC**: the TPC track of the candidate proton must have a proton-like particle-identification pull with magnitude satisfying < 3.5. The pull that a track bears to a particular particle type i is based on the truncated mean of the energy loss of the track as it crosses the TPC [119]. The pull is defined as $(c^{obs.} - c_i^{exp.})/\sigma_i^{obs.}$, where $c^{obs.}$ is the observed energy loss, $c_i^{exp.}$ is the expected energy loss given that the track was caused by particle type i and $\sigma_i^{obs.}$ is the uncertainty of the observed energy loss given that the track was caused by particle type i.
- **Has reasonable kinetic energy**: the reconstructed kinetic energy of the candidate proton track must not exceed 1 GeV, as high-momenta μ^+ and p are indistinguishable in the TPC.

6.3.2 Cut-based signal sample

The cut-based signal sample is essentially comprised of events with at least one proton-like track in the tracker but no other reconstructed objects anywhere else in the near-detector. In addition to those of the proton selection, the signal selection includes the following requirements:

- **Has no other tracks in the TPC with negative charge**: there must be no reconstructed objects in the event, other than the candidate protons, that have a TPC track with a negative charge. Specifically, this cut targets the removal of events with μ^-, indicative of a CC interaction.
- **Has no other tracks in the TPC with positive charge**: there must be no reconstructed objects in the event, other than the candidate protons and objects with a common

vertex with the candidate proton tracks, that have a TPC track with a positive charge. Objects are said to possess a common vertex with the leading proton candidate if their starting positions lie within 300 mm of each other. Specifically, this cut targets the removal of OOFV events where the proton track is wrongly reconstructed (e.g. the track of a proton that traverses TPC1, FGD1 and TPC2 in truth, rather than be correctly reconstructed as a single track, could be reconstructed as two separate tracks at the TPC1–FGD interface with one which appears to originate within the FGD1 fiducial volume).

- **Has no muon-like tracks in the TPC**: there must be no reconstructed objects in the event, other than the candidate proton tracks, that have μ-like particle-identification pull with magnitude satisfying ≤ 2.9. Specifically, this cut targets the removal of events with μ^-, indicative of a CC interaction.

- **Has no other reconstructed objects in the P0D, ECal and P0D-ECal**: there must be no reconstructed objects in the event, other than the candidate proton and objects with a common vertex with the candidate proton track, that have tracks in the P0D, ECal and P0D-ECal. Specifically, this cut targets the removal of OOFV events.

- **Has no other reconstructed objects in the FGD**: there must be no reconstructed objects in the event, other than the candidate protons, that have tracks in FGD1.

- **Has no Michel electrons**: the e^- produced when a μ^- decays at rest ($\mu^- \to \nu_\mu \bar{\nu}_e e^-$) is a *Michel electron*. These are identified in the FGD as charged clusters with more than six hits >100 ns after the bunch containing the selected proton. There must be no reconstructed objects in the event with Michel electrons. Specifically, this cut targets the removal of events with a μ^- produced via a CC interaction that is stopped by FGD1.

- **Proton-like track is not at the last FGD1 layer**: the z-axis component of the starting position of the candidate proton track must not exceed 400 mm with respect to the ND280 global coordinate system. Specifically, this cut targets the removal of objects produced by interactions of OOFV neutrons. These interactions are likely to produce low-energy protons which can only cross into the TPC—and hence pass the selection—if the interaction occurred at the most downstream layers of the FGD.

The cut flow for the signal selection is given in Table 6.3. As seen from Fig. 6.1, this sample is dominated by true signal, but neutrons produced from interactions outside of the fiducial volume then produce interactions in FGD1 form a significant background. The P0D

and SMRD control samples are windows to measure this contribution directly from the data, as it is not expected of the MC to predict the neutron rates correctly. The second highest background is CC events where the μ^- failed to be reconstructed; in a similar way, the Michel electron control sample is used to also measure this background directly.

Cut	Data	Total MC	Signal	NC bkgd	CC	OOFV Neutron	Other	Purity	Efficiency
+ve TPC charge	4187	5722	2446	324	1206	1534	212	0.43	0.20
TPC proton pull	4187	5376	2100	324	1206	1534	212	0.39	0.17
Proton E_{K} [MeV]	4187	5313	2036	324	1206	1534	212	0.38	0.16
-ve TPC track veto	4187	5098	1821	324	1206	1534	212	0.36	0.15
TPC muon veto	4187	4965	1689	324	1206	1534	212	0.34	0.14
ECal veto	3038	3552	1307	234	889	1027	95	0.37	0.11
FGD veto	2949	3447	1268	230	867	1010	72	0.37	0.10
P0D veto	2579	3031	1217	209	849	704	51	0.4	0.10
P0D ECal veto	2437	2848	1158	206	842	599	43	0.41	0.09
+ve TPC track veto	2437	2837	1147	206	842	599	43	0.4	0.09
Michel veto	1691	2049	1015	164	299	540	32	0.5	0.08
Reduced fiducial volume	1691	1768	734	164	299	540	32	0.41	0.06

Table 6.3 Number of events passing each cut in the selection of the cut-based signal sample.

6.3.3 P0D control sample

The P0D control sample includes all cuts placed on events of the signal selection, with the following exceptions:

- **Has a reconstructed object in the P0D**: there must be at least one reconstructed object in the P0D; the corresponding veto in the signal selection is inverted.
- **May or may not contain reconstructed objects in the P0D-ECal, ECal and FGD**: activity in other subdetectors is permitted; the vetoes placed on sub-detector activity in the signal sample are omitted.

The cut flow for the P0D control sample selection is given in Table 6.4. In Fig. 6.2, it is seen that this sample is dominated by OOFV neutrons that originate in the P0D.

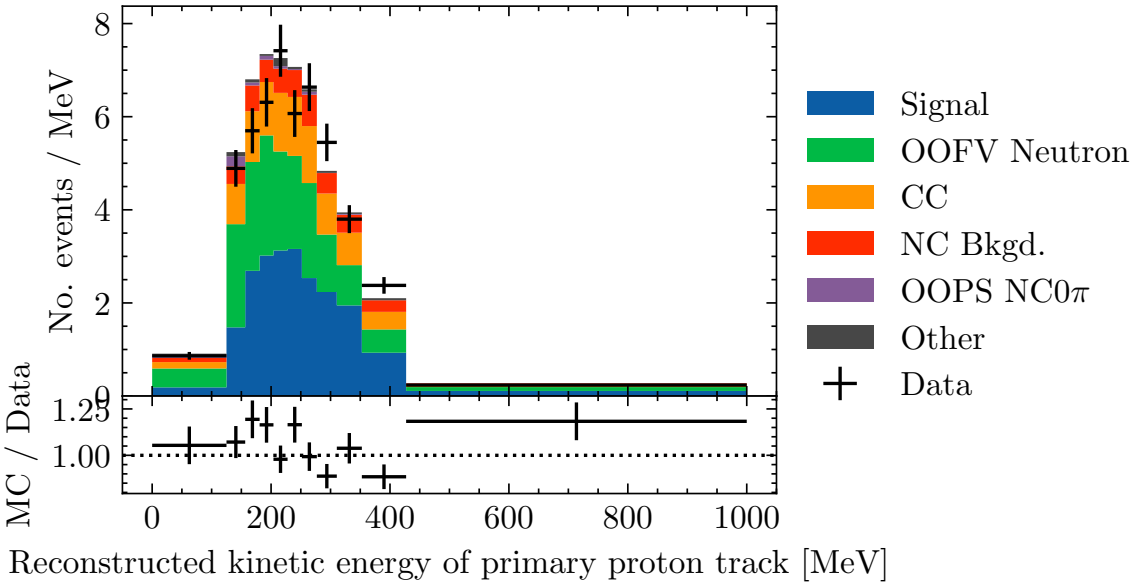

(a) Reconstructed kinetic energy.

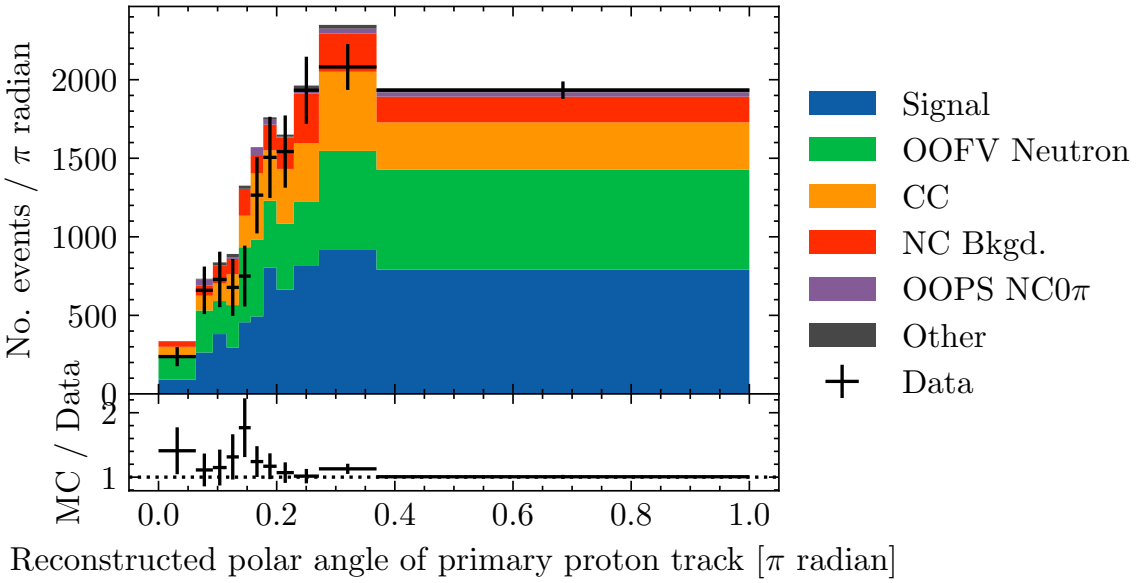

(b) Reconstructed polar angle.

Fig. 6.1 Reconstructed kinematics of primary proton candidates within the signal sample.

Cut	Data	Total MC	Signal	NC bkgd	CC	OOFV Neutron	Other
+ve TPC charge	4187	5722	2446	324	1206	1534	212
TPC proton pull	4187	5376	2100	324	1206	1534	212
Proton E_{K} [MeV]	4187	5313	2036	324	1206	1534	212
-ve TPC track veto	4187	5098	1821	324	1206	1534	212
TPC muon veto	4187	4965	1689	324	1206	1534	212
Inverted P0D veto	1009	1012	95	63	125	653	75
+ve TPC track veto	1009	1008	90	63	125	653	75
Michel veto	1009	999	82	63	125	653	75
Reduced fiducial volume	1009	979	62	63	125	653	75

Table 6.4 Number of events passing each cut in the selection of the P0D control sample.

6.3.4 SMRD control sample

The SMRD control sample includes all cuts of the signal selection, with the following exceptions and additions:

- **Has a reconstructed object in the SMRD**: there must be at least one reconstructed object in the SMRD.
- **SMRD object begins upstream**: the z-axis component of starting position of the reconstructed SMRD object must be less than 136 mm with respect to the ND280 global coordinate system.
- **May or may not contain reconstructed objects in the ECal, FGD**: activity in other subdetectors is permitted; the other vetoes placed on sub-detector activity in the signal sample are omitted.
- **Must not pass the signal selection**: the signal and SMRD sample must be disjoint, as required by the sideband fit.

The cut flow for the SMRD control sample selection is given in Table 6.5. In Fig. 6.3, it is seen that this sample is dominated by OOFV neutrons that originate in the near-detector magnet.

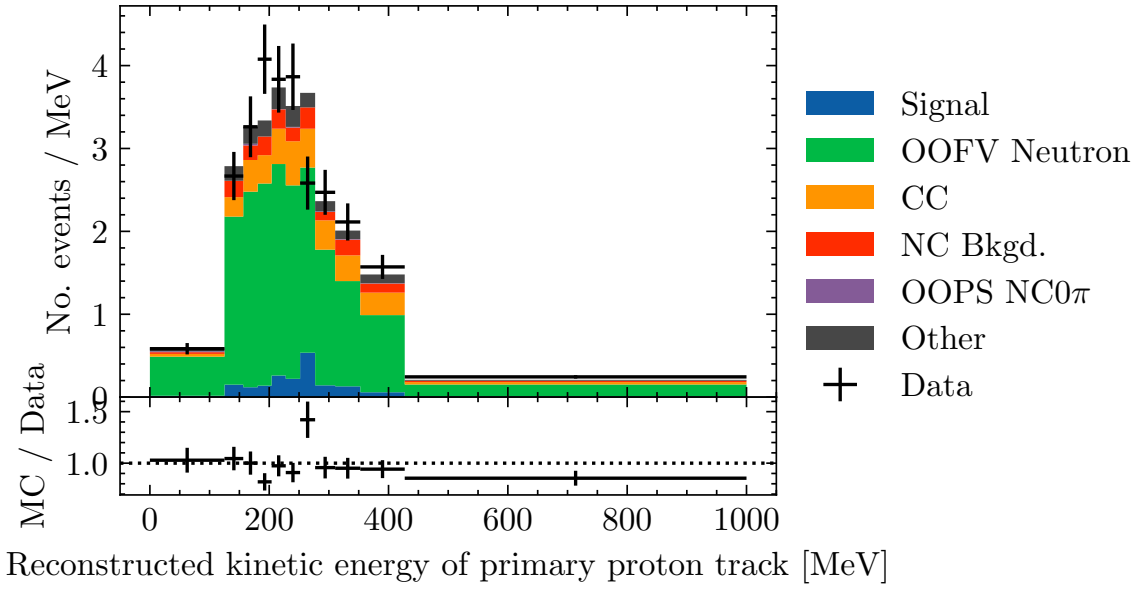

(a) Reconstructed kinetic energy.

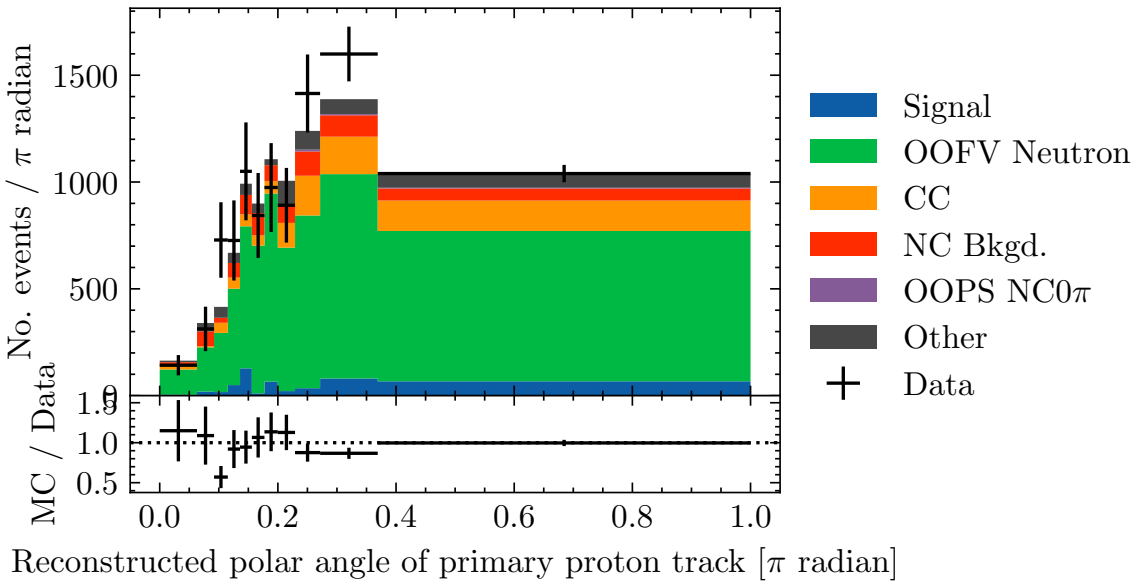

(b) Reconstructed polar angle.

Fig. 6.2 Reconstructed kinematics of primary proton candidates within the P0D control sample.

Cut	Data	Total MC	Signal	NC bkgd	CC	OOFV Neutron	Other
+ve TPC charge	4187	5722	2446	324	1206	1534	212
TPC proton pull	4187	5376	2100	324	1206	1534	212
Proton E_{K} [MeV]	4187	5313	2036	324	1206	1534	212
-ve TPC track veto	4187	5098	1821	324	1206	1534	212
TPC muon veto	4187	4965	1689	324	1206	1534	212
P0D veto	3178	3953	1594	261	1081	881	137
+ve TPC track veto	3178	3901	1542	261	1081	881	137
Michel veto	2432	3070	1367	219	537	822	125
Reduced fiducial volume	2432	2702	999	219	537	822	125
Inverted SMRD veto	1010	991	113	70	265	439	103
Is not signal?	741	726	57	55	239	282	94

Table 6.5 Number of events passing each cut in the selection of the SMRD control sample.

6.3.5 Michel electron control sample

The Michel electron control sample includes all cuts of the signal selection, with the following exception:

- **Has Michel electrons**: there must be at least one reconstructed Michel electron; the corresponding veto in the signal selection is inverted.

The cut flow for the Michel electron control sample selection is given in Table 6.6. In Fig. 6.4, it is seen that this sample is dominated by CC interactions that produce μ^- that stop in FGD1.

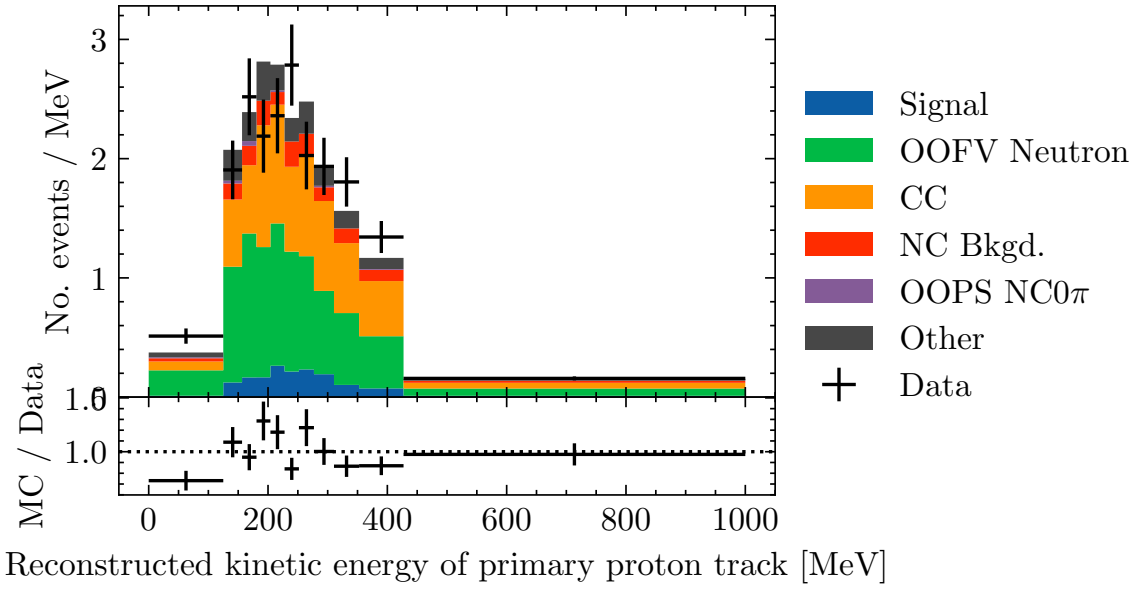

(a) Reconstructed kinetic energy.

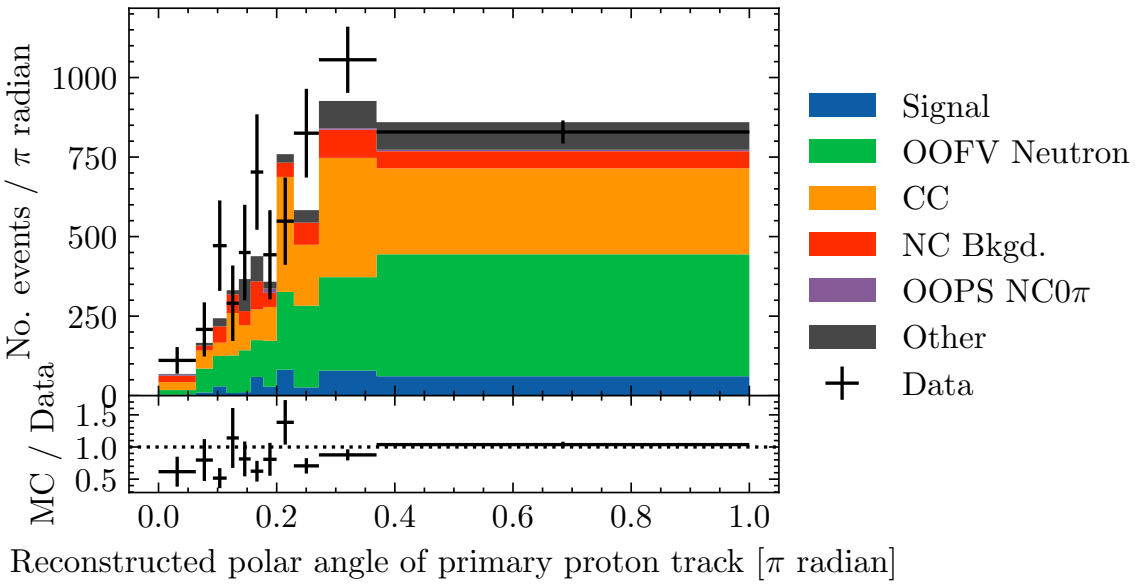

(b) Reconstructed polar angle.

Fig. 6.3 Reconstructed kinematics of primary proton candidates within the SMRD control sample.

Cut	Data	Total MC	Signal	NC bkgd	CC	OOFV Neutron	Other
+ve TPC charge	4187	5722	2446	324	1206	1534	212
TPC proton pull	4187	5376	2100	324	1206	1534	212
Proton E_{K} [MeV]	4187	5313	2036	324	1206	1534	212
-ve TPC track veto	4187	5098	1821	324	1206	1534	212
TPC muon veto	4187	4965	1689	324	1206	1534	212
ECal veto	3038	3552	1307	234	889	1027	95
FGD veto	2949	3447	1268	230	867	1010	72
P0D veto	2579	3031	1217	209	849	704	51
P0D ECal veto	2437	2848	1158	206	842	599	43
+ve TPC track veto	2437	2837	1147	206	842	599	43
Inverted Michel veto	746	788	132	42	543	59	11
Reduced fiducial volume	746	757	101	42	543	59	11
Is not SMRD sample?	746	757	101	42	543	59	11

Table 6.6 Number of events passing each cut in the selection of the Michel electron control sample.

6.3.6 Multivariate samples

The development of the multivariate signal classifier is detailed in the next chapter, Chapter 7. As mentioned, the multivariate signal classifier assigns a probability score ranging from zero to one to each event based on how closely the event resembles the signal, lower scores being background-like and higher scores being signal-like. There is a trade-off in selection performance based on the threshold applied: the higher the threshold applied, the higher the purity but the lower efficiency, and vice versa. The composition of the signal sample as a function of the signal probability threshold applied is shown in Fig. 6.5a. For the vast majority of possible thresholds, the various backgrounds overwhelm the sample hence the signal purity is minuscule; only at sufficiently high thresholds ($\lesssim 0.65$) does the selection performance become superior to that of the cut-based selection. The range of threshold values was chosen such that both an improved signal purity and efficiency would be achieved relative to the cut based selection. The minimum threshold value was 0.646, which yields a sample with an identical purity (0.401) and enhanced efficiency (0.313) to the cut based selection; the maximum threshold value was 0.705, which yields a sample with identical efficiency (0.173) and enhanced purity (0.552) to the cut based selection. Sixteen evenly spaced thresholds were chosen—the results of each multivariate analysis will be shown simultaneously with the results yielded from the cut based analysis where possible. In Fig. 6.6, the distribution of signal probability scores for events that pass the preselection are shown over this range; in Fig. 6.7, the achieved purities and efficiencies corresponding to these thresholds are shown.

The cut-based samples are guaranteed to be disjoint by the usage of vetoes; the multivariate samples, however, are not, as the multivariate model is trained on the variables that these vetoes concern. Put differently, the multivariate selection is free, for example, to identify as signal events accompanied by activity in the other subdetectors or by Michel electrons in the FGD. Given this, additional vetoes are applied to the control samples that ensure that they do not intersect with the multivariate signal sample for any signal probability threshold applied. As a result, the composition of the control samples also change with the threshold applied. The compositions of the multivariate analysis samples are shown in Fig. 6.8; their contributions by share of the total sample are shown in Fig. 6.9. In the next chapter, it is shown that the enhanced purity of the multivariate selection arises from an improved rejection of CC interactions and that the rejection of the dominant background—OOFV neutrons—is about the same as the cut-based selection. These observations are reflected in Fig. 6.9, in which it can be seen that, as a proportion of the selected sample, the number of CC events decreases whereas those of OOFV neutrons and NC background are approximately constant with respect to the threshold applied. As a proportion of the selected sample, a 55% reduction in CC background relative to that of the cut-based selection is achieved for the purest sample.

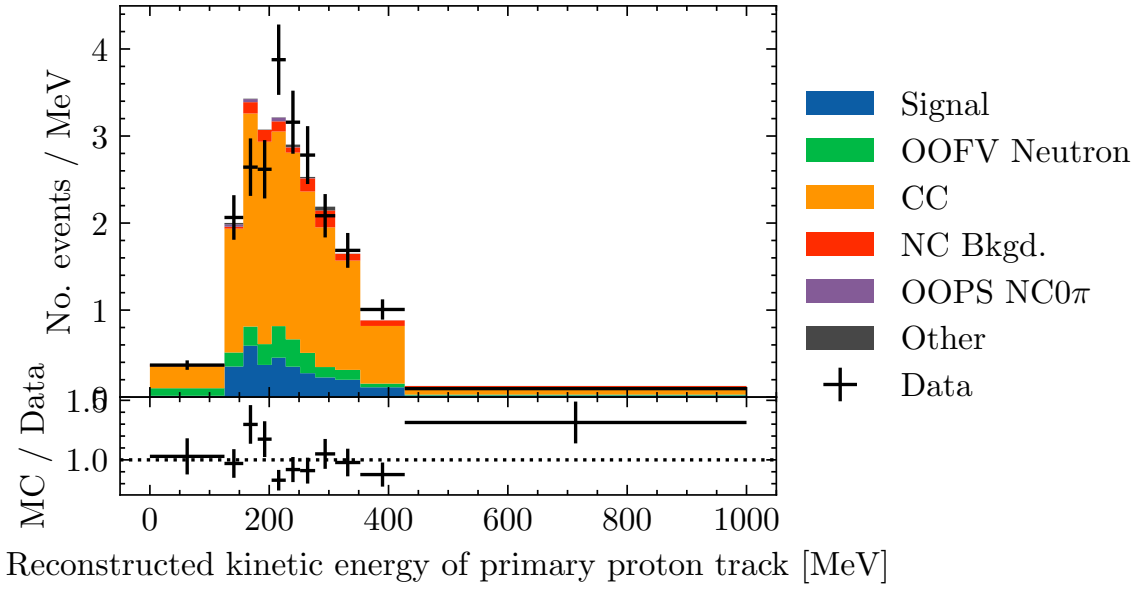

(a) Reconstructed kinetic energy.

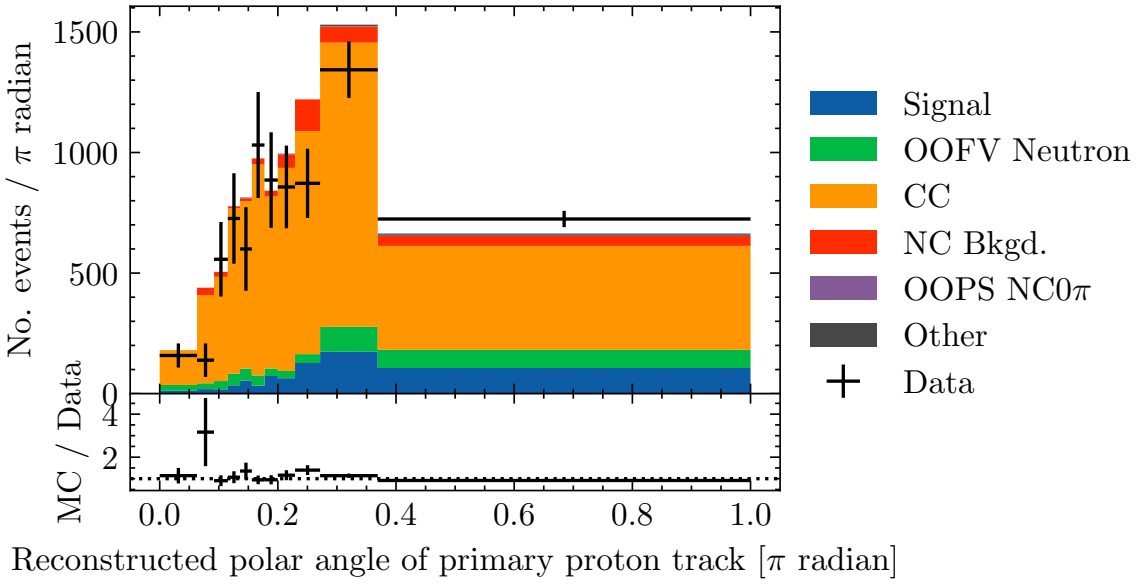

(b) Reconstructed polar angle.

Fig. 6.4 Reconstructed kinematics of primary proton candidates within the Michel electron control sample.

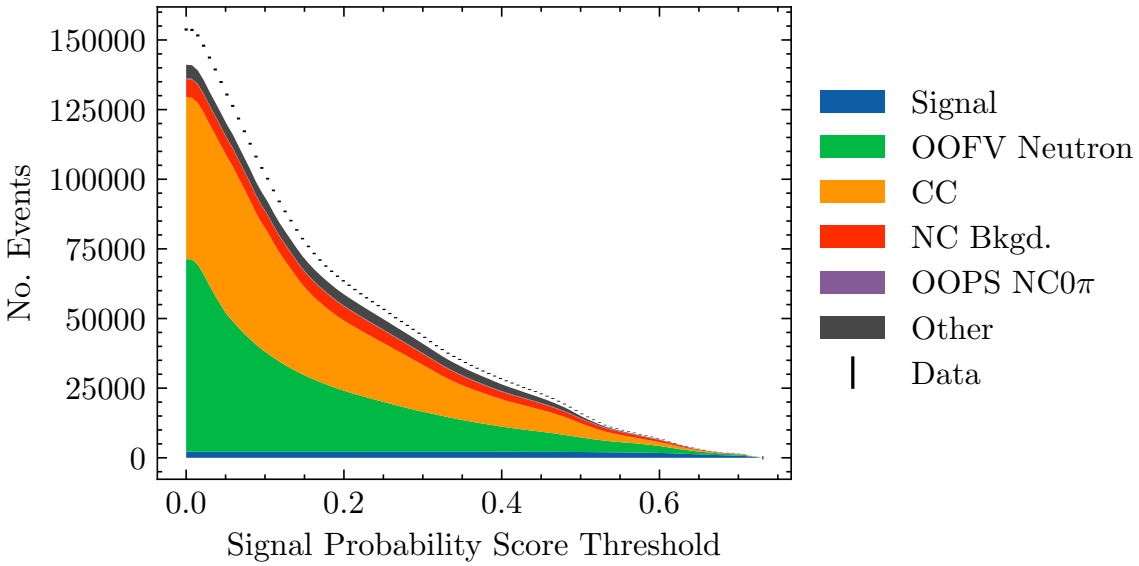

(a) Full spectrum.

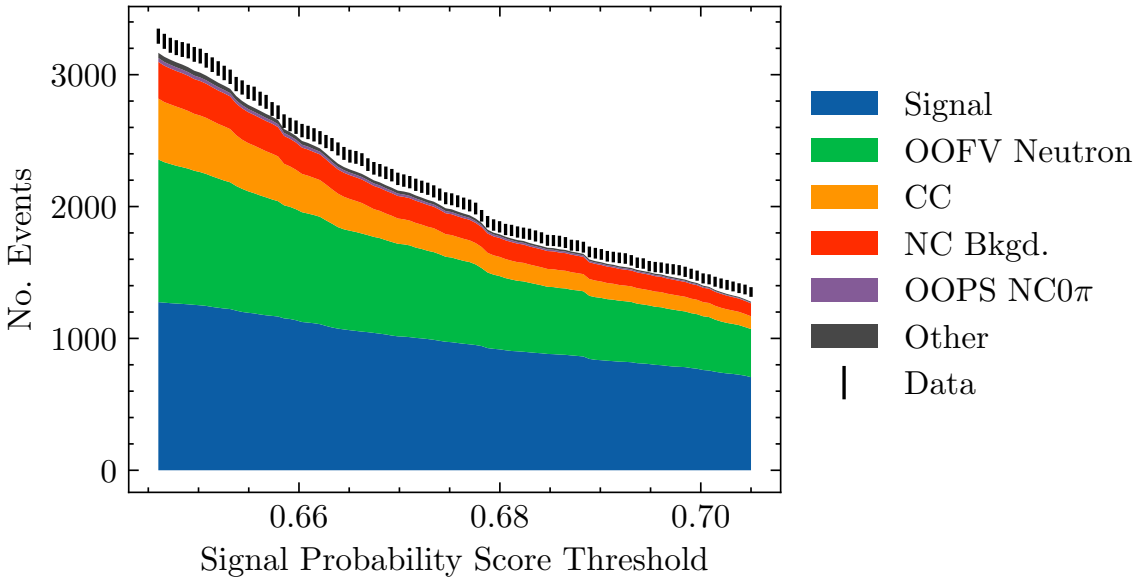

(b) Region of enhanced signal purity and/or efficiency.

Fig. 6.5 Composition of the signal sample as function of threshold applied signal probability score for events that pass the multivariate preselection. Fig. 6.5a shows the full spectrum, whereas Fig. 6.5a shows the range of thresholds that achieve either an improved purity or efficiency compared to the cut-based selection.

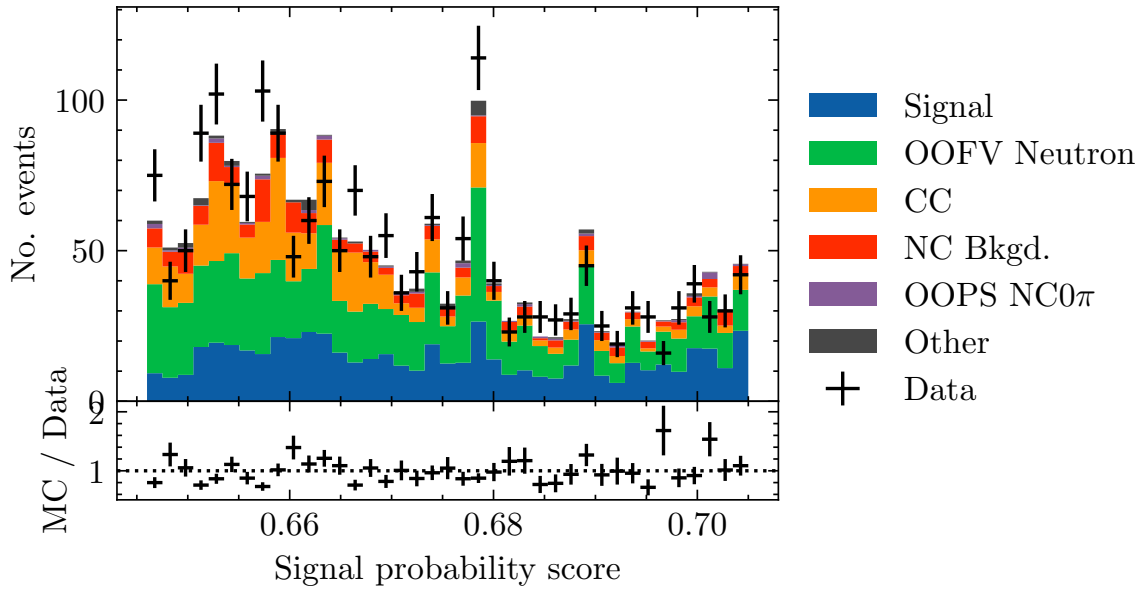

Fig. 6.6 Distribution of signal probability scores for events that pass the multivariate preselection.

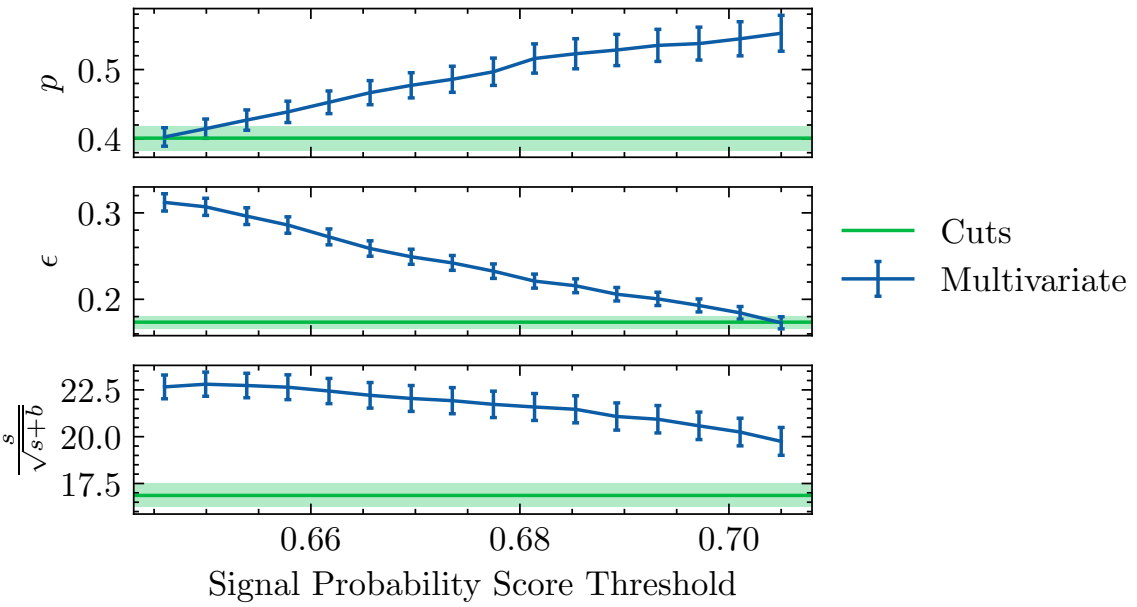

Fig. 6.7 Metrics of selection performance in the signal sample as a function of threshold applied to signal probability score compared with those achieved with the cut-based selection. The signal purity is represented by p, the signal efficiency by ϵ and the significance by $s/\sqrt{s+b}$, where s and b are the number of signal and background events.

This effect is greatest for the Michel electron sample, which is the control sample most contaminated by signal. With the signal selection set for high efficiency, the Michel electron sample is less contaminated with signal and purer with CC events; at thresholds set for high purity, the Michel electron sample aligns with that of the cut-based analysis. Likely at loose threshold values, events with an increasing amount of Michel electrons reconstructed are able to be identified as signal. This would undoubtedly bring about a weaker power to discriminate between signal and CC, therefore it is unsurprising that CC events are also being transferred from the Michel electron sample to the signal sample. The phase space shared by the SMRD selection and the cut-based signal, if it were not for a dedicated veto applied to the SMRD selection that rejects events selected by the signal sample, are events with activity in the SMRD or no activity in the FGD or P0D-ECal other than the two leading proton tracks. In the multivariate analysis, the composition of the SMRD sample grows with the tighter threshold applied and eventually stabilises; at loose threshold values, the sample tends to that of the cut-based selection. Likely with increasing thresholds applied, the multivariate signal selection is rejective of the ~ 200 events that would have otherwise been also selected by the SMRD, due to tighter restrictions made on the subdetector multiplicities. These overlapping events are mostly OOFV neutrons, though there is a small amount of signal too. The P0D control sample, containing mostly OOFV neutrons, is effectively identical in the cut-based and the multivariate analysis for all thresholds—this is because the multivariate selection is no more rejective of OOFV neutrons than the cut-based selection, as has been shown in the next chapter.

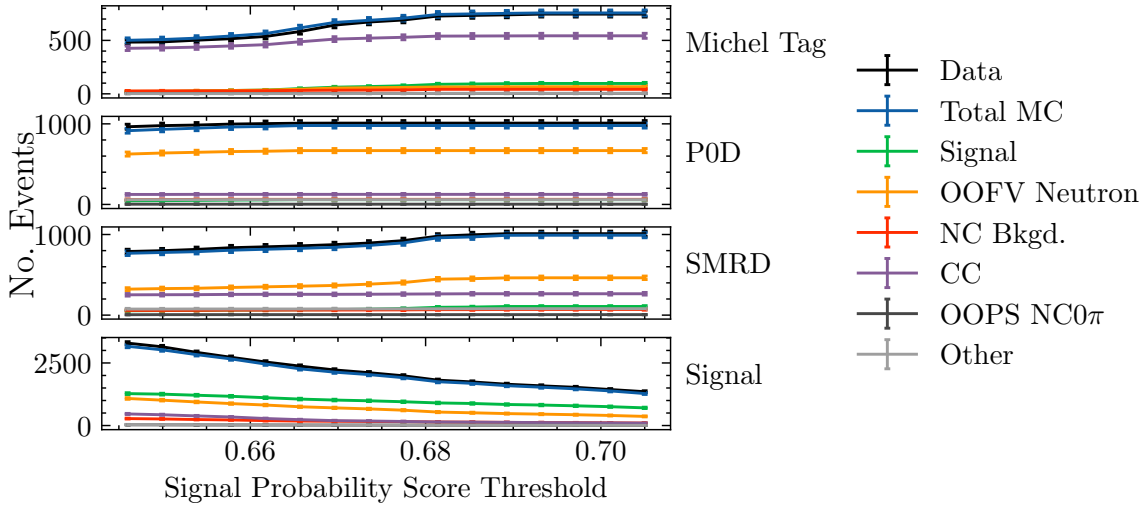

(a) By number of events in sample.

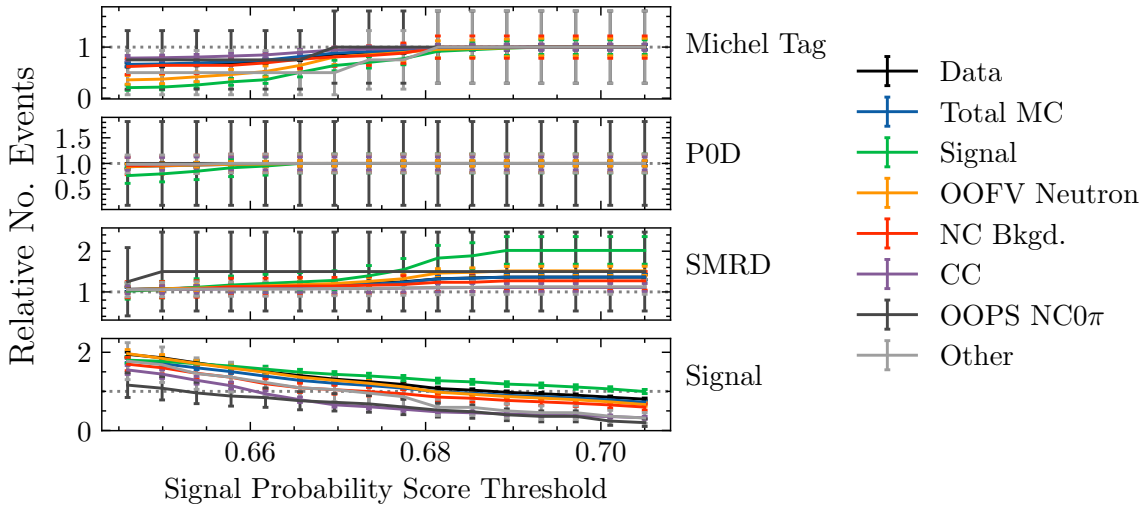

(b) By number of events in sample relative to that of the cut-based selection.

Fig. 6.8 Counts of event classes of each analysis sample a function of threshold applied to signal probability score.

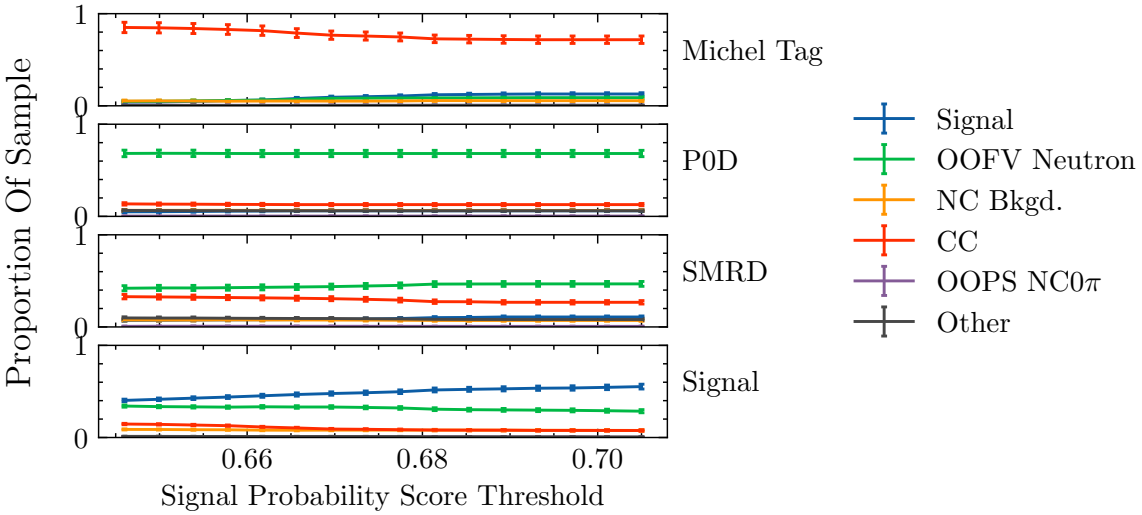

(a) By proportion of events within the total sample.

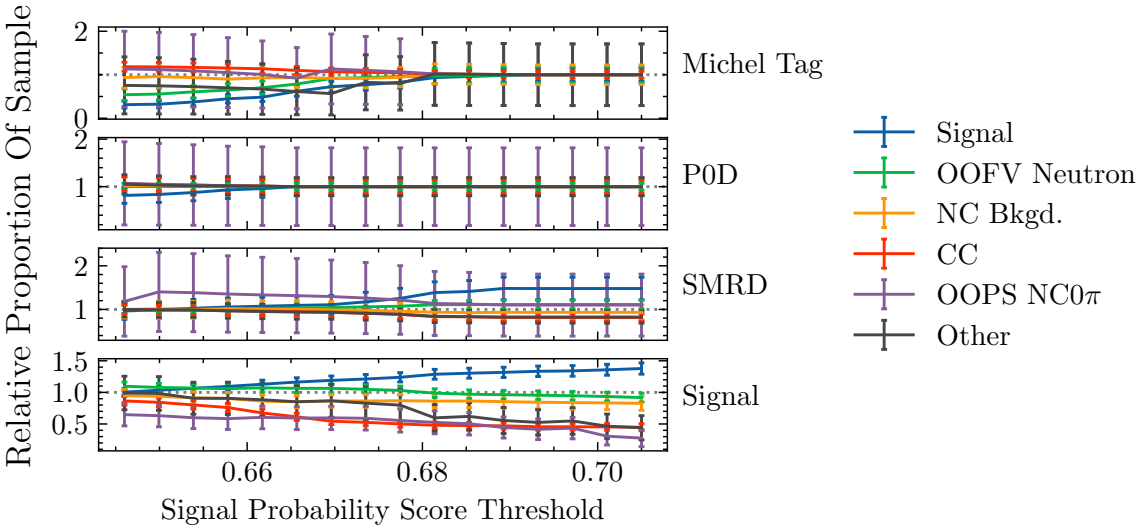

(b) By proportion of events within the total sample relative to that of the cut-based selection.

Fig. 6.9 Proportions of event classes within each analysis sample as a function of threshold applied to signal probability score.

CHAPTER 7

DEVELOPMENT OF MULTIVARIATE SIGNAL EVENT SELECTION

7.1 Introduction

Statistical models that undergo learning draw information from a sample dataset to inform predictions of unseen data, and automatically improve their predictions when supplied with new data. Supervised learning algorithms are to be contrasted with unsupervised learning algorithms. In the former, the model is led to target particular knowns supplied by the analyst; in the latter case, patterns are found in the training data of the model's own accord. Supervised learners may be further classed as either regressors, for prediction of continuous quantities, or classifiers, for prediction of discrete ones. Also to be contrasted somewhat are deep learning algorithms from classical machine learning algorithms (i.e. support vector machines, clustering, decision trees etc.). Their difference lies chiefly in that deep-learning algorithms deploy several layers to identify relatively high-level features in data.

In particle physics of late, there is a growing interest in the power of classification algorithms to perform event selections. An event selection is a pure example of the classical supervised classification scenario, which reduces further to that of binary classification when simply a discriminator between signal and background is desired. Classically, physics samples are formed by an aggregation of thresholds, or *cuts*, placed on the reconstructed quantities; the values of these cuts are intuited either from the physics principles or from the data itself. Given that these cuts typically consider a single variable at a time, this approach is necessarily coarse—the boundary of the acceptance region defined within the reconstructed phase-space is hyper-rectangular. To human eyes, any non-linearities in the boundaries of the signal region in truth are obscured by the dimensionality of the dataset;

multivariate tools, however, may consider non-linearities between any number of variables—therefore, the signal region can be as fine-grained as desired. A caveat in the deployment of machine-learning tools is that some degree of labour must be undertaken to ensure their generalisation. A properly trained model will be sensitive enough to the training data to have absorbed the general features in the data with some degree of precision, but not so sensitive as to have learnt the statistical noise present in the data. The usage of *ensemble* algorithms—meta-algorithms that combine the predictive power of many learners—offer defences from such problems of *overfitting*.

Loosely speaking, the kinds of algorithms in circulation for event selections belong either to those based on decision trees (examples include DØ [120] and MiniBooNE [121]) or those using neural networks (examples include DUNE [122], IceCube [123], NOvA [124]). For this analysis, we have considered only the former, as these methods are widely recognised as robust, as they do not require heavy fine-tuning, and intuitive, as their decision logic is entirely intelligible. Though there are many machine learning software libraries available, this study exclusively uses scikit-learn [125][1]—an open-source package for the Python programming language.

7.2 Classification Trees

A classification tree is a set of conditional statements structured as a binary tree. The most elementary form of the classification tree is the *stump*—a single conditional statement that discriminates between two classes.

A few algorithms are available for growing decision trees from a labelled training dataset—such as ID3 [126], C4.5 [127] and C5.0—the scikit-learn package uses a version of CART (Classification and Regression Trees) [128] which will be briefly described. The variable and split value of each conditional statement are those that best split the data, having considered all variables and all of their possible split values. Characteristic to CART, the quality of a split is given as the gain in Gini impurity. Given a sample, the Gini impurity G is calculated as

$$G = \sum_i p_i (1 - p_i) \tag{7.1}$$

where i indexes the target classes and p_i is the probability of having randomly picked an event belonging to the class with index i. A sample that contains only events belonging to

[1] Version 0.24.2

one class has an impurity of 0. Given a particular split, the gain in Gini impurity is found by calculating the difference in the impurity of the sample before the split and the weighted sum of the Gini impurities of the two samples produced from the split according to the sample sizes. Conditional statements are recursively created in this manner until either the dataset is exhausted or a stopping condition is met. Some common examples of these conditions include limits on the number of conditions, the depth of the tree, or the minimum number of samples at each node. Another approach, known as *pruning*, is to let the tree grow large enough to exhaust the dataset and remove low-quality splits after the fact.

7.3 Ensemble Methods

Classification trees, by themselves, are particularly sensitive to noise present in the training dataset. This problem of high variance is easily overcome by the usage of ensemble algorithms. The training process and the prediction logic vary by the type of ensemble: some of the most popular ensembles—bagging, boosting and the random forest—will be briefly described.

7.3.1 Bagging

Bagging—or, bootstrap aggregation—creates bootstrap samples (samples drawn with replacement) from the training dataset to train each of its constituent trees. After training, the predicted class probability of a given event is taken as the proportion of the predicted classes calculated by the trees.

While trees in the ensemble may be overfitted, they are uncorrelated with each other as the bootstrap sampling creates independent training samples. Their averaging produces a predictive model with reduced variance. The number of trees, along with their hyper-parameters, are the only hyper-parameters to be tuned for the bagging ensemble.

7.3.2 Random forest

The random forest trains its trees with bootstrap samples of the original training dataset and aggregates their predictions as does bagging; however, a further element of randomisation is introduced: the variables considered for each conditional statement in each tree is a random subset of the variables available in the original training dataset. In addition to the number of trees, and their hyper-parameters, the number of variables in that subset is a hyper-parameter to be tuned.

7.3.3 Boosting

Boosting trains its constituent trees sequentially, adapting and improving the training conditions with each iteration. The boosting effect is promoted by increasing the training weights of events misclassified in the current iteration; this approach motivates the tree of the next iteration to correct for the shortcomings of the previous tree, thus improving the predictive power gradually. Again, the predicted class of a given event is conducted by a majority vote.

There exist a few algorithms which enact different treatments of the boosting process and system of voting, such as XGBoost [129] and LogitBoost [130]; the most popular by far, however, is AdaBoost (Adaptive Boosting) and its SAMME (Stagewise Additive Modelling) and SAMME.R [131] (Stagewise Additive Modelling Real) implementations. SAMME and SAMME.R are both multi-class treatments of adaptive boosting but differ in their outputs: SAMME gives the label of the predicted class whereas SAMME.R gives the class probabilities; the latter is also generally favoured to converge with fewer boosting iterations. AdaBoost modifies the weights of training events at every iteration according to

$$w_i' = w_i e^{\alpha \mathbb{1}(y_i = \hat{y}_i)} \tag{7.2}$$

where w_i is the weight of event with index i, $\mathbb{1}$ is the indicator function which evaluates to unity when the predicted class of event i, y_i, matches the class in truth, $\hat{y}_i$, and is zero otherwise, α is the tree's *amount of say* and is calculated by

$$\alpha_m = \beta \ln\left(\frac{\omega_m}{1-\omega_m}\right) \tag{7.3}$$

where β is a hyper-parameter to be tuned known as the *boosting rate*, and ω_m is the previous tree's training accuracy score calculated as

$$\omega_m = \frac{\sum_i w_i \mathbb{1}\left(y_i = \hat{y}_i\right)}{\sum_i w_i}. \tag{7.4}$$

The predicted class is a majority vote of all learners in the ensemble, weighted by the amount of say possessed by each tree.

	Training	Testing
No. signal events	7642 (1.6%)	3275 (1.6%)
No. background events	481474 (98.4%)	206347 (98.4%)
Sum of signal weights	52728 (50%)	652 (1.4%)
Sum of background weights	52728 (50%)	44694 (98.6%)

Table 7.1 Signal and background counts in the training and testing datasets. The bracketed quantities are the proportions of the left quantity within the total sample.

7.4 Training and Testing Dataset

The training and testing datasets are drawn from the MC sample used in the cross section analysis in Chapter 8. Events selected for the training and testing datasets are of good quality and have at least one reconstructed proton; further cuts are also applied to most training variables that remove events with extreme values, as these events are recognised to impede the performance of the model. The complete preselection is given in Table 7.2. Of the events that pass the preselection 70% are assigned as training data; the remaining 30% are assigned as testing data. The target class of events to be identified are true NC0π events with any number of final state protons, with truth-level kinematics within the phase space given by the acceptance of the detector, whereby the highest momentum proton has a true kinetic energy of $>$125 MeV and a true polar angle of $\cos\theta > 0.4$; the remainder of events are classed as background. The statistics of the training and testing datasets are presented in Table 7.1, where it can be seen that signal is overwhelmingly in the minority. To improve signal identification, the sum of signal and background weights are balanced in the training sample by increasing the weights of signal events. The number of signal and background events are left unbalanced for the visualisations and scorings so that results from the testing and training dataset may be compared.

Table 7.2 List of cuts applied as a pre-selection to the multivariate training and test samples.

Variable	Cut
Reconstructed no. proton tracks	> 0
TPC quality of primary proton track	$= 1$
Reconstructed TPC proton particle identification pull of primary proton track	-1×10^4 to 1×10^4
Reconstructed kinetic energy of primary proton track	$< 1 \times 10^4$ MeV
Reconstructed TPC muon particle identification pull of primary proton track	-500 to 500
Muon particle identification pull of most muon-like global track	-1×10^4 to 1×10^4
X-axis component of reconstructed primary proton track end position	-2.5×10^3 to 2.5×10^3 cm
X-axis component of reconstructed secondary proton track end position	-2.5×10^3 to 2.5×10^3 cm
Y-axis component of reconstructed primary proton track end position	-2.5×10^3 to 2.5×10^3 cm
Y-axis component of reconstructed secondary proton track end position	-2.5×10^3 to 2.5×10^3 cm
Z-axis component of reconstructed primary proton track end position	-2×10^3 to 5×10^3 cm
Z-axis component of reconstructed secondary proton track end position	-2×10^3 to 5×10^3 cm
Reconstructed momentum of secondary proton track	$< 1 \times 10^4$ MeV
X-axis component of reconstructed primary proton track momentum	-2.5×10^3 to 2.5×10^3 MeV
X-axis component of reconstructed secondary proton track momentum	-2.5×10^3 to 2.5×10^3 MeV
Y-axis component of reconstructed primary proton track momentum	-2.5×10^3 to 2.5×10^3 MeV
Y-axis component of reconstructed secondary proton track momentum	-2.5×10^3 to 2.5×10^3 MeV
Z-axis component of reconstructed primary proton track momentum	$> -1.5 \times 10^3$ MeV
Z-axis component of reconstructed secondary proton track momentum	$> -1.5 \times 10^3$ MeV
Reconstructed TPC electron particle identification pull of primary proton track	$> -1 \times 10^2$

Continued on next page

Variable	Cut
Reconstructed TPC electron particle identification pull of secondary proton track	$> -1 \times 10^2$
Reconstructed TPC muon particle identification pull of secondary proton track	-1×10^2 to 1×10^2
Reconstructed TPC proton particle identification pull of secondary proton track	-1×10^2 to 1×10^2

7.5 Hyper-Parameter Tuning

The best performing model out of bagging, boosting and the random forest was selected. Accordingly, the hyper-parameters belonging to each model were optimised in order to compare their performances. In each case, the hyper-parameters were optimised by performing the training many times, each time with a randomised set of hyper-parameters, and selecting the set which maximises the classification performance.

After training, these models assign to each event in the testing dataset a probability score ranging from zero to one for each target class—in which a score of zero is background-like and a score of one is signal-like. The sample of selected events is yielded by applying an arbitrary threshold to these probability scores, where the threshold value is to be optimised according to some metric. Ultimately, the threshold value is to be optimised for minimisation of the fractional uncertainty on the extracted cross section; however, how the cross section measurement responds to modifying the composition of the signal sample was not known a priori. Therefore, at this stage in the analysis, the classifier's broad classification ability should be assessed and maximised, in such a way that takes all possible threshold values into consideration. In high energy physics, the most common metrics to evaluate selection performance are signal *purity*[2], which is defined as

$$p = \frac{n(y \cap \hat{y})}{n(y)} \tag{7.5}$$

and *efficiency*[3], which is defined as

$$\varepsilon = \frac{n(y \cap \hat{y})}{n(\hat{y})} \tag{7.6}$$

where n is the number of events, y is the set of selected signal events and $\hat{y}$ is the set of true signal events. The scores received by both of these metrics must both be reasonable for any viable physics analysis, therefore both the purity and efficiency can be integrated into a single value by taking $p \times \varepsilon$[4]. There is necessarily a trade-off between selection purity and efficiency according to the threshold value applied—at low threshold values, the efficiency is high but the purity is vanishing, and vice versa. A suitable metric to measure a model's purity and efficiency over the full spectrum of possible threshold values is the average purity. The average purity is calculated as

[2] Otherwise known as *precision* outside of the particle physics community.

[3] Otherwise known as *recall* outside of the particle physics community.

[4] The maximisation of $p \times \varepsilon$ is equivalent to minimisation of the relative statistical error on the amount of signal in the sample.

$$\bar{p} = \sum_i (\varepsilon_i - \varepsilon_{i-1})\, p_i \tag{7.7}$$

where i indexes the set of possible score threshold values, p_i is the purity achieved with the i^{th} threshold and ε_i is the efficiency achieved with the i^{th} threshold. Another viable metric is certainly the area under the purity–efficiency curve; however, this tends to overestimate the classification performance particularly in regions of low efficiency[5]. The average purity is favoured as a more conservative estimate of the model performance.

The three candidate models were optimised for their hyper-parameters based on the cross-validated average purity achieved on the test dataset. Cross-validation is an alternative to the train–validation–test dataset paradigm used to evaluate the degree of under or overfitting during the hyper-parameter optimisation process. Instead of training and validating the model once, the process is performed a number of times with resampled versions of the training dataset. The *K-fold* approach to cross-validation in particular generates k training datasets of size $(k-1)/k$ and k hold-out datasets of size $1/k$. The data is partitioned randomly into k subsets; for all subsets, each subset is taken as validation data and the aggregate of all remaining subsets is used as the associated training data. A *stratified* K-fold approach ensures that the amount of signal events is consistent in each partition. Five-fold stratified cross-validation was used to assess the model performances.

The variables included in the training and test datasets are the 13 of the cut-based signal selection listed in Table 7.3 and reasoned later in Section 6.3.2; loosely speaking, they describe properties of the primary proton candidate and the corresponding amount of activity in the subdetectors. Given these, the average purity scores are shown as a function of their hyper-parameters for the boosting, bagging and random forest models in Figs. 7.1 to 7.3.

[5] In regions of low efficiency, the vanishing sample size provokes wild fluctuations in the sample purity. Given also that there are necessarily fewer threshold values available in this region, calculating the area under the curve via the trapezoid method—which interpolates linearly between points—becomes problematic.

Table 7.3 Univariate selection variables.

Alias	Description	Unit
fgd1_num_michel_tag	Reconstructed no. Michel electrons	—
global_most_muon_like_pull	Muon PID of most muon-like global track	—
num_other_with_ecal	Reconstructed no. other ECal tracks	—
num_other_with_fgd	Reconstructed no. other FGD1 tracks	—
num_other_with_p0d	Reconstructed no. other P0D tracks	—
num_other_with_p0decal	Reconstructed no. other P0D-ECal tracks	—
num_other_with_tpc_negative_tracks	Reconstructed no. other negative TPC tracks	—
num_other_with_tpc_positive_tracks	Reconstructed no. other positive TPC tracks	—
p0charge	Reconstructed charge of primary proton track	—
p0ek	Reconstructed kinetic energy of primary proton track	MeV
p0pos.Z()	Z-axis component of reconstructed primary proton track front position	cm
p0tpcpullmuon	Reconstructed TPC muon particle identification pull of primary proton track	—
p0tpcpullproton	Reconstructed TPC proton particle identification pull of primary proton track	—

In the boosting algorithm, a low boosting rate, high number of leaf nodes per tree and high number of boosting iterations appears to be the ideal set of parameters. In Fig. 7.1b[6], a preference for a high number of boosting iterations can be seen when the boosting rate is kept low (< 0.1). In Fig. 7.1c, a preference for a high number of leaves can be seen when the boosting rate is kept low. In Fig. 7.1a[7], the relationship between the maximum number of leaves per tree and number of boosting iterations is relatively weak, so it can be concluded that a low boosting rate is essential. It is also somewhat worth noting that, since boosting is an iterative training process and therefore cannot be parallelised, the fit times at high model complexities could become so large as to become impractical. If the hyper-parameter search were to be conducted again, the range of maximum number of leaves should be extended as the optimised model is at the upper limit of the sampling. Also, the boosting rate would be better off sampled with a logarithmic probability density function rather than a uniform one, given that low values are favoured. In the random forest algorithm (Fig. 7.3), convergence to the optimised value can be achieved with a lower model complexity than the bagging algorithm. In both the bagging and random forest algorithms (Figs. 7.2 and 7.3), the performance is dependent primarily on the maximum number of leaves per tree given that the number of trees is not extremely low.

Given these hyper-parameters, the performances can certainly be improved further by introducing further signal events into the training dataset, as these are overwhelmingly in the minority. For all models, the fractional standard deviation of the scores calculated for each training dataset fold is small—therefore the model variance is deemed low. Also for all models, the average purity scores are moderately consistent—therefore the model's bias to their training dataset is also deemed low. For the cross section analysis performed on blinded data, some amount of bias is tolerable as the training and test datasets are partitions of the deployment dataset; for the unblinded analysis, achieving a low training bias is more important as it will show up as a discrepancy in data and MC rates since the deployment dataset—real detector data—is unseen during the training. Given these, it can be concluded that the optimised models are neither significantly under-fit nor over-fit.

The maximum scores achieved for each model are given in Table 7.4; the set of hyper-parameters corresponding to each of these scores are given in Table 7.5. It is seen from these scores and Figs. 7.4 and 7.5 that the boosting algorithm is the highest performing.

[6]The hyper-parameter sampling for the AdaBoost model was conducted in two stages, hence the distribution of data points observed in Fig. 7.1b. After the first stage, a preference for a high number of boosting iterations was observed; therefore, a second stage was run with a reduced range of boosting rates and a higher number of boosting iterations. In retrospect, this was not ideal for the visualisations.

[7]As per [6], it should be acknowledged in this plot that data points with the number of boosting iterations exceeding 200 have a reduced range of boosting rates.

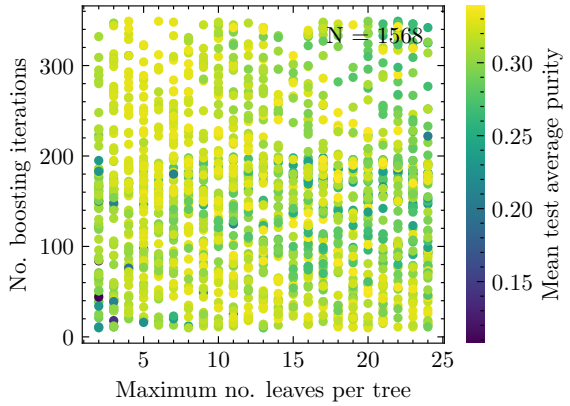

(a) Number of boosting iterations against maximum number of leaves in trees.

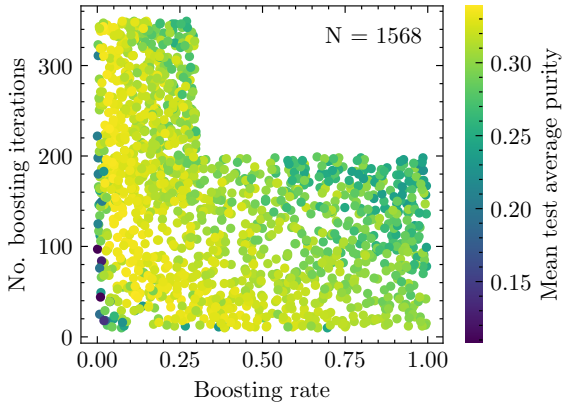

(b) Number of boosting iterations against learning rate.

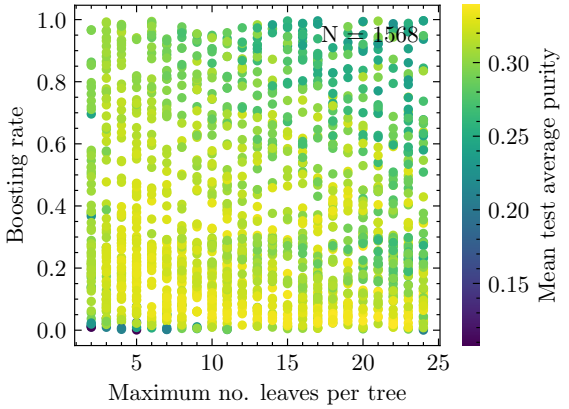

(c) Maximum number of leaves in tree against learning rate.

Fig. 7.1 Average purity on test dataset as a function of hyper-parameters relevant to the AdaBoost ensemble algorithm. The quantity N corresponds to the number of data points.

The bagging ensemble is the worst performing and is capable of no better purity than the cut-based selection. The performances of the random forest and bagging ensembles are competitive in the high-efficiency region, though the random forest has some superiority in terms of its peak purity. The boosting algorithm, on the other hand, achieves superior purity and efficiency scores for all threshold values—achieving a $\sim$5% increase in purity relative to the cut-based selection for an equivalent efficiency score.

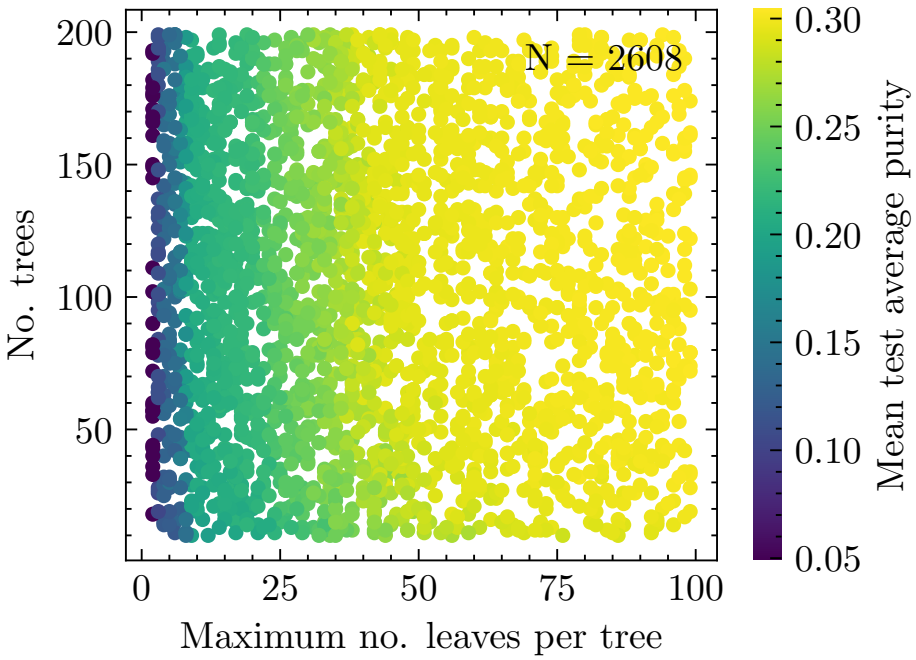

Fig. 7.2 Average purity on test dataset as a function of number of trees and their sizes within Bagging ensemble algorithm. The quantity N corresponds to the number of data points.

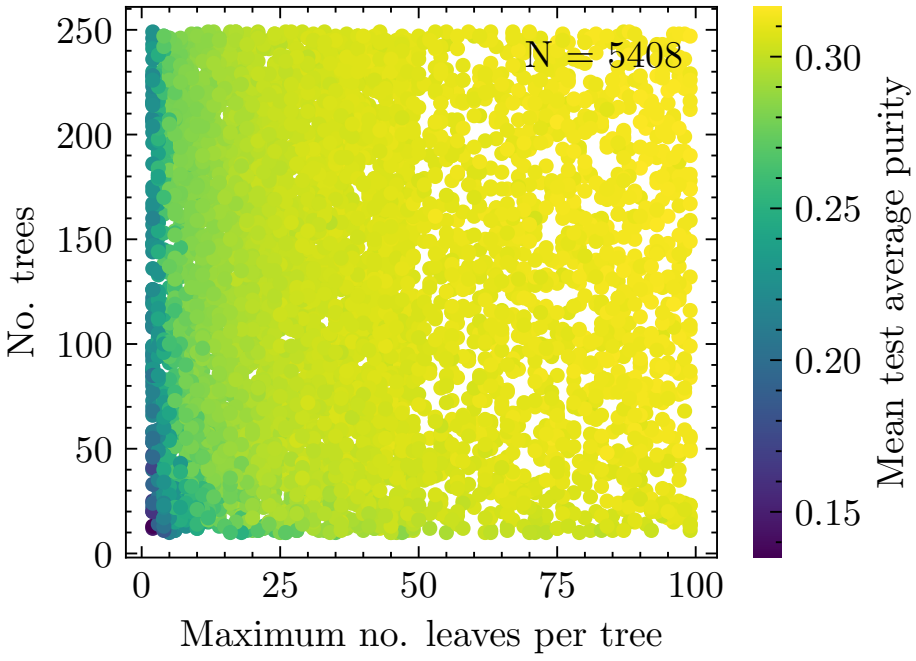

Fig. 7.3 Average purity on test dataset as a function of number of trees and their sizes within Random Forest ensemble algorithm. The quantity N corresponds to the number of data points.

	AdaBoost	Bagging	Random Forest
Training score	0.354 ± 0.003	0.312 ± 0.003	0.321 ± 0.004
Testing score	0.339 ± 0.013	0.304 ± 0.010	0.316 ± 0.014

Table 7.4 Average purity scores on the test and training datasets for the AdaBoost, Bagging and Random Forest ensembles. The value and error given is the mean and standard deviation calculated across the scores of each cross-validation fold.

AdaBoost		Bagging		Random Forest	
No. boosting iterations	173	No. trees	122	No. trees	238
Boosting rate	0.026	Maximum no. leaves per tree	99	Maximum no. leaves per tree	89
Maximum no. leaves per tree	24				

Table 7.5 Optimised parameters for the AdaBoost, Bagging and Random Forest ensembles.

7.6 Selection of Training Variables

At this stage in the study, the inferior bagging and random forest models are abandoned and the training variables for the AdaBoost model are optimised. Of the 46 reconstructed variables available in the dataset—which are listed in Table 7.6—only a select few are helpful for telling signal from background; the majority are either redundant or carry little to no of such information and hence may be discarded with a negligible penalty to the selection performance. Reducing the dimensionality of the training dataset reveals the most important variables and reduces computation times. An algorithm was developed to remove redundant or unimportant variables from the dataset.

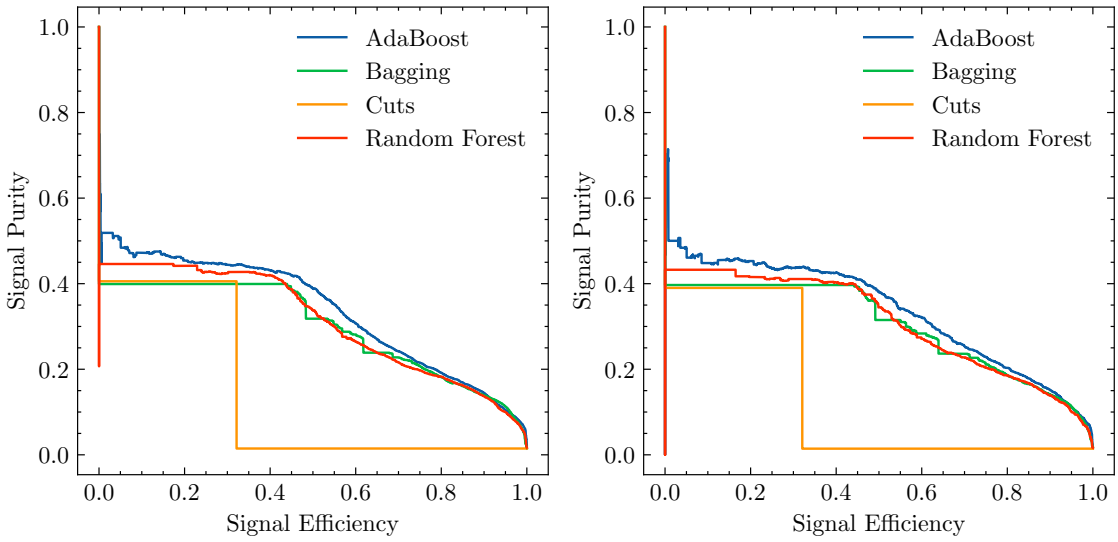

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.4 Purity–efficiency curves for the optimised multivariate models and the cut-based signal selection as calculated on the training dataset (Fig. 7.4a) and the testing datasets (Fig. 7.4b). The training variables for the multivariate models are the cut-based variables.

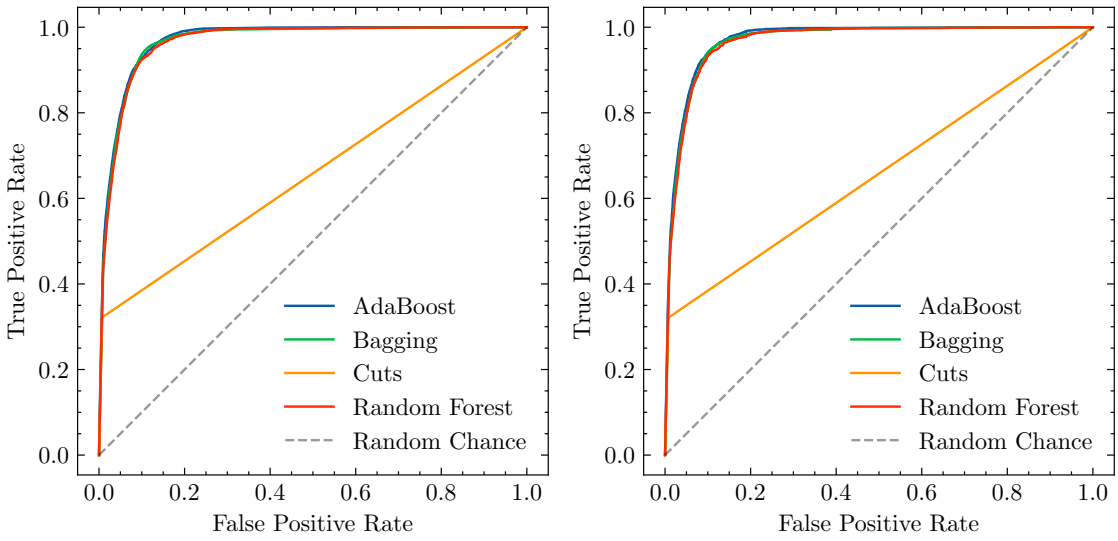

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.5 Receiver Operating Characteristic (ROC) curves for the optimised multivariate models and the cut-based signal selection as calculated on the training dataset (Fig. 7.5a) and the testing datasets (Fig. 7.5b). The training variables for the multivariate models are the cut-based variables.

Table 7.6 List of variables selected from multivariate training and test samples.

Alias	Description	Unit
fgd1_num_michel_tag	Reconstructed no. Michel electrons	—
global_most_muon_like_pull	Muon PID of most muon-like global track	—
num_other_with_ecal	Reconstructed no. tracks other than the primary proton track with an ECal component	—
num_other_with_fgd	Reconstructed no. tracks other than the primary proton track with an FGD1 component	—
num_other_with_p0d	Reconstructed no. tracks other than the primary proton track with an P0D component	—
num_other_with_p0decal	Reconstructed no. tracks other than the primary proton track with an P0D-ECal component	—
num_other_with_tpc_negative_tracks	Reconstructed no. tracks other than the primary proton track with an negatively charged TPC component	—
num_other_with_tpc_positive_tracks	Reconstructed no. tracks other than the primary proton track with an positively charged TPC component	—
num_other_with_upstream_smrd_tracks	Reconstructed no. tracks other than the primary proton track with an upstream SMRD component	—
numproton	Reconstructed no. proton tracks	—
p0charge	Reconstructed charge of primary proton track	—
p0ek	Reconstructed kinetic energy of primary proton track	MeV

Continued on next page

Alias	Description	Unit
p0endpos.X()	X-axis component of reconstructed primary proton track end position	cm
p0endpos.Y()	Y-axis component of reconstructed primary proton track end position	cm
p0endpos.Z()	Z-axis component of reconstructed primary proton track end position	cm
p0mom	Reconstructed momentum of primary proton track	MeV
p0mom3vec.X()	X-axis component of reconstructed primary proton track momentum	MeV
p0mom3vec.Y()	Y-axis component of reconstructed primary proton track momentum	MeV
p0mom3vec.Z()	Z-axis component of reconstructed primary proton track momentum	MeV
p0pos.X()	X-axis component of reconstructed primary proton track front position	cm
p0pos.Y()	Y-axis component of reconstructed primary proton track front position	cm
p0pos.Z()	Z-axis component of reconstructed primary proton track front position	cm
p0theta	Reconstructed polar angle of primary proton track	π radian
p0tpccharge	Reconstructed TPC charge of primary proton track	—
p0tpcnhits	Reconstructed no. TPC hits of primary proton track	—
p0tpcpullelectron	Reconstructed TPC electron particle identification pull of primary proton track	—

Continued on next page

Alias	Description	Unit
p0tpcpullmuon	Reconstructed TPC muon particle identification pull of primary proton track	—
p0tpcpullproton	Reconstructed TPC proton particle identification pull of primary proton track	—
p1charge	Reconstructed charge of secondary proton track	—
p1endpos.X()	X-axis component of reconstructed secondary proton track end position	cm
p1endpos.Y()	Y-axis component of reconstructed secondary proton track end position	cm
p1endpos.Z()	Z-axis component of reconstructed secondary proton track end position	cm
p1mom	Reconstructed momentum of secondary proton track	MeV
p1mom3vec.X()	X-axis component of reconstructed secondary proton track momentum	MeV
p1mom3vec.Y()	Y-axis component of reconstructed secondary proton track momentum	MeV
p1mom3vec.Z()	Z-axis component of reconstructed secondary proton track momentum	MeV
p1pos.X()	X-axis component of reconstructed secondary proton track front position	cm
p1pos.Y()	Y-axis component of reconstructed secondary proton track front position	cm

Continued on next page

Alias	Description	Unit
p1pos.Z()	Z-axis component of reconstructed secondary proton track front position	cm
p1theta	Reconstructed polar angle of secondary proton track	π radian
p1tpccharge	Reconstructed TPC charge of secondary proton track	—
p1tpcnhits	Reconstructed no. TPC hits of secondary proton track	—
p1tpcpullelectron	Reconstructed TPC electron particle identification pull of secondary proton track	—
p1tpcpullmuon	Reconstructed TPC muon particle identification pull of secondary proton track	—
p1tpcpullproton	Reconstructed TPC proton particle identification pull of secondary proton track	—
p1tpcquality	TPC quality of secondary proton track	—

The degree to which each variable in the dataset influences the classification performance may be gauged by the permutation method [132]. In this method, the importance of a particular variable is quantified by the decrease in score, given a particular metric, calculated when the variable's relationship to the target class is broken—this is achieved by randomly shuffling the data of the variable of interest. The importances calculated by this method on all reconstructed variables are shown in Fig. 7.6. This method will evidently overestimate the importance of variables that are significantly correlated with other variables in the dataset and hence provide redundant information. With this flaw considered, the dimensionality reduction process occurs in two stages: firstly, the dataset is rid of redundancies by collecting correlated variables in clusters and selecting only one from each, and secondly the unimportant variables are identified and removed. The optimum number of clusters formed in the first stage to pass to the second stage cannot be known a priori, therefore the second stage is run for every possible way of clustering. The variables selected were taken as the variable list yielded by the algorithm that minimises both their number and the decrease in performance beyond a particular threshold.

Variables that express similar information in the dataset are collected together by performing a cluster analysis on the correlation matrix. The correlation matrix for events passing the preselection is given in Fig. 7.7. A hierarchy is constructed according to the strength of correlations between variables in each cluster: clusters merged low in the hierarchy are likely to contain a number of similar (and therefore redundant) variables, whereas clusters merged high in the hierarchy contain more dissimilar variables. In the agglomerative approach, the hierarchy is constructed from the bottom-up: all variables initially occupy their own cluster and are merged one by one until all variables are contained within a single cluster. The pair of clusters to be merged next in the hierarchy is decided by the *linkage* between pairs of clusters, which provides a measure of their similarity; the linkage is, in turn, a function of the *distance metric* calculated between each pair of variables across the pair of clusters. Ward's method [133], which merges the pair of clusters that minimises the variance of the distance metric contained within the merged cluster, is used as the linkage function. The distance metric used is the Euclidean distance. The hierarchy created from this cluster analysis is shown in Fig. 7.8. At the bottom of the hierarchy, it can be seen that derived quantities are merged first as expected. For example, proton kinetic energy is derived from its momenta measurement hence these are the first merged variables; also, the TPC data quality flag is derived from the number of TPC hits hence these are merged second. By applying an arbitrary threshold to the distances calculated between each cluster, or the *clustering strength*, yielded is a number of groups of variables that are inter-correlated to a degree according to the strength value.

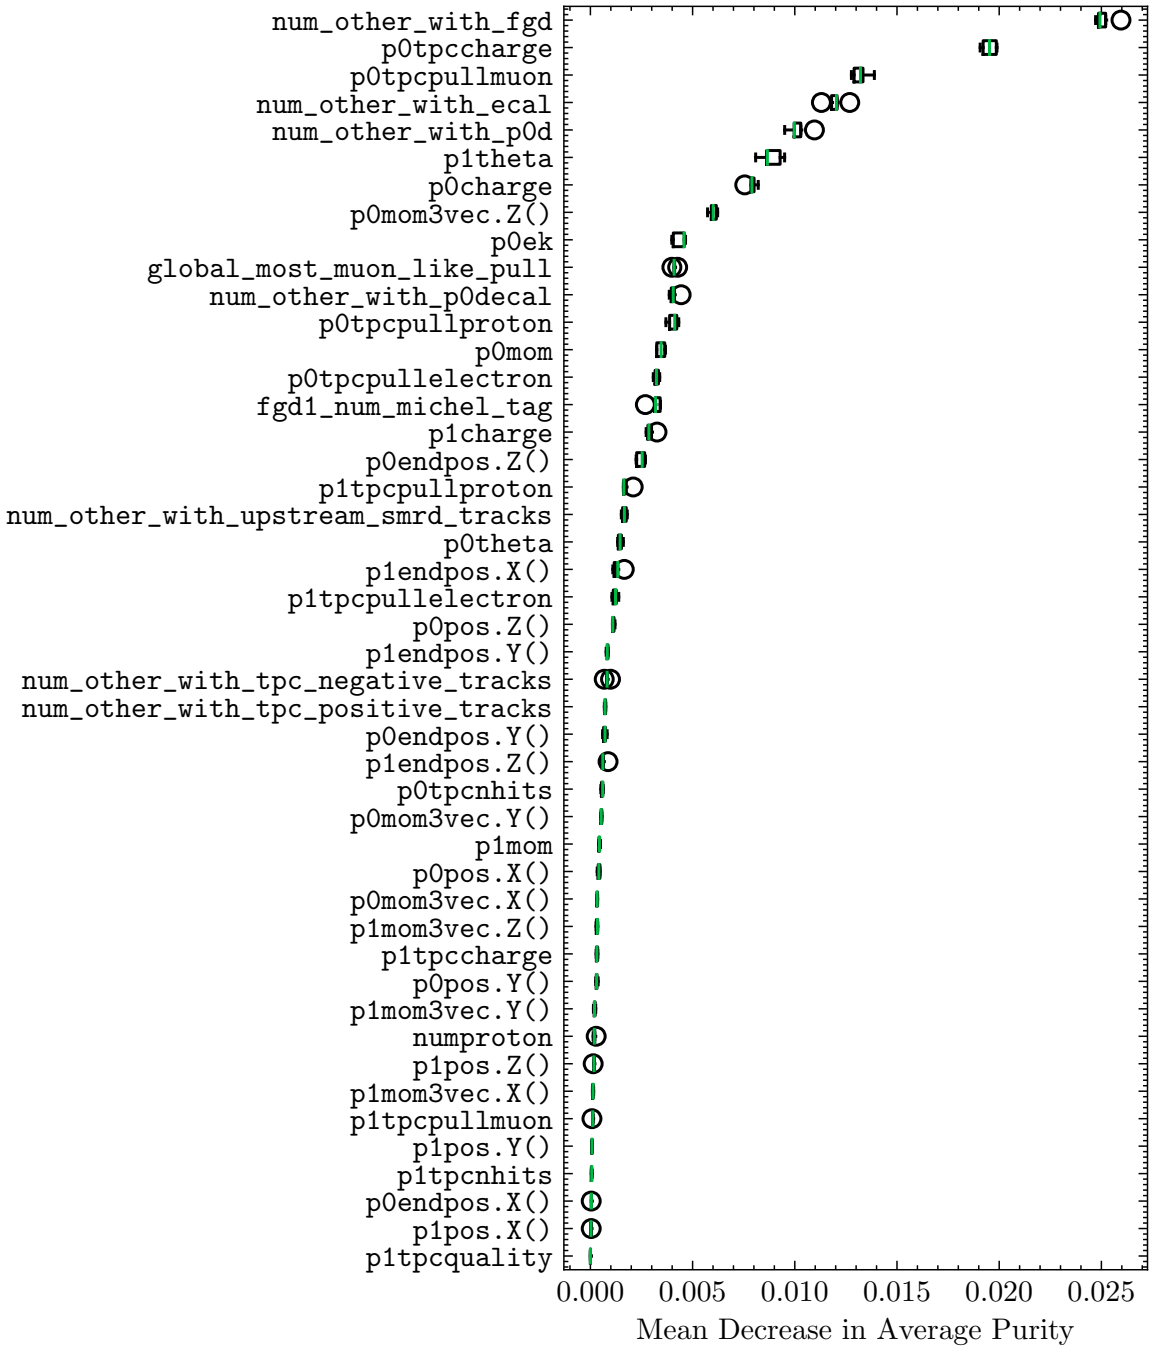

Fig. 7.6 Relative importances of the 46 variables as calculated via the permutation method on the training dataset. Descriptions of these variables are given in Table 7.6. The importances are calculated five times for each variable, these results are represented by a box-and-whisker.

Fig. 7.7 Matrix of Spearman's rank pairwise correlations on the training dataset including all 46 reconstructed variables. Descriptions of these variables are given in Table 7.6.

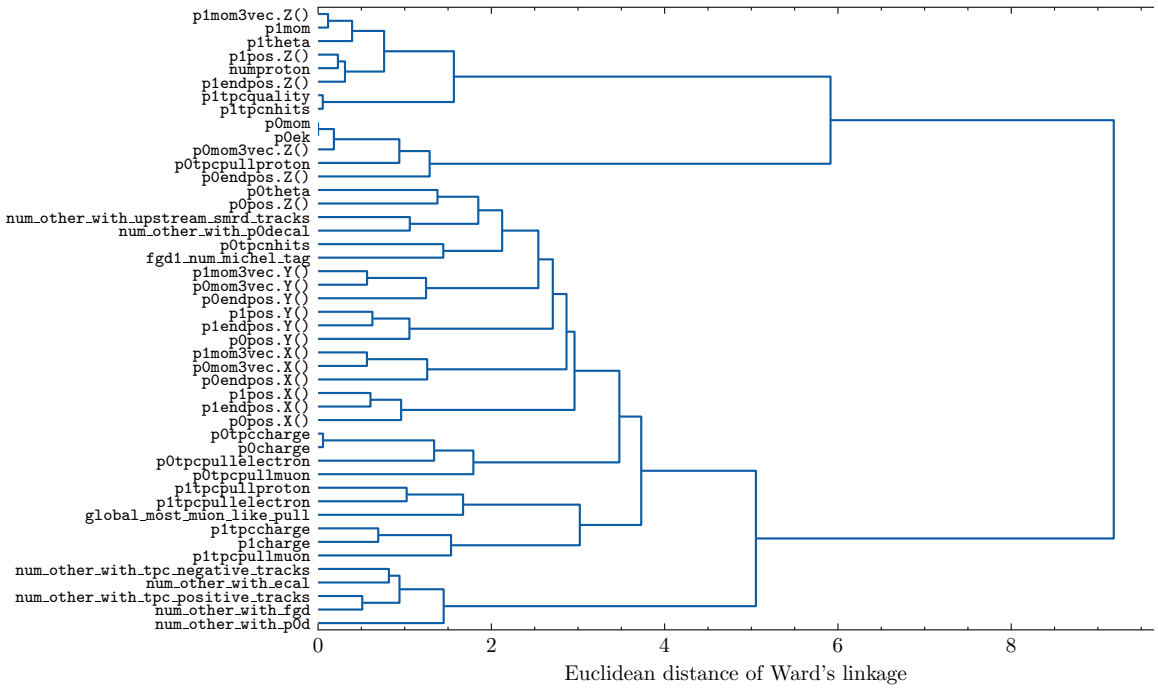

Fig. 7.8 Dendrogram constructed of the Euclidean distances achieved with Ward's linkage metric as calculated on the training dataset that includes all 46 reconstructed variables. Descriptions of the variables are given in Table 7.6.

These variable groups contain redundant information, therefore a single variable is then chosen to represent the information contained within each group found at each possible clustering strength. Accordingly, the hierarchy is solved for the most important variables of the child clusters belonging to each possible cluster. This is approached recursively from the bottom-up, whereby the most important of the two variables presented by the left and right branches of the head node of each cluster is taken as the best of the cluster. The most important variable is taken as the variable of the highest permutation importance as calculated on the testing dataset after the model was trained with the other variable removed; the variable of lesser importance is identified as the redundant one. Having performed this process for the entirety of the hierarchy, the result is a list of training variables with redundancies removed for each possible clustering strength.

In the second stage, each list of training variables yielded by the first stage in the algorithm is reduced by importance. Given a list of training variables yielded by a given clustering strength applied, one by one the least important variable in the list is removed until the average purity calculated on the testing dataset declines beyond a particular threshold. A limit of 0.01 loss in score was placed relative to the average purity achieved by training the model on all reconstructed variables.

No. selected variables	No. clusters	No. eliminated	Decrease in average purity
18	0	28	0.009
17	1	28	0.009
18	2	26	0.008
16	3	27	0.010
16	4	26	0.010
17	5	24	0.006
15	6	25	0.008
16	7	23	0.007
15	8	23	0.008
15	9	22	0.008
15	10	21	0.008
15	11	20	0.008
15	12	19	0.008
15	13	18	0.008
15	14	17	0.008
15	15	16	0.008
15	16	15	0.008
15	17	14	0.008
15	18	13	0.008

Table 7.7 Number of variables selected by the variable reduction algorithm according to the number of clusters formed in the first stage of the algorithm.

The minimised number of variables selected by the reduction algorithm according to the number of clusters formed in the first stage are given in Table 7.7. It can be seen that the algorithm with clustering can remove up to three more variables compared with no clustering. The algorithm selected 15 training variables by forming eight clusters in the first stage, corresponding to an average purity loss of 0.008 compared with training on all reconstructed variables. The selected variables are listed in Table 7.8, where it can be seen that these coincide much with the cut-based variables listed in Table 7.3 but additionally consider the charge and polar angle of the secondary proton candidate.

Table 7.8 List of variables selected for multivariate training and test samples.

Alias	Description	Unit
fgd1_num_michel_tag	Reconstructed no. Michel electrons	—
global_most_muon_like_pull	Muon PID of most muon-like global track	—
num_other_with_ecal	Reconstructed no. other ECal tracks	—
num_other_with_fgd	Reconstructed no. other FGD1 tracks	—
num_other_with_p0d	Reconstructed no. other P0D tracks	—
num_other_with_p0decal	Reconstructed no. other P0D-ECal tracks	—
num_other_with_upstream_smrd_tracks	Reconstructed no. other upstream SMRD tracks	—
p0ek	Reconstructed kinetic energy of primary proton track	MeV
p0endpos.Z()	Z-axis component of reconstructed primary proton track end position	cm
p0theta	Reconstructed polar angle of primary proton track	π radian
p0tpccharge	Reconstructed TPC charge of primary proton track	—
p0tpcpullelectron	Reconstructed TPC electron particle identification pull of primary proton track	—
p0tpcpullmuon	Reconstructed TPC muon particle identification pull of primary proton track	—
p1charge	Reconstructed charge of secondary proton track	—
p1theta	Reconstructed polar angle of secondary proton track	π radian

The purity–efficiency curves for the optimised AdaBoost model trained on the 13 cut-based variables, the 15 selected variables and all 46 reconstructed variables are shown in Fig. 7.9. It can be seen that the performance degradation of the training performed with the selected variables compared to training on all reconstructed variables is minor. Including the charge and angle of the secondary candidate proton in the training can achieve a significant increase in purity—this is due to an enhanced rejection of some backgrounds.

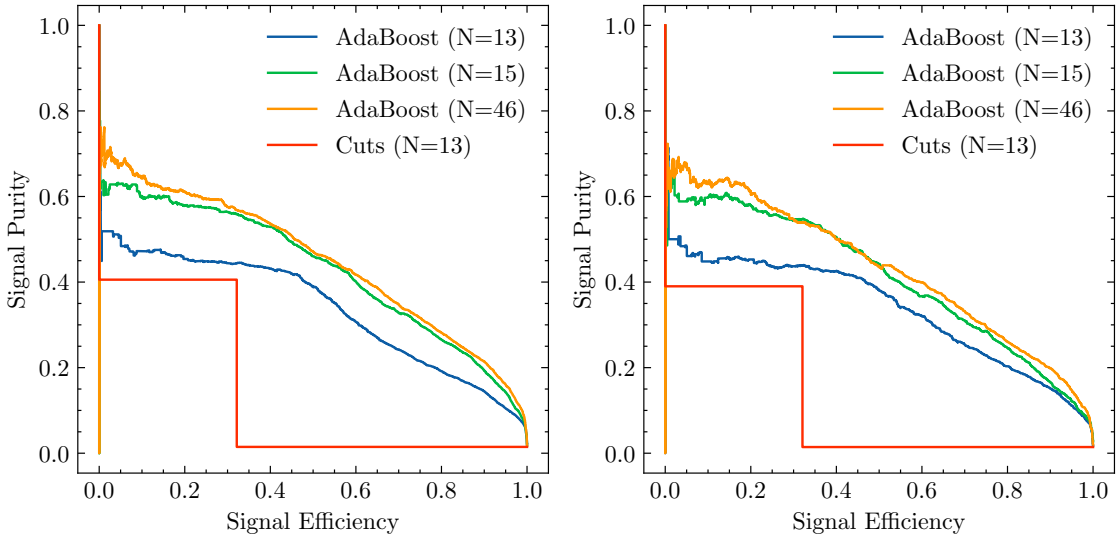

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.9 Purity–efficiency curves for the optimised AdaBoost model varied by number of training variables and the cut-based signal selection as calculated on the training dataset (Fig. 7.9a) and the testing datasets (Fig. 7.9b). The curve with 13 variables is trained with the variables that form the cut-based selection; with 15 variables, the reduced variables; with 46 variables, all reconstructed variables.

The relative importances of these selected variables given in Fig. 7.10 show the polar angle of the secondary particle is among the most important; likely the model is using correlations between the angles reconstructed of the first and second proton to discriminate between signal and background. The two dominant backgrounds in the signal region are interactions from neutrons that originate from outside of the fiducial volume and CC interactions within the fiducial volume.

As seen in Fig. 7.12, a marginal enhancement in rejection power to out-of-fiducial-volume neutrons is observed when secondary particle information was trained upon; however, as shown in Fig. 7.11 a significant increase is seen for CC background. For CC interactions with two candidate reconstructed proton tracks, the second-highest momentum track is most likely a proton, μ^- or a π^+. The increase in rejection power is greatest for CC interactions with

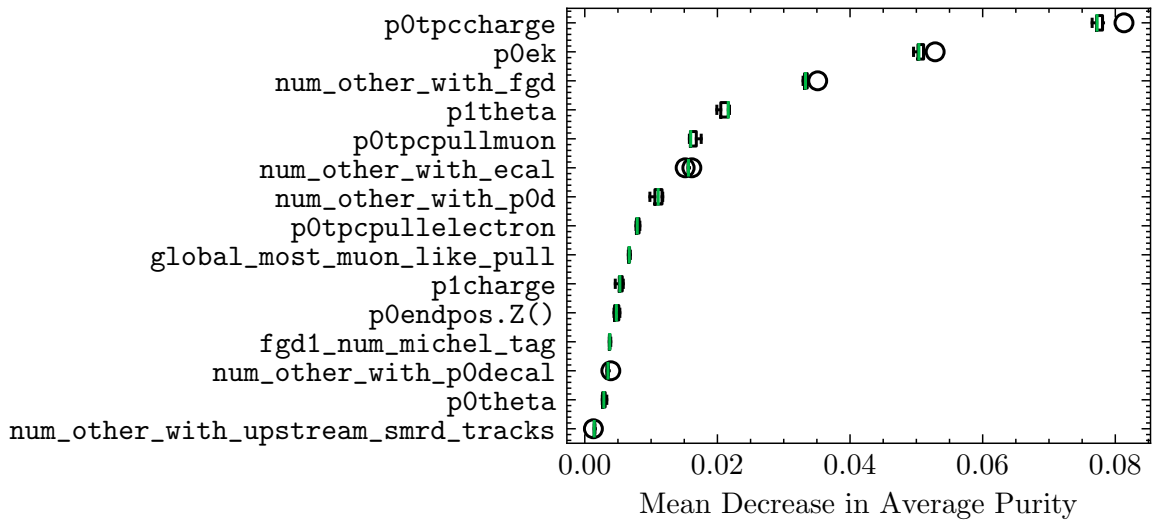

Fig. 7.10 Relative importances of the 15 selected variables as calculated via the permutation method on the training dataset. Descriptions of these variables are given in Table 7.8.

a muon as the secondary particle when the charge and angle of the secondary particle are trained upon; some enhanced rejection is seen for secondary π^+, but little was observed for secondary protons. Also, a subdominant contribution to the background in the signal region is provided by NC interactions with pions or an absence of protons present in the final state. As shown in Fig. 7.13, some increased rejection of these events is seen when secondary particle information is trained upon. But despite knowing about the secondary proton, the model has no special sensitivity to signal events with two protons in the final state—as seen in Fig. 7.14, the efficiency is increased by a constant factor for any number of final state protons relative to the cut-based selection. Because how the model partitions the 15-dimensional phase-space into the designated signal and background is both highly-dimensional and irregular, it is difficult to gain intuition into the logic of the predictor. In Fig. 7.15, the pairwise scatters and distributions for the five most important training variables are shown, split by true positives, true negatives, false negatives and false positives. The one-dimensional histograms reveal a number of sensible acceptance regions are designated by the predictor. For example, that the highest-momentum track of a signal event will have a positive charge in the TPC, that there will be no objects other than the two proton candidates in the FGD when a signal event occurs, or that the primary proton candidate will not be muon-like in the TPC.

The AdaBoost model was trained on the 15 selected variables and classified on the nominal MC and the detector data. A discussion of the event composition of the sample selected by this classifier may be found in Section 6.3.6

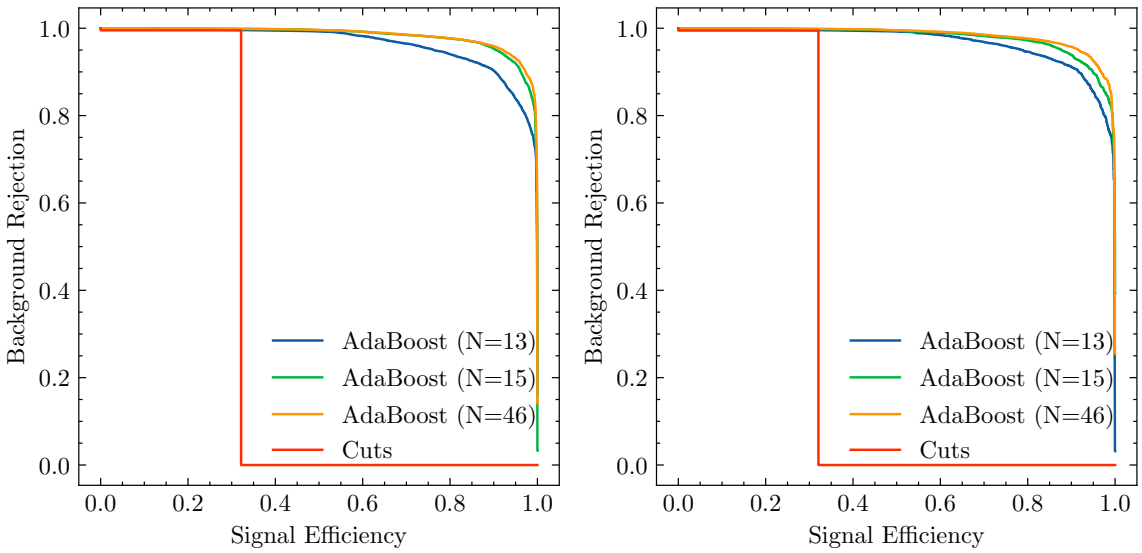

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.11 Charged current background rejection against signal efficiency evaluated on the training (Fig. 7.11a) and testing (Fig. 7.11b) datasets. The curve with 13 variables is trained with the variables that form the cut-based selection; with 15 variables, the reduced variables; with 46 variables, all reconstructed variables.

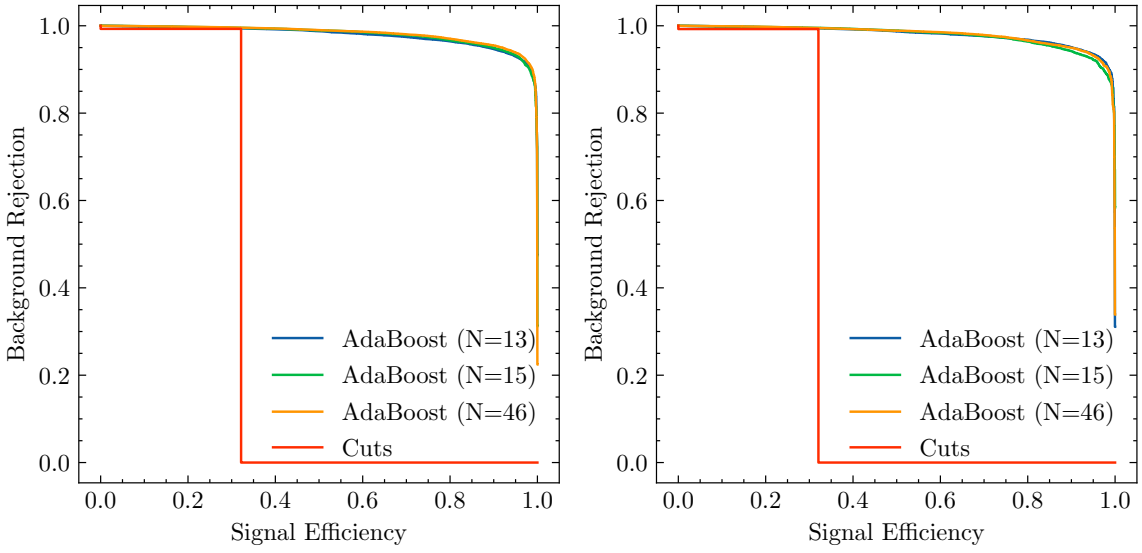

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.12 Out-of-fiducial-volume neutron background rejection against signal efficiency evaluated on the training (Fig. 7.12a) and testing (Fig. 7.12b) datasets. The curve with 13 variables is trained with the variables that form the cut-based selection; with 15 variables, the reduced variables; with 46 variables, all reconstructed variables.

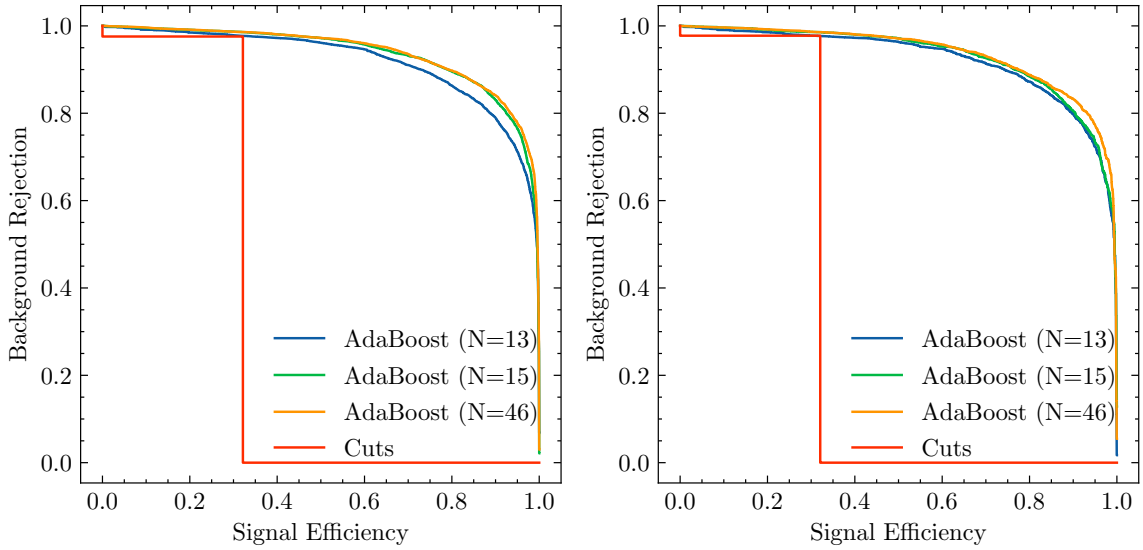

(a) For the training dataset.

(b) For the testing dataset.

Fig. 7.13 Neutral current background rejection against signal efficiency evaluated on the training (Fig. 7.13a) and testing (Fig. 7.13b) datasets. The curve with 13 variables is trained with the variables that form the cut-based selection; with 15 variables, the reduced variables; with 46 variables, all reconstructed variables.

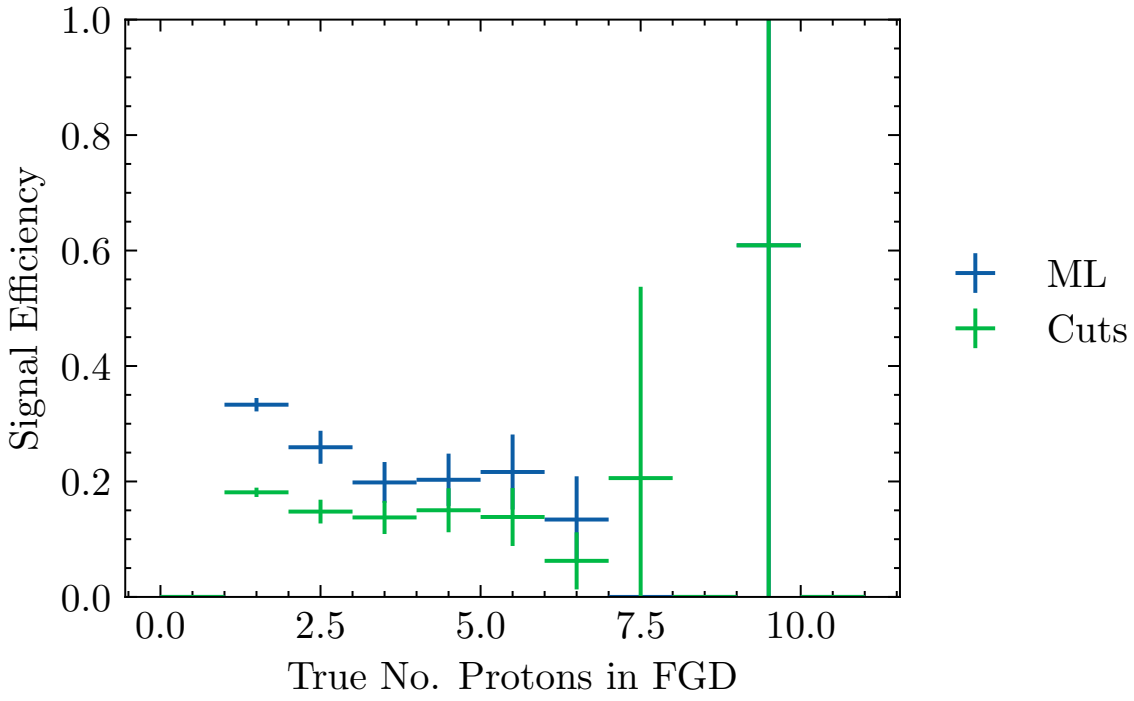

Fig. 7.14 Signal efficiency as a function of true proton multiplicity for the multivariate and cut-based signal selections.

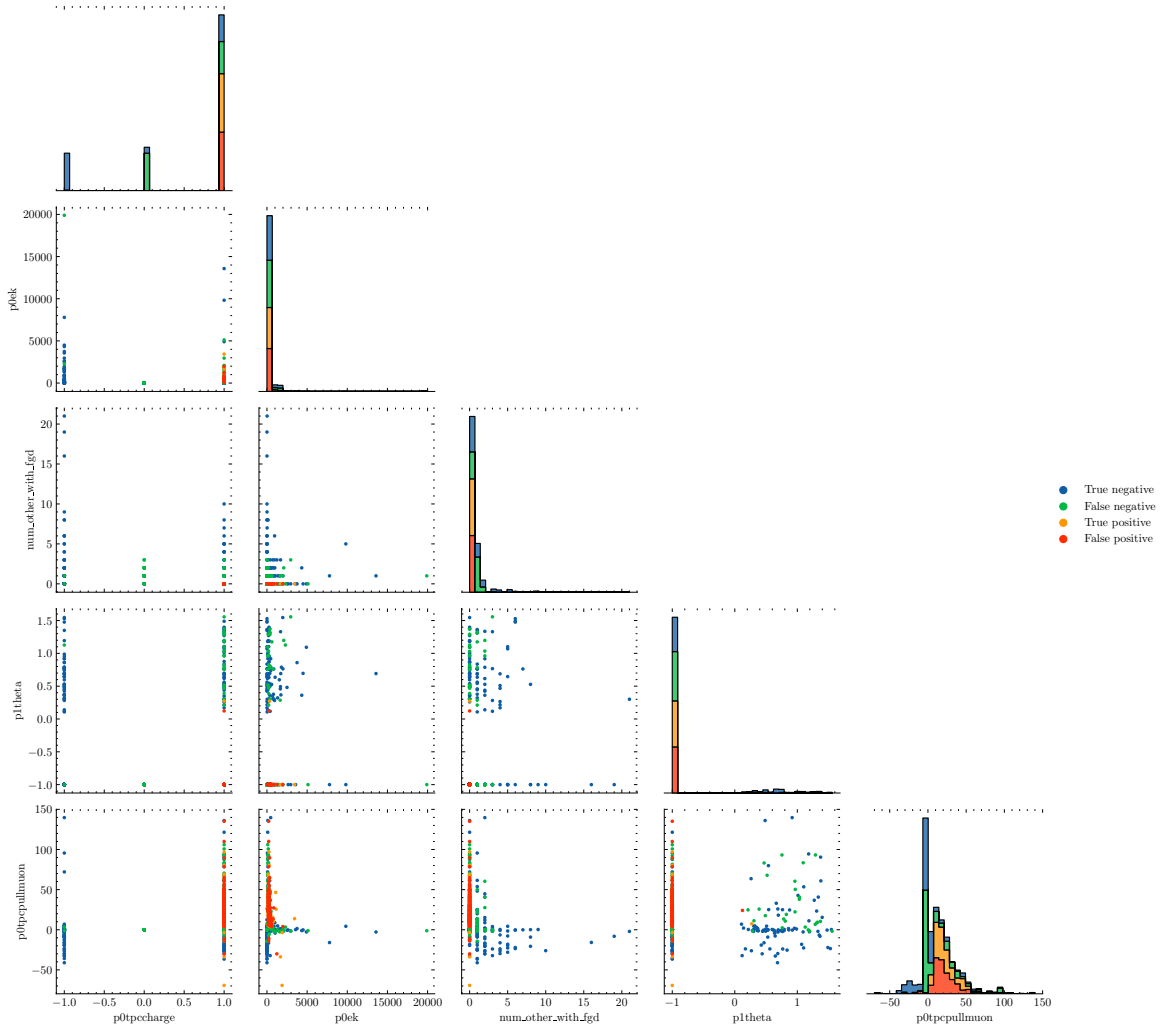

Fig. 7.15 Pairwise scatter diagrams and distributions for the five most important variables for the multi-variate signal selection, categorised by true negatives, false negatives, true positives and false positives. The population per category are made equal.

CHAPTER 8

MEASUREMENT OF THE NC0π CROSS SECTION

8.1 Systematic Uncertainties

In the next section, Section 8.2, the background contamination in the signal sample is determined by a simultaneous fit of MC to data in the signal and control samples. Sources of discrepancy between data and simulation are each represented in this fit by a parameter which modifies the simulated spectra accordingly. These systematic parameters can be considered *nuisance parameters* in the sense that their specific values are not of immediate interest, but rather serve to propagate systematic uncertainties through to the extracted background rates. All systematic parameters are assigned a Gaussian prior probability distribution with central values set to their nominal values and widths set to their prior uncertainty. The fitting routine requires that the simulated spectra can be evaluated for arbitrary values of these parameters; to achieve this, the simulated spectra are evaluated for parameter values set at $P + n\delta P$, where P and δP is the nominal value and prior uncertainty of the parameter and $\{n \in \mathbb{Z}, -5 \leq n \leq 5\}$ and linearly interpolated between. Studying the effect of uncertainties on simulation outputs requires that the data is known for arbitrary perturbations of its physics inputs; however, running the simulation multiple times is often unfeasible in terms of computational resources. Depending on the type of uncertainty, the simulated spectra are generated for these parameter values either by obtaining a covariance matrix or by the re-weighting approach.

The uncertainties considered by this analysis are either associated with the neutrino flux model, the interaction model or the modelling of the detector physics. These uncertainties and their implementation will be described.

8.1.1 Flux uncertainties

The fluxes of each neutrino species[1] within 0.0 GeV to 30.0 GeV at the near-detector are readily available to T2K analysers[2]. These are calculated with simulations of proton–target interactions and their transport through the target, horns and decay volume; the uncertainties—of which hadron interactions are dominant—are evaluated with external data from CERN's NA61/SHINE experiment and in-situ measurements of the beam properties [101]. For ν_μ, there are 11 bins with edges

$$[0.0, 0.4, 0.5, 0.6, 0.7, 1.0, 1.5, 2.5, 3.5, 5.0, 7.0, 30.0]\,\text{GeV} \tag{8.1}$$

for $\bar{\nu}_\mu$, there are 7 bins with edges

$$[0.0, 0.7, 1.0, 1.5, 2.5, 30.0]\,\text{GeV} \tag{8.2}$$

for ν_e, there are 5 bins with edges

$$[0.0, 0.5, 0.7, 0.8, 1.5, 2.5, 4.0, 30.0]\,\text{GeV} \tag{8.3}$$

and for $\bar{\nu}_e$, there are 2 bins with edges

$$[0.0, 2.5, 30.0]\,\text{GeV}. \tag{8.4}$$

A systematic parameter is created for each of these bins that scale its nominal value. The diagonal of the covariance matrix with respect to neutrino energy and species is taken as their prior uncertainties; the prior uncertainties of these 25 parameters can be seen in Table 8.2.

[1] With the exception of ν_τ, as the flux is assumed to be negligible at the near-detector's baseline.

[2] Version 13a1.1 of the pre-BANFF-fit flux matrix is used here.

8.1.2 Interaction model uncertainties

Uncertainties arising from the interaction model are propagated using GENIE's event re-weighting framework[3], interfaced through the T2KReweight package[4]. Event re-weighting is an approach to uncertainty propagation in which additional weights are generated and applied to each event to emulate the simulation outputs according to a modification of its inputs. In GENIE, each physics input P to the event generator is controlled by a *dial*, x_P, that modifies P according to

$$P' = P\left(1 + x_P \frac{\delta P}{P}\right) \tag{8.5}$$

where P is the nominal value of the physics input and δP is its prior uncertainty. A dial x_P set at zero sets P at its nominal value; a dial set at unity sets P at one standard deviation away from its nominal value. The complete list of dials implemented in GENIE can be found in Ref. [134].

Given a value of x_P, the weight assigned to each event is calculated differently for cross section dials, hadronisation and resonance decay dials, and intranuclear hadron transport dials. The cross section dials simply scale the neutrino interaction probability directly and are therefore calculated as

$$w = \frac{d^n\sigma'/dK^n}{d^n\sigma/dK^n} \tag{8.6}$$

where $d^n\sigma/dK^n$ is the nominal differential cross section with respect to the kinematic phase space K^n associated with the event class[5] and $d^n\sigma'/dK^n$ is the same except calculated with modified dials. The procedure is less straightforward for the remainder of the dials therefore Ref. [134] should be referred to.

The set of GENIE dials considered in this analysis can be found in Table 8.1. The set of weights for each parameter are evaluated at x_P values of $\{x_P \in \mathbb{Z} - 5 \leq x_P \leq 5\}$. Three additional parameters are added that scale the rates of the two dominant backgrounds in the signal sample. Normalisation uncertainties are each assigned to OOFV neutrons and CC interactions with prior uncertainties of 100%—loose priors are given to allow these parameters to be constrained by data in the control samples rather than by assumptions drawn from the MC. The last parameter controls the OOFV background rates based on the parent nucleus of the neutron. A weight w is assigned to these events calculated as

[3] Release 2.8.6.

[4] Release 1.27.2.

[5] For CC and NCE, K^n is simply the momentum transfer Q^2 [134].

$$w = 1 + c\frac{P}{\delta P}\left(Z_T - Z_{\mathrm{Fe}}\right) \tag{8.7}$$

where c is a constant, P is the parameter value, δP is the prior uncertainty of P, and Z_T and Z_{Fe} are the number of protons carried by the parent nucleus and Fe respectively. The value of c is set to 0.0178 such that a 1σ variation of this parameter cover MC-data discrepancies in the C ($\sim$30%) and Pb ($\sim$90%) differential cross sections observed by MINERνA [135, 136].

Table 8.1 List of neutrino interaction cross section systematic parameters considered in the sideband-fit as implemented in GENIE's event re-weighting package. The alias corresponds to the name of the parameter within T2KReWeight. For parameters without either an assigned nominal value or fractional uncertainty, the reweighting strategy may be found in Ref. [134].

Alias	Description	Nominal Value	Fractional Uncertainty
GXSec_MaNCEL	Axial mass for NC elastic	0.99 GeV	0.25
GXSec_EtaNCEL	Strange axial form factor	0.12	0.3
GXSec_MaCCQE	Axial mass for CC quasi-elastic	0.99 GeV	0.25
GXSec_MaCCRES	Axial mass for CC resonance production	0.99 GeV	0.2
GXSec_MvCCRES	Vector mass for CC resonance production	0.84 GeV	0.1
GXSec_MaNCRES	Axial mass for NC resonance production	0.99 GeV	0.2
GXSec_MvNCRES	Vector mass for NC resonance production	0.84 GeV	0.1
GXSec_RvpCC1pi	Scale for non-resonant background in ν–p CC1π	1.00	0.5
GXSec_RvpCC2pi	Scale for non-resonant background in ν–p CC2π	1.00	0.5
GXSec_RvpNC1pi	Scale for non-resonant background in ν–p NC1π	1.00	0.5
GXSec_RvpNC2pi	Scale for non-resonant background in ν–p NC2π	1.00	0.5
GXSec_RvnCC1pi	Scale for non-resonant background in ν–n CC1π	1.00	0.5
GXSec_RvnCC2pi	Scale for non-resonant background in ν–n CC2π	1.00	0.5
GXSec_RvnNC1pi	Scale for non-resonant background in ν–n NC1π	1.00	0.5
GXSec_RvnNC2pi	Scale for non-resonant background in ν–n NC2π	1.00	0.5
GXSec_RvbarpCC1pi	Scale for non-resonant background in $\bar{\nu}$–p CC1π	1.00	0.5
GXSec_RvbarpCC2pi	Scale for non-resonant background in $\bar{\nu}$–p CC2π	1.00	0.5

Continued on next page

Alias	Description	Nominal Value	Fractional Uncertainty
GXSec_RvbarpNC1pi	Scale for non-resonant background in $\bar{\nu}$–p NC1π	1.00	0.5
GXSec_RvbarpNC2pi	Scale for non-resonant background in $\bar{\nu}$–p NC2π	1.00	0.5
GXSec_RvbarnCC1pi	Scale for non-resonant background in $\bar{\nu}$–n CC1π	1.00	0.5
GXSec_RvbarnCC2pi	Scale for non-resonant background in $\bar{\nu}$–n CC2π	1.00	0.5
GXSec_RvbarnNC1pi	Scale for non-resonant background in $\bar{\nu}$–n NC1π	1.00	0.5
GXSec_RvbarnNC2pi	Scale for non-resonant background in $\bar{\nu}$–n NC2π	1.00	0.5
GXSec_AGKY_xF1pi	Pion Feynman-x PDF for N states in AGKY (hadronisation model)	—	—
GXSec_AGKY_pT1pi	Pion transverse momentum PDF for N states in AGKY (hadronisation model)	—	—
GSystNucl_CCQEPauliSupViaKF	CCQE Pauli suppression factor k_F	—	0.35
GRDcy_BR1gamma	Branching ratio for radiative resonance decays	—	0.5
GRDcy_BR1eta	Branching ratio for single-η resonance decays	—	0.5
GRDcy_Theta_Delta2NPi	Choice of pion angular distribution in $\Delta \to \pi N$ (isotropic / Rein-Sehgal)	—	—
GINuke_MFP_pi	Pion mean free path (total rescattering probability)	—	0.2
GINuke_FrCEx_pi	Pion charge exchange probability	—	0.5
GINuke_FrInel_pi	Pion inelastic reaction probability	—	0.4
GINuke_FrAbs_pi	Pion absorption probability	—	0.2
GINuke_FrPiProd_pi	Pion π-production probability	—	0.2
GINuke_MFP_N	Nucleon mean free path (total rescattering probability)	—	0.2

Continued on next page

Alias	Description	Nominal Value	Fractional Uncertainty
GINuke_FrCEx_N	Nucleon charge exchange probability	—	0.5
GINuke_FrInel_N	Nucleon inelastic reaction probability	—	0.4
GINuke_FrAbs_N	Nucleon absorption probability	—	0.2
GINuke_FrPiProd_N	Nucleon π-production probability	—	0.2

8.1.3 Detector physics uncertainties

The detector systematic parameters—excluding the last two systematic uncertainties in this subsection—each scale the event rate in a given bin for a given analysis sample. The uncertainty for each of these systematic parameters is contained within a covariance matrix of size $N_{samples} \times N_{bins}$, where $N_{samples}$ is the number of analysis samples and N_{bins} is the number of bins that the background subtraction is performed under. Instead of passing each of these matrices to the fit, they are integrated into a single covariance matrix representing the total detector uncertainty. The covariance matrix representing the total detector uncertainty is calculated as the sum over all detector systematic uncertainties plus a diagonal matrix representing the MC statistical uncertainty in each bin. The covariance matrix representing the MC statistical uncertainty M_{ij} is simply calculated as $M_{ij} = \delta_{ij}/N_i$, where δ_{ij} is the Kronecker delta function and N_i is the number of MC events in analysis bin i.

Common near-detector tracker uncertainties

Most of the systematic uncertainties considered in this analysis are shared with those in other T2K analyses and are readily available for calculation in T2K's dedicated analysis framework HighLAND2[6]. The systematic uncertainties are propagated by *pseudo-experiments*—the selection is performed multiple times in which random values of the systematic parameters are drawn from their assigned probability density functions, the systematic is the root-mean-square of the number of selected events over all pseudo-experiments. In general, the probability density function is almost always Gaussian.

In all cases, the systematic errors are propagated into a covariance matrix with pseudo-experiments. The method with which a particular systematic uncertainty is propagated is appropriate to the nature of the uncertainty itself; in HighLAND2, uncertainties are categorised as one of three classes from which their propagation method follows. These classes, and error propagation methods, are:

(a) **Efficiency-like**: encompasses systematic uncertainties tied with reconstruction efficiencies at the track-level—these uncertainties with respect to a particular analysis sample are estimated by comparing data and MC in well-understood control samples. The following relation is assumed to hold

$$\varepsilon^{data} = \frac{\varepsilon_0^{data}}{\varepsilon_0^{MC}} \varepsilon^{MC} \tag{8.8}$$

[6]Version 2.67, built with version 12.31 of the ND280 software.

where ε_0^{data} and ε_0^{MC} are the efficiencies in the MC and data control samples, and ε^{MC} and $\Delta\varepsilon^{data}$ are the efficiencies in the analysis sample predicted by the MC and measured in the data respectively; the efficiency in the MC analysis sample can be calculated by comparing the reconstruction with truth-level information. The variation in predicted efficiency for the data analysis sample is given by

$$\Delta\varepsilon^{data} = \frac{\varepsilon_0^{data} + \delta_0^{data} \cdot \sigma_0^{data}}{\varepsilon_0^{MC} + \delta_0^{MC} \cdot \sigma_0^{MC}} \tag{8.9}$$

where σ_0^{data} and σ_0^{MC} are the statistical uncertainties on the efficiency calculated on the data and MC control samples, and δ_0^{data} and δ_0^{MC} are the variations in the number of standard deviations in the data and MC control samples.

The track-level efficiencies of a given event are consolidated into an event-level efficiency by a process of re-weighting. If the track is correctly reconstructed then the event is assigned a weight calculated as

$$W = \frac{\Delta\varepsilon^{data}}{\varepsilon^{MC}}. \tag{8.10}$$

Conversely, if the track is incorrectly reconstructed then the event is assigned a weight calculated as

$$W = \frac{1 - \Delta\varepsilon^{data}}{1 - \varepsilon^{MC}}. \tag{8.11}$$

The detector-level systematic uncertainties considered in this analysis that are labelled efficiency-like are:

- **TPC charge identification**: a charged particle may be reconstructed with the wrong sign charge given limited information about a particle's trajectory. The uncertainty is calculated from the discrepancy in charge misidentification rates between data and MC using a sample consisting of tracks that cross all three TPCs.
- **TPC tracking efficiency**: accounts for the reconstruction fail rate of tracks with a TPC component. The uncertainty is calculated, again, with a sample of through-going μ^-. The reconstruction algorithm is high-performing: the efficiency is upwards of 99% for events with more than 18 hits and bears no dependence on the particle kinematics.

- **TPC–FGD track matching efficiency**: particle tracks that cross a TPC–FGD interface can be wrongly reconstructed as two separate tracks by the reconstruction algorithm. In the case that the track begins inside the FGD and crosses into the TPC, this effect reduces the signal efficiency; in the case that the track crosses into an FGD from a TPC, this effect contributes to the OOFV background. The systematic uncertainty on the matching failure rates between data and MC is estimated with a sample of through-going μ^-. The matching efficiency was found to be upwards of 99% except in regions of low momenta.
- **TPC cluster efficiency**: a cluster, within the TPC, is a group of neighbouring single hits which corresponds to the smallest resolved subsection of the ionisation trace. The efficiency with which the reconstruction algorithm is able to identify clusters is particularly relevant to the TPC track quality. The systematic uncertainty is the relative difference in efficiency between data and MC using a sample of mostly μ^-.
- **Michel electron tagging efficiency**: Michel electrons are identified by delayed activity in the FGDs, at least 100 ns after the neutrino interaction. Given they are indicative of CC interactions, the misidentification of Michel electrons can contaminate a sample with CC events. The systematic uncertainty on the efficiency is calculated using a sample of cosmic μ^- stopping in FGD1.

(b) **Observable-like**: encompasses systematic uncertainties tied with discrepancies in scale or resolution between data and simulation of reconstructed observables. These are calculated by the *variation* method, in which the observable is varied according to

$$x' = x + \Delta x + \delta \cdot \sigma_{\Delta x} \tag{8.12}$$

where x is the value of the observable, Δx is the correction necessary to match the observed value in MC to data, $\sigma_{\Delta x}$ is the statistical uncertainty associated with Δx and δ is the variations in number of standard deviations.

The detector-level systematic uncertainties considered in this analysis that are labelled observable-like are:

- **Magnetic field distortions**: ionisation electrons induced by particles crossing the TPC drift towards the readout plane in the direction of the electric and magnetic field lines within the chamber. Non-uniformities in the field established by the near-detector magnet, however, can alter the course of charged particles in the

plane transverse to their drift direction—ultimately biasing the measurement of their momenta when they arrive at the readout plane. These field distortions can be corrected for empirically by comparing simulation data of drift electrons with photo-electron calibration data collected in the TPC. As the cause of these distortions is not greatly understood, these corrections are not applied directly to the MC; instead, these corrections are used to form a systematic error based on the discrepancy in reconstructed momenta per track with and without the corrections applied.

- **TPC momentum scale**: measurements of momenta depend on prior knowledge of the magnetic field's strength and direction throughout the near-detector basket—this was measured directly before the near-detector was constructed.
- **TPC momentum resolution**: measurements of momenta are limited to the detector granularity. To estimate this systematic, a sample was constructed consisting of μ that cross more than one TPC (and hence at least one FGD).
- **TPC particle identification**: particle types are classified in the TPC according to their energy loss; how closely the energy loss of a reconstructed track resembles the characteristic energy loss of a particular particle class is represented by a pull. The TPC particle identification systematic characterises the discrepancy in particle pull distributions between data and MC. To this end, two high-purity samples were constructed—one consisting of protons, the other of sand muons; the systematic uncertainty was calculated as the relative difference between data and MC in central values of a Gaussian function fitted to the pull distribution.

(c) **Normalisation-like**: encompasses systematic uncertainties tied with a discrepancy in rate, or normalisation, of some process between data and MC; this is dealt with by re-weighting at the event-level. The weight applied to the event is given by

$$W = W_0 (1 + \delta \cdot \sigma_W) \tag{8.13}$$

where W_0 is the nominal weight, before any systematic uncertainties have been applied, σ_W is the systematic error on the normalisation and δ is the variation in number of standard deviations.

The detector-level systematic uncertainties considered in this analysis that are labelled normalisation-like are:

- **Proton secondary interactions**: protons induced by neutrino interactions can undergo interactions in the detector material. A systematic uncertainty for these

rates is estimated by comparing predictions from GEANT4 with proton scattering data.

- **Pion secondary interactions**: pions also undergo secondary interactions. A systematic uncertainty for these rates is also estimated by comparing GEANT4 predictions with external data.

Pile-up on veto cut efficiency

Detector pile-up causes a deficit in the sub-detector veto efficiencies in the data compared to MC—in the signal sample, the deficit is greatest in the ECal veto efficiency (3%). The efficiencies are calculated with the fraction of bunches passing the veto cuts relative to the number of bunches. The fractional difference between these efficiencies for data and MC is implemented as a systematic uncertainty which is taken as 100% correlated in the signal and Michel tag samples, given that they apply the ECal veto cut, and uncorrelated between all other samples.

SMRD object reconstruction efficiency

The reconstruction efficiency of the SMRD was studied with a sample of beam-induced muons. A sample overwhelmingly composed of muons that pass into the SMRD was constructed with the following event selection: *i)* the event is of good quality, *ii)* the event has a negative charged TPC track matched to a barrel ECal object, *iii)* the ECal object has hits in every layer (ensuring the track eclipsed the ECal and passed into the magnet), *iv)* the TPC object at its endpoint is headed for an SMRD module, *v)* the tracks have a polar angle that satisfies $\theta - \pi/2 < 1/4$ and a momenta within the range 0.75 GeV to 5 GeV. The reconstruction efficiency is defined to be the fraction of events that have an object reconstructed in the SMRD module the TPC object was headed for. The discrepancy in efficiency between data and MC varies by SMRD module but is largest in the top module with a discrepancy of 5%. The fractional difference between these efficiencies is implemented as a normalisation uncertainty applied to events in the SMRD control sample.

Neutron secondary interactions

Comparisons of neutron interaction cross sections between GEANT4 and external data show some discrepancy [137]. Given that $\sim$20% of the OOFV neutron background undergo secondary interactions, this is accounted for by the addition of a normalisation uncertainty on events that were products of neutron secondary interactions. This parameter is assigned a prior uncertainty of 10% in accordance with the size of this discrepancy in the energy range of interest to this analysis.

P0D object reconstruction efficiency

The reconstruction efficiency of the P0D has been studied by multiple studies comparing sand and cosmic muon kinematics in data and MC [138, 139]. Excellent agreement between MC and data was found in their reconstruction efficiencies. Perfect efficiency is achieved in MC, whereas an upper limit of $\leq$2% inefficiency was measured of the data. This inefficiency limit is implemented as a normalisation uncertainty applied to events with at least one reconstructed object in the P0D.

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
0	`f_numu_0`	ν_μ flux with energy 0.0 GeV to 0.4 GeV.	1.000 ± 0.066
1	`f_numu_1`	ν_μ flux with energy 0.4 GeV to 0.5 GeV.	1.000 ± 0.057
2	`f_numu_2`	ν_μ flux with energy 0.5 GeV to 0.6 GeV.	1.000 ± 0.054
3	`f_numu_3`	ν_μ flux with energy 0.6 GeV to 0.7 GeV.	1.000 ± 0.048
4	`f_numu_4`	ν_μ flux with energy 0.7 GeV to 1.0 GeV.	1.000 ± 0.070
5	`f_numu_5`	ν_μ flux with energy 1.0 GeV to 1.5 GeV.	1.000 ± 0.084
6	`f_numu_6`	ν_μ flux with energy 1.5 GeV to 2.5 GeV.	1.000 ± 0.062
7	`f_numu_7`	ν_μ flux with energy 2.5 GeV to 3.5 GeV.	1.000 ± 0.066
8	`f_numu_8`	ν_μ flux with energy 3.5 GeV to 5.0 GeV.	1.000 ± 0.084
9	`f_numu_9`	ν_μ flux with energy 5.0 GeV to 7.0 GeV.	1.000 ± 0.102
10	`f_numu_10`	ν_μ flux with energy 7.0 GeV to 30.0 GeV.	1.000 ± 0.119
11	`f_nue_0`	ν_e flux with energy 0.0 GeV to 0.5 GeV.	1.000 ± 0.102
12	`f_nue_1`	ν_e flux with energy 0.5 GeV to 0.7 GeV.	1.000 ± 0.072
13	`f_nue_2`	ν_e flux with energy 0.7 GeV to 0.8 GeV.	1.000 ± 0.069
14	`f_nue_3`	ν_e flux with energy 0.8 GeV to 1.5 GeV.	1.000 ± 0.077
15	`f_nue_4`	ν_e flux with energy 1.5 GeV to 2.5 GeV.	1.000 ± 0.075
16	`f_nue_5`	ν_e flux with energy 2.5 GeV to 4.0 GeV.	1.000 ± 0.056
17	`f_nue_6`	ν_e flux with energy 4.0 GeV to 30.0 GeV.	1.000 ± 0.056
18	`f_numubar_0`	$\bar{\nu}_\mu$ flux with energy 0.0 GeV to 0.7 GeV.	1.000 ± 0.052
19	`f_numubar_1`	$\bar{\nu}_\mu$ flux with energy 0.7 GeV to 1.0 GeV.	1.000 ± 0.062

Continued on next page

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
20	f_numubar_2	$\bar{\nu}_\mu$ flux with energy 1.0 GeV to 1.5 GeV.	1.000 ± 0.087
21	f_numubar_3	$\bar{\nu}_\mu$ flux with energy 1.5 GeV to 2.5 GeV.	1.000 ± 0.087
22	f_numubar_4	$\bar{\nu}_\mu$ flux with energy 2.5 GeV to 30.0 GeV.	1.000 ± 0.099
23	f_nuebar_0	$\bar{\nu}_e$ flux with energy 0.0 GeV to 2.5 GeV.	1.000 ± 0.100
24	f_nuebar_1	$\bar{\nu}_e$ flux with energy 2.5 GeV to 30.0 GeV.	1.000 ± 0.137
25	DetSyst_MichelTag_Ek_0	Detector covariance bin with index 0 for Michel electron control sample.	—
26	DetSyst_MichelTag_Ek_1	Detector covariance bin with index 1 for Michel electron control sample.	—
27	DetSyst_MichelTag_Ek_2	Detector covariance bin with index 2 for Michel electron control sample.	—
28	DetSyst_MichelTag_Ek_3	Detector covariance bin with index 3 for Michel electron control sample.	—
29	DetSyst_MichelTag_Ek_4	Detector covariance bin with index 4 for Michel electron control sample.	—
30	DetSyst_P0D_Ek_0	Detector covariance bin with index 0 for P0D control sample.	—
31	DetSyst_P0D_Ek_1	Detector covariance bin with index 1 for P0D control sample.	—

Continued on next page

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
32	DetSyst_P0D_Ek_2	Detector covariance bin with index 2 for P0D control sample.	—
33	DetSyst_P0D_Ek_3	Detector covariance bin with index 3 for P0D control sample.	—
34	DetSyst_P0D_Ek_4	Detector covariance bin with index 4 for P0D control sample.	—
35	DetSyst_SMRD_Ek_0	Detector covariance bin with index 0 for SMRD control sample.	—
36	DetSyst_SMRD_Ek_1	Detector covariance bin with index 1 for SMRD control sample.	—
37	DetSyst_SMRD_Ek_2	Detector covariance bin with index 2 for SMRD control sample.	—
38	DetSyst_SMRD_Ek_3	Detector covariance bin with index 3 for SMRD control sample.	—
39	DetSyst_SMRD_Ek_4	Detector covariance bin with index 4 for SMRD control sample.	—
40	DetSyst_Signal_Ek_0	Detector covariance bin with index 0 for signal sample.	—
41	DetSyst_Signal_Ek_1	Detector covariance bin with index 1 for signal sample.	—
42	DetSyst_Signal_Ek_2	Detector covariance bin with index 2 for signal sample.	—
43	DetSyst_Signal_Ek_3	Detector covariance bin with index 3 for signal sample.	—

Continued on next page

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
44	DetSyst_Signal_Ek_4	Detector covariance bin with index 4 for signal sample.	—
45	Neutron SI	Rate of secondary interactions of neutrons.	1.000 ± 0.100
46	P0D Recon Efficiency	Reconstruction efficiency of objects that traverse the P0D.	1.000 ± 0.002
47	GenieSyst_GHadrAGKY_pT1pi	Pion transverse momentum PDF for N states in AGKY (hadronisation model).	0.000 ± 1.000
48	GenieSyst_GHadrAGKY_xF1pi	Pion Feynman-x PDF for N states in AGKY (hadronisation model).	0.000 ± 1.000
49	GenieSyst_GINuke_FrAbs_N	Nucleon absorption probability.	0.000 ± 1.000
50	GenieSyst_GINuke_FrAbs_pi	Pion absorption probability.	0.000 ± 1.000
51	GenieSyst_GINuke_FrCEx_N	Nucleon charge exchange probability.	0.000 ± 1.000
52	GenieSyst_GINuke_FrCEx_pi	Pion charge exchange probability.	0.000 ± 1.000
53	GenieSyst_GINuke_FrInel_N	Nucleon inelastic reaction probability.	0.000 ± 1.000
54	GenieSyst_GINuke_FrInel_pi	Pion inelastic reaction probability.	0.000 ± 1.000
55	GenieSyst_GINuke_FrPiProd_N	Nucleon π-production probability.	0.000 ± 1.000
56	GenieSyst_GINuke_FrPiProd_pi	Pion π-production probability.	0.000 ± 1.000
57	GenieSyst_GINuke_MFP_N	Nucleon mean free path (total rescattering probability).	0.000 ± 1.000
58	GenieSyst_GINuke_MFP_pi	Pion mean free path (total rescattering probability).	0.000 ± 1.000
59	GenieSyst_GRDcy_BR1eta	Branching ratio for single-η resonance decays.	0.000 ± 1.000
60	GenieSyst_GRDcy_BR1gamma	Branching ratio for radiative resonance decays.	0.000 ± 1.000

Continued on next page

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
61	GenieSyst_GRDcy_Theta_Delta2Npi	Choice of pion angular distribution in $\Delta \to \pi N$ (isotropic / Rein-Sehgal).	0.000 ± 1.000
62	GenieSyst_GSystNucl_CCQEPauliSupViaKF	CCQE Pauli suppression factor k_F.	0.000 ± 1.000
63	GenieSyst_GXSec_EtaNCEL	Strange axial form factor η for NC elastic.	0.000 ± 1.000
64	GenieSyst_GXSec_MaCCQE	Axial mass for CC quasi-elastic (0.99 GeV).	0.000 ± 1.000
65	GenieSyst_GXSec_MaCCRES	Axial mass for CC resonance production (0.99 GeV).	0.000 ± 1.000
66	GenieSyst_GXSec_MaNCEL	Axial mass for NC elastic (0.99 GeV).	0.000 ± 1.000
67	GenieSyst_GXSec_MaNCRES	Axial mass for NC resonance production (0.99 GeV).	0.000 ± 1.000
68	GenieSyst_GXSec_MvCCRES	Vector mass for CC resonance production (0.84 GeV).	0.000 ± 1.000
69	GenieSyst_GXSec_MvNCRES	Vector mass for NC resonance production (0.84 GeV).	0.000 ± 1.000
70	GenieSyst_GXSec_RvbarnCC1pi	Scale for non-resonant background in $\bar{\nu}$–n CC1π.	0.000 ± 1.000
71	GenieSyst_GXSec_RvbarnCC2pi	Scale for non-resonant background in $\bar{\nu}$–n CC2π.	0.000 ± 1.000
72	GenieSyst_GXSec_RvbarnNC1pi	Scale for non-resonant background in $\bar{\nu}$–n NC1π.	0.000 ± 1.000
73	GenieSyst_GXSec_RvbarnNC2pi	Scale for non-resonant background in $\bar{\nu}$–n NC2π.	0.000 ± 1.000
74	GenieSyst_GXSec_RvbarpCC1pi	Scale for non-resonant background in $\bar{\nu}$–p CC1π.	0.000 ± 1.000
75	GenieSyst_GXSec_RvbarpCC2pi	Scale for non-resonant background in $\bar{\nu}$–p CC2π.	0.000 ± 1.000
76	GenieSyst_GXSec_RvbarpNC1pi	Scale for non-resonant background in $\bar{\nu}$–p NC1π.	0.000 ± 1.000
77	GenieSyst_GXSec_RvbarpNC2pi	Scale for non-resonant background in $\bar{\nu}$–p NC2π.	0.000 ± 1.000
78	GenieSyst_GXSec_RvnCC1pi	Scale for non-resonant background in ν–n CC1π.	0.000 ± 1.000
79	GenieSyst_GXSec_RvnCC2pi	Scale for non-resonant background in ν–n CC2π.	0.000 ± 1.000

Continued on next page

Table 8.2 List of systematic parameters.

Index	Name	Description	Prior
80	`GenieSyst_GXSec_RvnNC1pi`	Scale for non-resonant background in ν–n NC1π.	0.000 ± 1.000
81	`GenieSyst_GXSec_RvnNC2pi`	Scale for non-resonant background in ν–n NC2π	0.000 ± 1.000
82	`GenieSyst_GXSec_RvpCC1pi`	Scale for non-resonant background in ν–p CC1π.	0.000 ± 1.000
83	`GenieSyst_GXSec_RvpCC2pi`	Scale for non-resonant background in ν–p CC2π.	0.000 ± 1.000
84	`GenieSyst_GXSec_RvpNC1pi`	Scale for non-resonant background in ν–p NC1π.	0.000 ± 1.000
85	`GenieSyst_GXSec_RvpNC2pi`	Scale for non-resonant background in ν–p NC2π.	0.000 ± 1.000
86	CC Background Normalisation	Scale for rate of CC interactions in FGD1.	1.000 ± 1.000
87	OOFV Neutron Background Normalisation	Scale for rate of OOFV neutron interactions in FGD1.	1.000 ± 1.000
88	OOFV Neutron A Dependence	Scale for the relative rate of OOFV neutron production depending on target nucleus.	5.000 ± 1.200

8.2 Background Subtraction

The background contamination in the data of each analysis sample is determined by a maximum likelihood fit[7] of MC to data in all samples simultaneously. The parameters in the fit are the systematic parameters as described in the previous section (Section 8.1) and an additional parameter which scales the signal rates. This extra parameter, which is the only parameter assigned a uniform prior uncertainty distribution, is included not for physics purposes, but rather to allow the small contamination of signal in the control samples to vary during the fit. The extraction of the signal cross section is described later in Section 8.5. This extra parameter is hence ignored when evaluating the uncertainty of the background rates extracted from the signal sample.

For a given binned observable x, the number of observed events n_i in each bin i adheres to Poisson statistics. The likelihood that the set of parameters $\vec{\theta}$ has generated the observed distribution of x is, therefore, the product of the probability in each bin that n_i events should be observed given that μ_i events were expected, where μ_i is controlled by θ. The likelihood function is therefore

$$\mathscr{L}(\theta) = \prod_i \frac{\mu_i^{n_i}\left(\vec{\theta}\right)}{n_i!} e^{-\mu_i(\vec{\theta})} \tag{8.14}$$

where i, in this analysis, indexes the 5 bins in reconstructed kinetic energy of the leading proton candidate. The edges of these bins are

$$E_K = [0.00, 180.71, 204.01, 310.56, 427.14, 1000.00]\,\text{MeV}. \tag{8.15}$$

The parameters should obey their prior distributions, hence the parameters with Gaussian priors contribute a penalty term to the likelihood function that increases as their values stray from their prior central values. The complete log-likelihood function is therefore given by

$$-2\ln\mathscr{L} = (\vec{\theta}^G - \hat{\theta}^G)^T V^{-1} (\vec{\theta}^G - \hat{\theta}^G) + 2\sum_j \sum_i \left[\mu_i(\vec{\theta}) - n_i + n_i \ln \frac{n_i}{\mu_i(\vec{\theta})}\right]_j \tag{8.16}$$

[7] The minimisation is performed by the MIGRAD routine of the MINUIT2 package [140].

where $\vec{\theta}^G$ is the set of parameters with Gaussian priors and $\hat{\theta}^G$ are their pre-fit central values, V is the covariance matrix of $\vec{\theta}$, i indexes the kinetic energy and j indexes the four analysis samples (i.e. the signal sample and the P0D, SMRD and Michel electron control samples).

In the previous section, it was described that the MC spectra can be generated for arbitrary perturbations of each parameter by interpolating between spectra evaluated at fixed parameter values. To generate spectra for arbitrary perturbations of a set of parameters, however, the nominal spectra are re-weighted. The weight that a particular parameter bears on a particular bin in a particular event class within a sample is the ratio of the bin value with the perturbed parameter with the nominal bin value. The total weight of the bin for a given sample, event class and set of parameters is the product of the weights calculated for each parameter perturbation.

The fit to the sidebands is used to extract 7 categories of events: *i) signal*: NC0π interactions within the fiducial volume and the detector phase space requirement, *ii) OOFV neutrons*: interactions of neutrons that cross into the fiducial from other subdetectors, *iii) NC background*: neutral current events in the fiducial volume with reconstructed π or other particles in the final state, *iv) CC*: charged current interactions within the fiducial volume, *v) Out-Of-Phase-Space (OOPS) NC*0π: NC0π interactions within the fiducial volume but outside of the detector phase space requirement, *vi) other*: all interactions that do not belong to any of the previous categories.

The pulls of the best-fit values of the post-fit systematic parameters are found in Figs. 8.1 and 8.2 for all analyses. The pull of a particular parameter P is defined as $(P_{\text{post}} - P_{\text{pre}})/\delta P_{\text{pre}}$, where P_{pre} and P_{post} are the prior and post-fit values of P and δP_{pre} is its prior uncertainty. In short, all post-fit values are consistent with their prior values to within one standard deviation. A common worry with the inclusion of flux parameters in cross section measurements is that large uncertainties on the interaction model parameters risk their over-constraining, leading to underestimated model systematic uncertainties. The parameter pulls demonstrate this effect is not at play here as the post-fit uncertainties of the flux parameters are consistent with their prior ones. Two parameters that are particularly constrained by the fit are those which parametrise the CC and OOFV neutron rates in the FGD (which are indexed as parameters 86 and 87), as they are assigned a large prior uncertainty.

The correlations between the systematic parameters before and after the fit may be observed in Figs. 8.3 to 8.6. The correlation matrices for all multivariate analyses are similar, hence only the matrix from the analysis conducted with the highest efficiency signal sample is shown. Before the fit, the parameters within the three groups of systematic parameters—flux, detector and theory—are uncorrelated with all parameters from different groups. Within their groups, the flux parameters and detector parameters are correlated; the theory parameters,

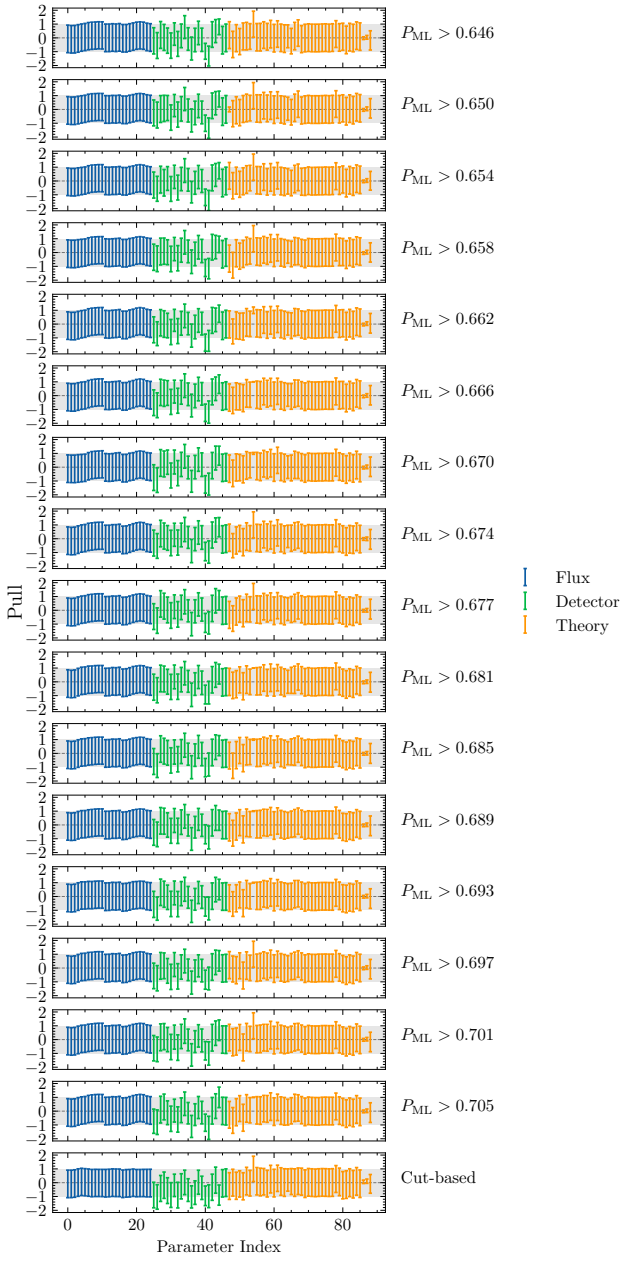

Fig. 8.1 Pulls of the best-fit values determined by the sideband fit for each parameter category. Each panel is an independent analysis in which the fit was performed with a signal sample as indicated by the panel label (panels labelled $P_{ML} > x$ indicate the signal probability threshold x of a multivariate selector was applied). Pull is defined as $(P_{\mathrm{post}} - P_{\mathrm{pre}})/\delta P_{\mathrm{pre}}$, where P_{pre} and P_{post} are the prior and post-fit values of parameter P and δP_{pre} is its prior uncertainty. The data in all analysis samples are real detector data. The parameters are indexed in Table 8.2.

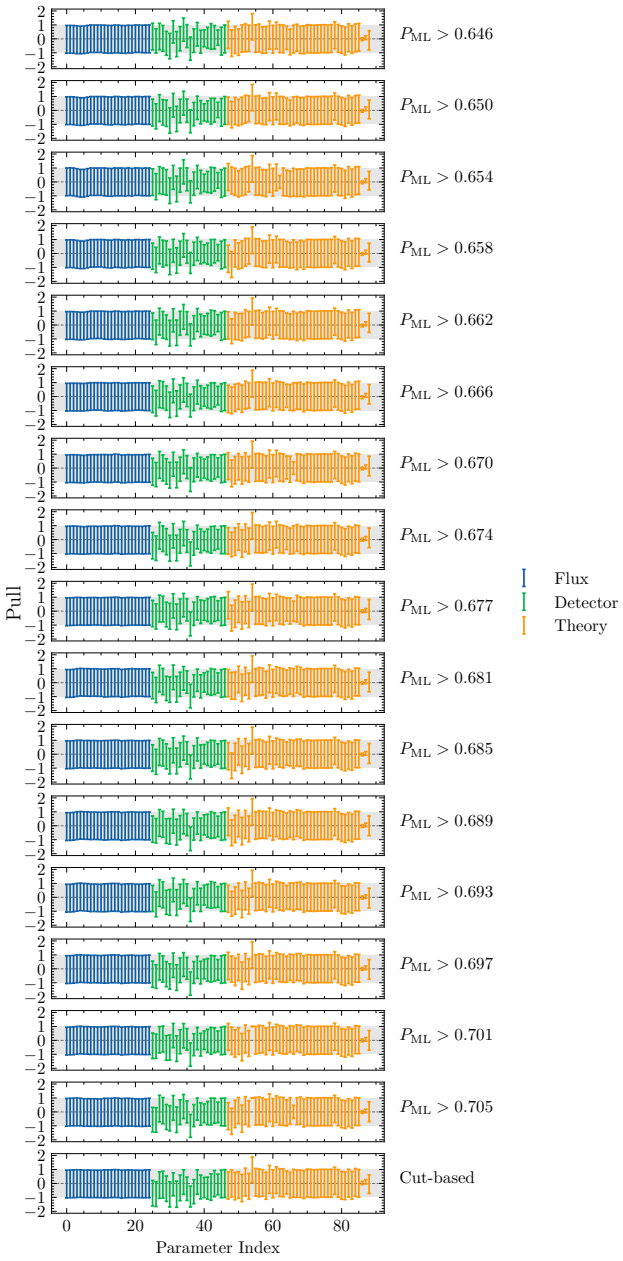

Fig. 8.2 Pulls of the best-fit values determined by the sideband fit for each parameter category. Each panel is an independent analysis in which the fit was performed with a signal sample as indicated by the panel label (panels labelled $P_{ML} > x$ indicate the signal probability threshold x of a multivariate selector was applied). Pull is defined as $(P_{\mathrm{post}} - P_{\mathrm{pre}})/\delta P_{\mathrm{pre}}$, where P_{pre} and P_{post} are the prior and post-fit values of parameter P and δP_{pre} is its prior uncertainty. The data in the control samples is real detector data whereas the data in the signal sample is substituted for the nominal Monte-Carlo. The parameters are indexed in Table 8.2.

however, are assumed to be uncorrelated with each other. Most notably, the fit introduces anti-correlations between the CC and OOFV neutron rate parameters and the flux and detector parameters. It will later be seen that these anti-correlations cause the individual contributions by each systematic parameter to the total fractional uncertainty on the background to not add in quadrature.

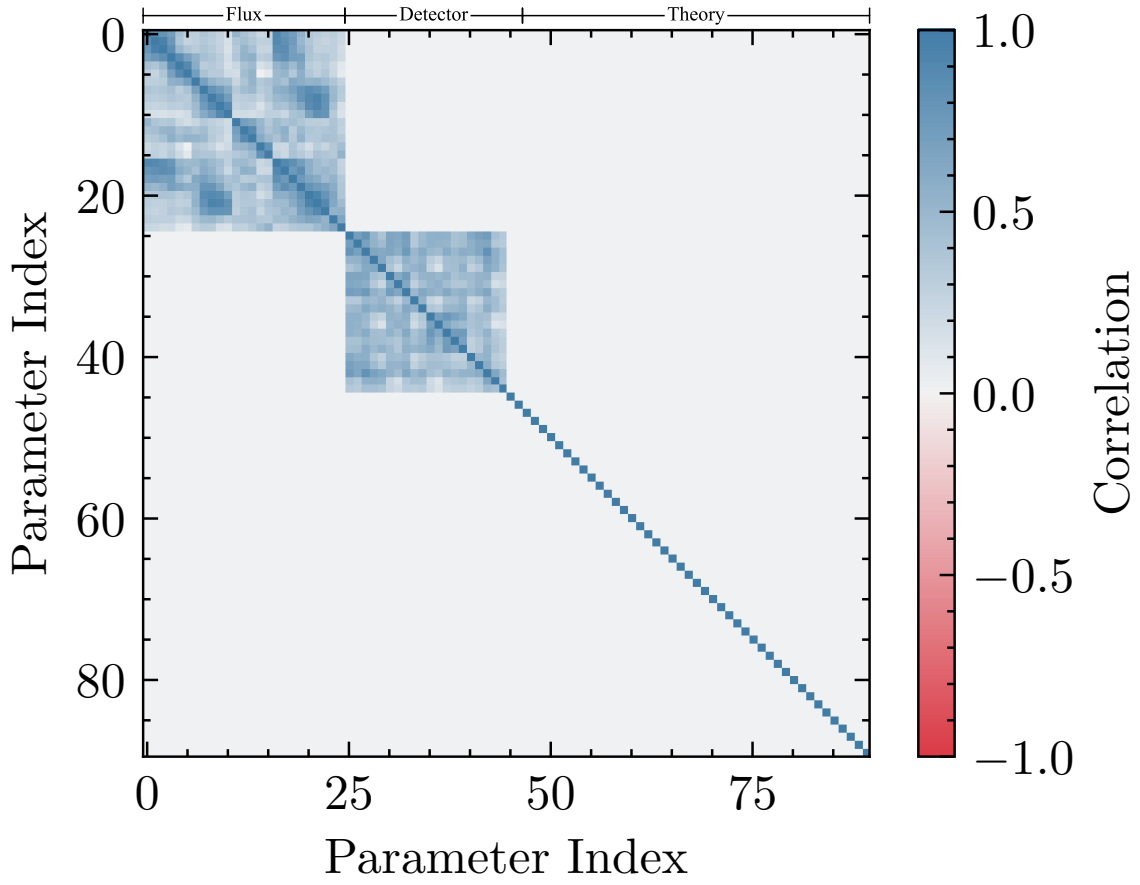

Fig. 8.3 Correlation matrix of systematic parameters before the sideband fit for the cut-based analysis.

In the multivariate analysis, the total number of events in the signal sample is underestimated by the MC by around 3% to 5% for all applied thresholds; in the cut-based analysis, the same quantity is overestimated by the MC by the same amount. The statistical significance of this discrepancy is between 1–2σ in both cases—the cause of which is unknown. The number of events in the SMRD and P0D is also underestimated by the MC by around $\sim$2% to 3% and $\sim$3% to 5% respectively, though the discrepancy is less significant. In the Michel electron control sample, the number of events is overestimated by the MC. These discrepancies between MC and data may be observed for all analysis samples in Fig. 8.7.

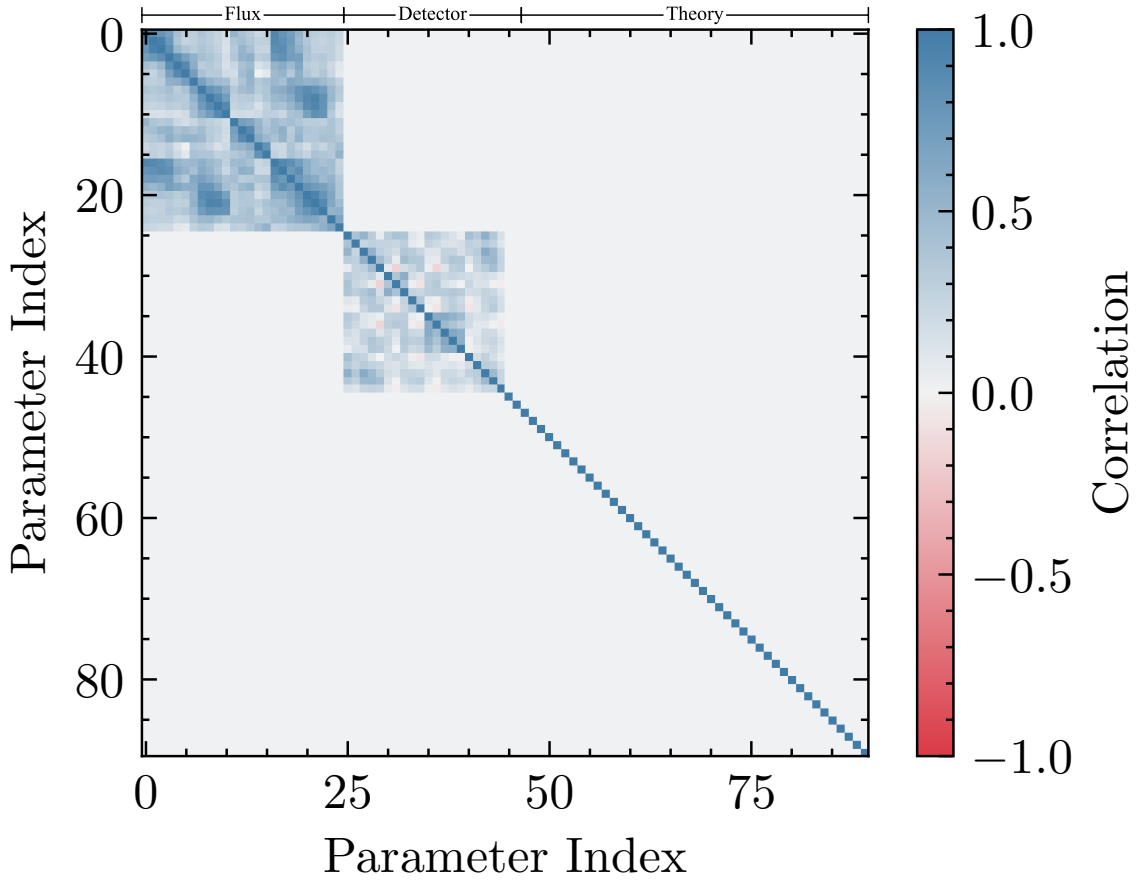

Fig. 8.4 Correlation matrix of systematic parameters before the sideband fit for the multivariate analysis.

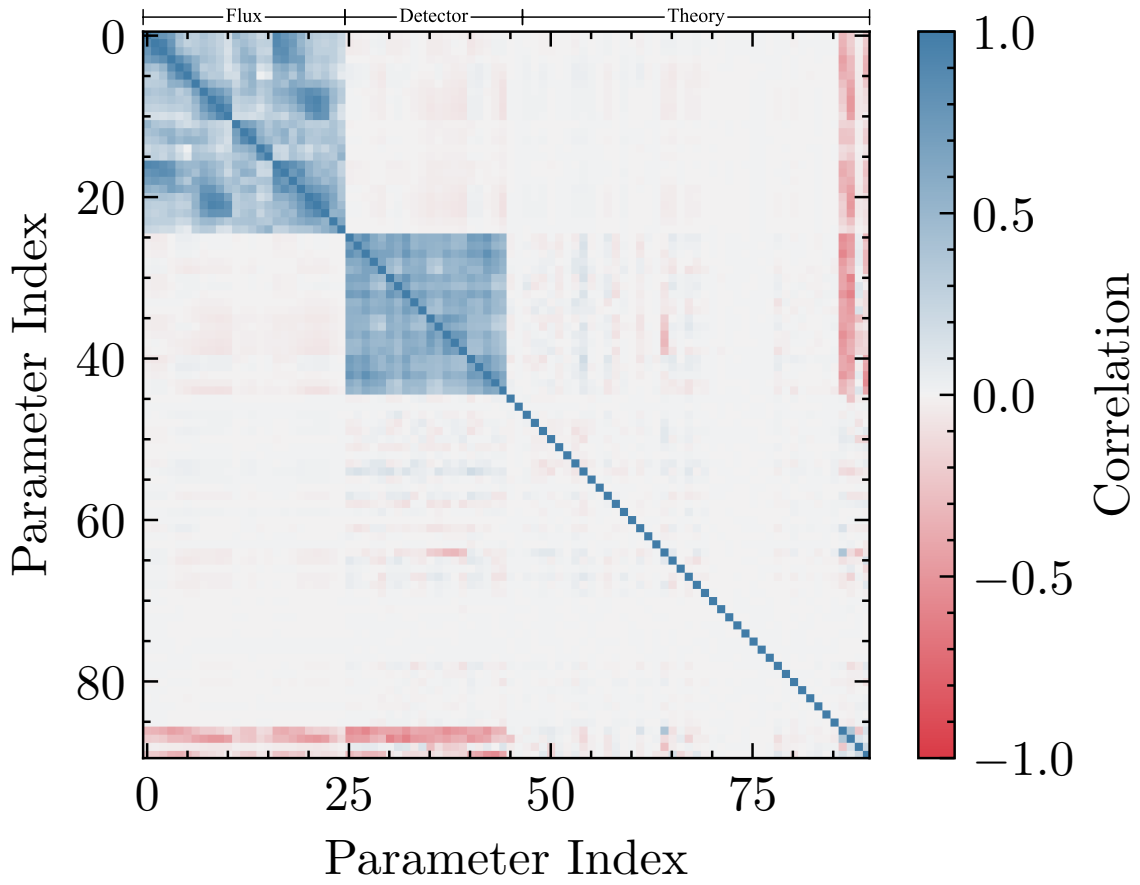

Fig. 8.5 Correlation matrix of systematic parameters after the sideband fit for the cut-based analysis.

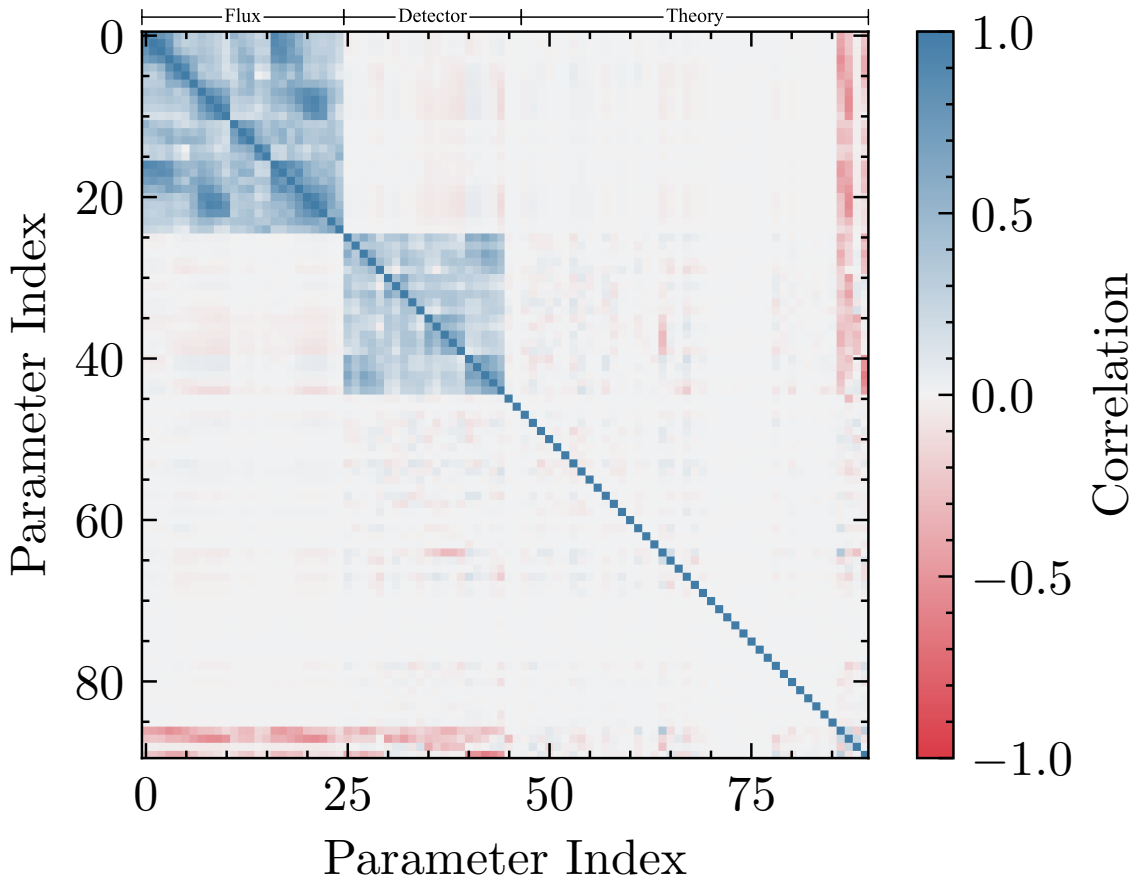

Fig. 8.6 Correlation matrix of systematic parameters after the sideband fit for the multivariate analysis.

After the fit, however, data and MC are made to agree to <1% for all samples and thresholds applied.

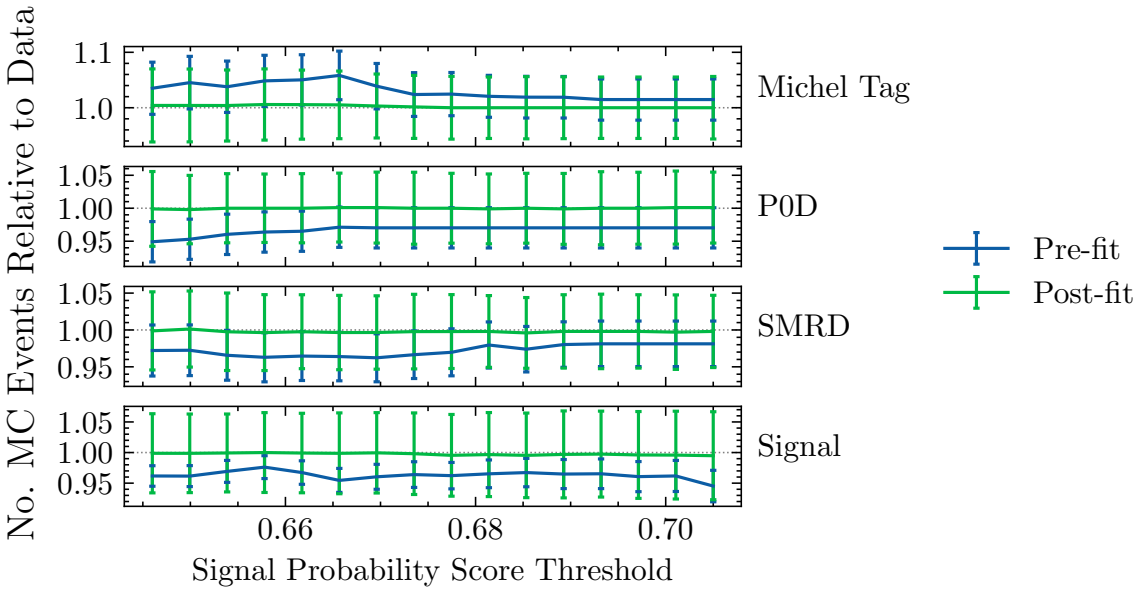

Fig. 8.7 Ratio of total number of MC events to data as a function of threshold applied to signal probability score.

Most importantly, the uncertainty on the number of background events in the signal sample is substantially reduced. This is evident by comparing Fig. 8.8a and Fig. 8.8b, and also Fig. 8.10a and Fig. 8.10b. Before the fit, the fractional uncertainty resides around 45% to 48% for all analyses with the multivariate signal samples; after the fit, the same quantity is dependent on the composition of the sample. The fractional uncertainty on the background in the signal sample is ∼7.5% for the most efficient signal selection and ∼10% for the purest—hence, the reduction in uncertainty achieved by the fit is around ∼80% on average. The uncertainty on the background is dominated by theoretical uncertainties in all analyses—the highest of which are the uncertainty on the OOFV neutron rates and the GENIE interaction model. The flux uncertainties are invariant to the composition of the signal sample and reside at the 5% level.

In Figs. 8.9 and 8.11, the uncertainty contributions achieved with the multivariate analyses are compared to those of the cut-based analysis. With the highest efficiency signal sample, the fractional uncertainty is about the same as that of the cut-based analysis; with the purest signal sample, the fractional uncertainty is increased by ∼25%. Though subdominant, the detector uncertainties are reduced in all samples for the multivariate analysis. In the signal sample, the reduction is ∼25% and appears to be largely driven by a reduction in the TPC PID uncertainty by an order of magnitude. In the multivariate analysis, the model assigns

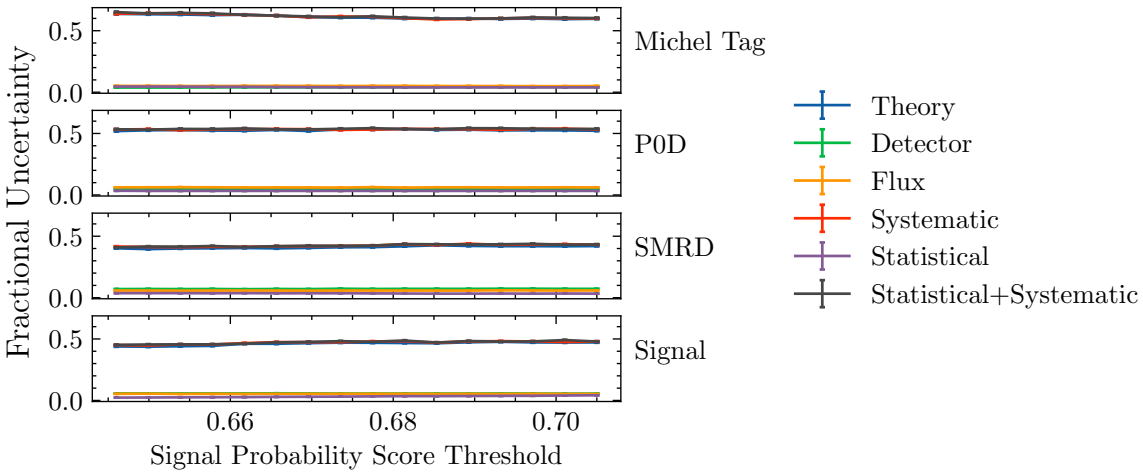

(a) Pre-fit.

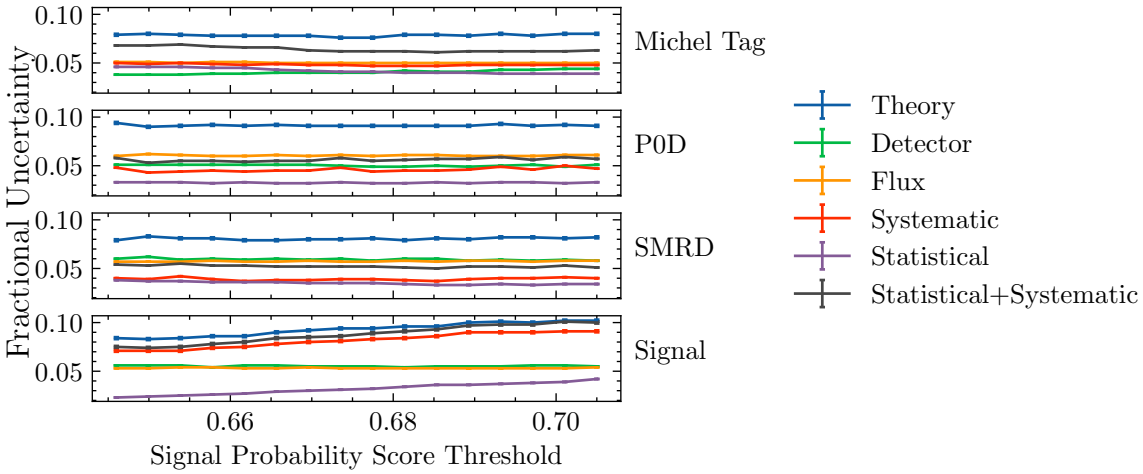

(b) Post-fit.

Fig. 8.8 Fractional uncertainty on the background contamination in the signal sample by source of uncertainty as a function of threshold applied to signal probability score before (Fig. 8.8a) and after (Fig. 8.8b) the sideband fit. Events in the signal sample are real detector data. The uncertainty contributions after the fit are not generally expected to add in quadrature as the fit parameters can become anti-correlated.

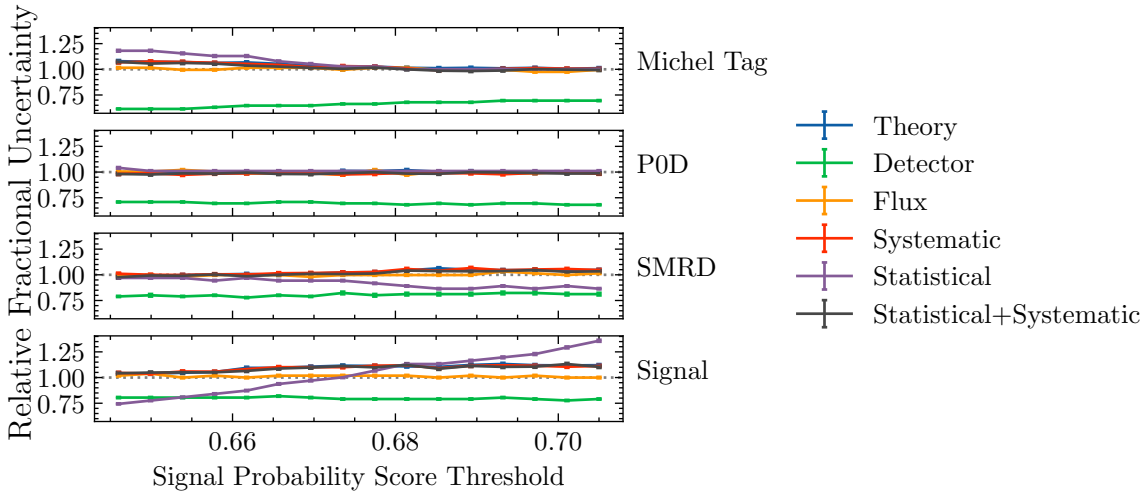

(a) Pre-fit.

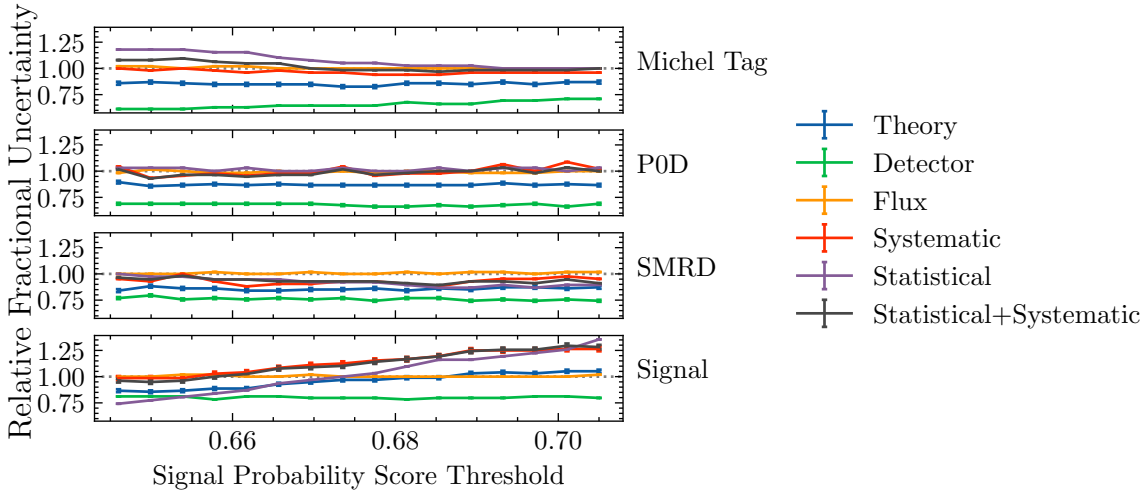

(b) Post-fit.

Fig. 8.9 Relative fractional uncertainty on the background contamination in the signal sample by source of uncertainty as compared to those achieved by the cut-based selection as a function of threshold applied to signal probability score before (Fig. 8.9a) and after (Fig. 8.9b) the sideband fit. Events in the signal sample are real detector data. The uncertainty contributions after the fit are not generally expected to add in quadrature as the fit parameters can become anti-correlated.

classification scores to the nominal dataset except with $\pm 1\sigma$ perturbations applied to their PID values; then, the PID distributions within the perturbed signal samples are compared with those achieved with the dataset with the nominal PID values. It could well be assumed that this reduction in systematic error is a marker of the model's generalisation.

The increased fractional error on the background is not expected to effect an increased uncertainty on the calculated cross section, however. Neglecting the uncertainty on the signal efficiency, flux and number of targets, the uncertainty on the calculated cross section obeys the relation

$$\Delta\sigma \propto \Delta N^{true}_{signal} = \frac{\Delta N^{sel.}_{signal}}{\varepsilon} \tag{8.17}$$

where ΔN^{true}_{signal} is the number of true signal events, $\Delta N^{sel.}_{signal}$ is the number of selected signal events and ε is the signal efficiency. Given that the number of selected events $N^{sel.}$ is given by

$$N^{sel.} = N^{sel.}_{signal} + N^{sel.}_{bkgrd.} \tag{8.18}$$

and assuming that $\Delta N^{sel.} \approx \sqrt{N^{sel.}}$, Eq. (8.17) can be written in terms of the fractional error on the background $\alpha \equiv \Delta N^{sel.}_{bkgrd.}/N^{sel.}_{bkgrd.}$ and the signal purity p and efficiency ε as

$$\Delta\sigma \propto \frac{\sqrt{\frac{\varepsilon N^{true}_{signal}}{p} + \left(\frac{\alpha\varepsilon N^{true}_{signal}(1-p)}{p}\right)^2}}{\varepsilon} \tag{8.19}$$

using the identity $N^{sel.} = \varepsilon N^{true}_{signal}/p$. Assuming that the statistical uncertainty of $N^{sel.}$ is subdominant compared to the uncertainty on the background rate, then Eq. (8.19) can further be approximated as

$$\Delta\sigma \propto \frac{\alpha\,(1-p)}{p}. \tag{8.20}$$

Thus, the uncertainty on the cross section is expected to decrease with the signal probability score threshold x if

$$\frac{\alpha'(x)}{\alpha(x)} < \frac{p(x)}{p(x)'^2\,(1-p(x))}. \tag{8.21}$$

At this stage in the analysis, given the purities and efficiencies of the multivariate signal samples and their derived background rate uncertainties, a number of predictions about the uncertainty on the calculated cross section may be made. For blinded data, the analysis ran

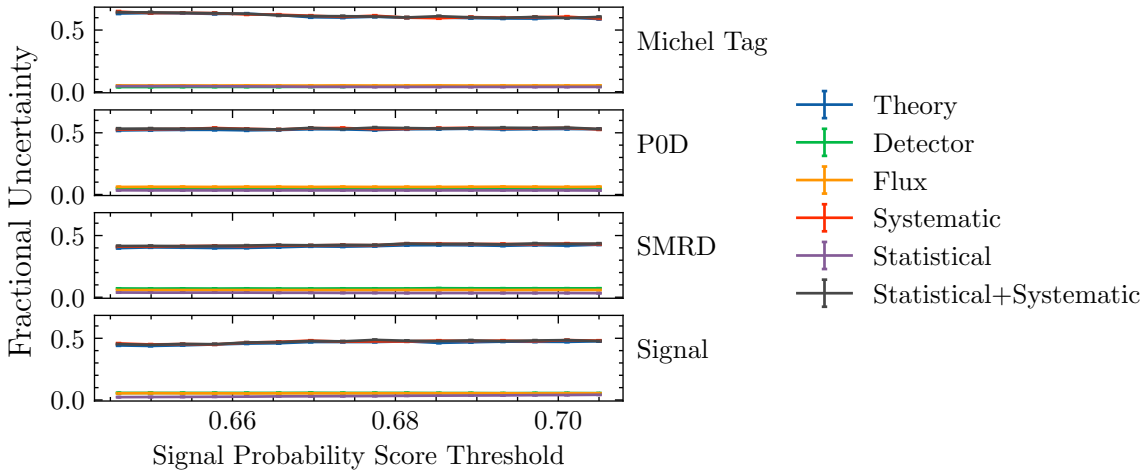

(a) Pre-fit.

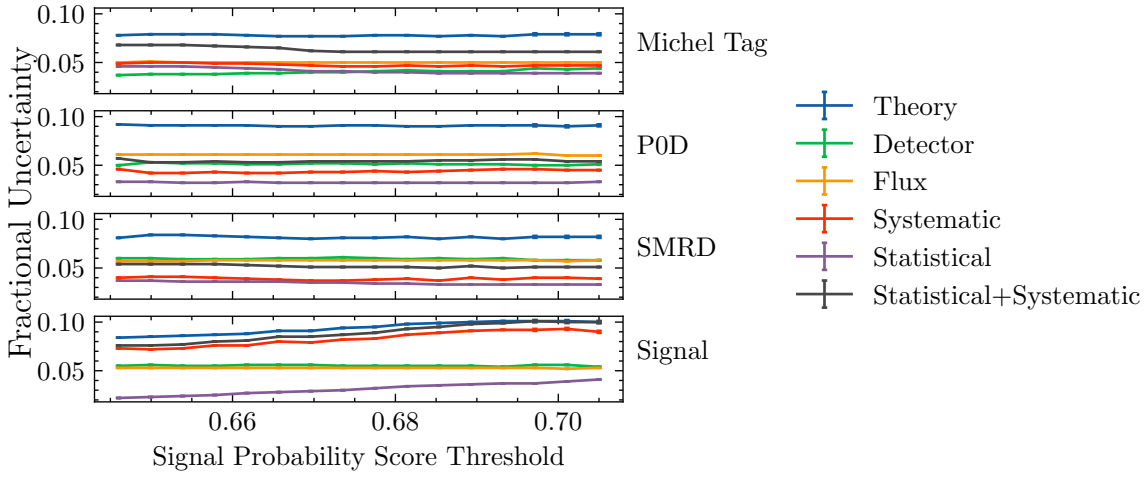

(b) Post-fit.

Fig. 8.10 Fractional uncertainty on the background contamination in the signal sample by source of uncertainty as a function of threshold applied to signal probability score before (Fig. 8.10a) and after (Fig. 8.10b) the sideband fit. Events in the signal sample are substituted for the nominal MC. The uncertainty contributions after the fit are not generally expected to add in quadrature as the fit values can become anti-correlated.

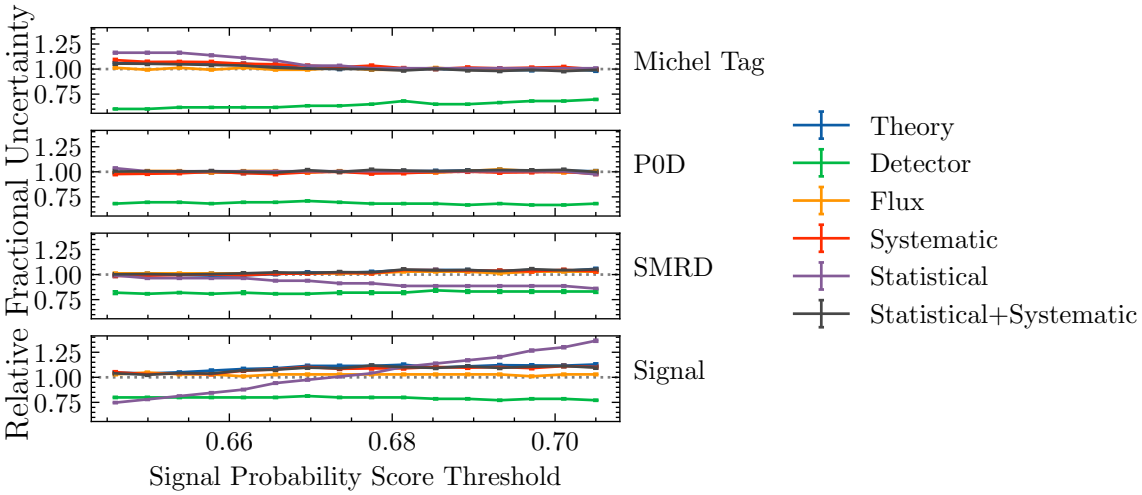

(a) Pre-fit.

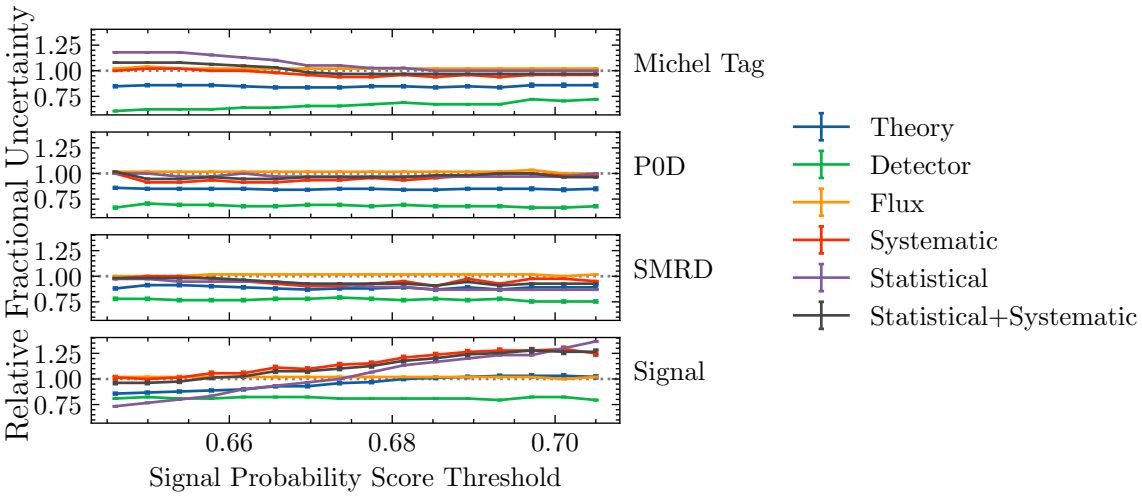

(b) Post-fit.

Fig. 8.11 Relative fractional uncertainty on the background contamination in the signal sample by source of uncertainty as compared to those achieved by the cut-based selection as a function of threshold applied to signal probability score before (Fig. 8.11a) and after (Fig. 8.11b) the sideband fit. Events in the signal sample are substituted for the nominal MC. The uncertainty contributions after the fit are not generally expected to add in quadrature as the fit values can become anti-correlated.

with the multivariate signal samples is expected to achieve, at minimum, a 9% reduction in fractional uncertainty on the calculated cross section relative to that achieved with the cut-based signal sample; at maximum, the reduction is projected to be 31%. For unblinded data, at minimum, a 23% reduction is expected relative to the cut-based signal sample; at maximum, a reduction of 39% is expected. The projections for blinded data are lower than that of unblinded data as the purities of the signal sample tend to also be lower. For both blinded and unblinded data, however, the least precise measurements to be made with the multivariate signal sample is the one which maximises the signal efficiency and matches the purity of the cut-based analysis. The most precise measurement is provided by a sample with a much-enhanced purity, however, not maximised. The 11^{th} signal probability threshold corresponds to a kink in the post-fit purity–efficiency curve, beyond which—in data points 11–14—the purity stays approximately constant but the efficiency drops by 17%. Over these points, the fractional uncertainty on the background also increases by $\sim$5%, driven by an increase in theoretical uncertainties. Thus, the 11^{th} threshold (0.685) is a minimum with respect to the uncertainty contributed to the cross section by the uncertainty on the background rate.

The number of events per event class for all analysis samples after the fit may be studied in Figs. 8.12 and 8.14. It can be seen in Figs. 8.16 and 8.17 that for all event classes, the post-fit event rates are consistent with their pre-fit ones to within one standard deviation with the exception of the signal. The fit appears to compensate for the underestimation of the total number of events in the signal sample and the P0D and SMRD control samples largely by increasing the number of signal events. The fractional uncertainties on the number of extracted events are shown in Figs. 8.13 and 8.15. As expected, the fractional uncertainty on the number of CC events is reduced in the multivariate analysis, given its enhanced rejection power. The number of OOPS NC0π and *other* events are very low ($\lesssim$ 10 events) and thus are subject to large statistical fluctuations.

8.3 Cross-Check of the Out-Of-Fiducial-Volume Neutron Background Constraint with Timing Information

Neutrons produced outside of the fiducial volume necessarily have to traverse some distance before they can interact in FGD1; accordingly, the protons produced by OOFV neutrons arrive in FGD1 around $\sim$5 ns later than those produced by true NC0π events. The sensitivity to this background offered by timing information is used as an independent cross-check of

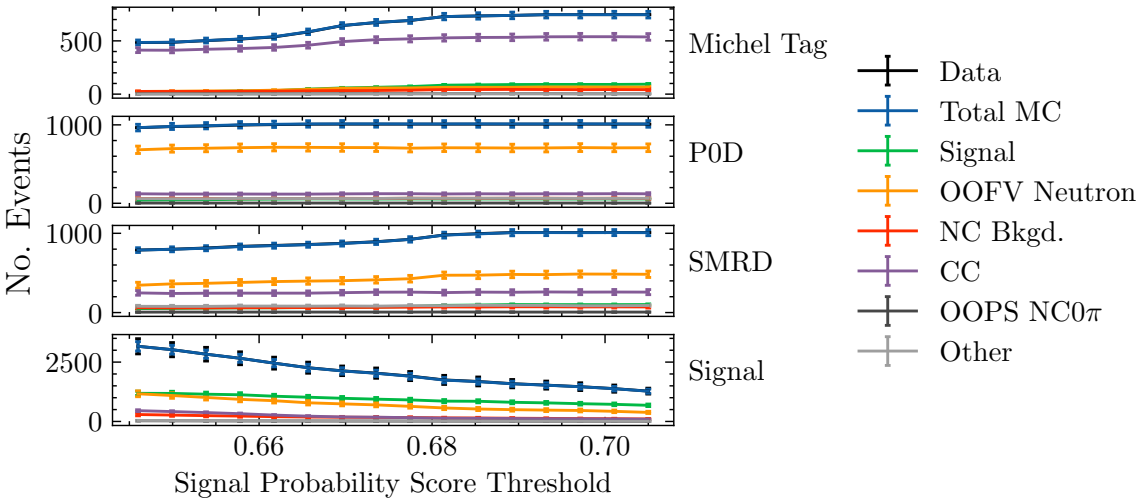

(a) Event rates.

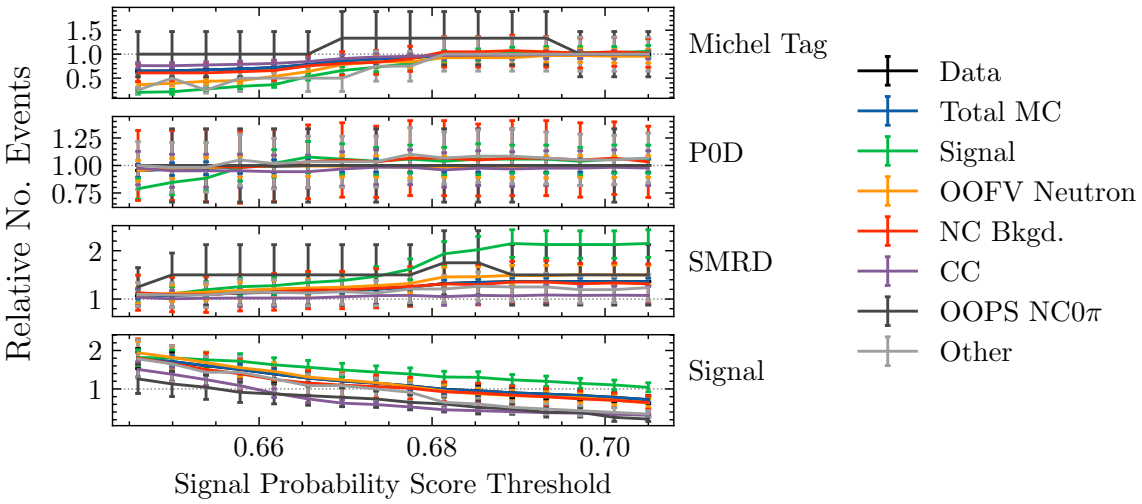

(b) Event rates relative to those achieved with the cut-based selection.

Fig. 8.12 Rates of event class contributions to the signal sample as a function of threshold applied to signal probability score. Events in the signal sample are substituted for the nominal MC.

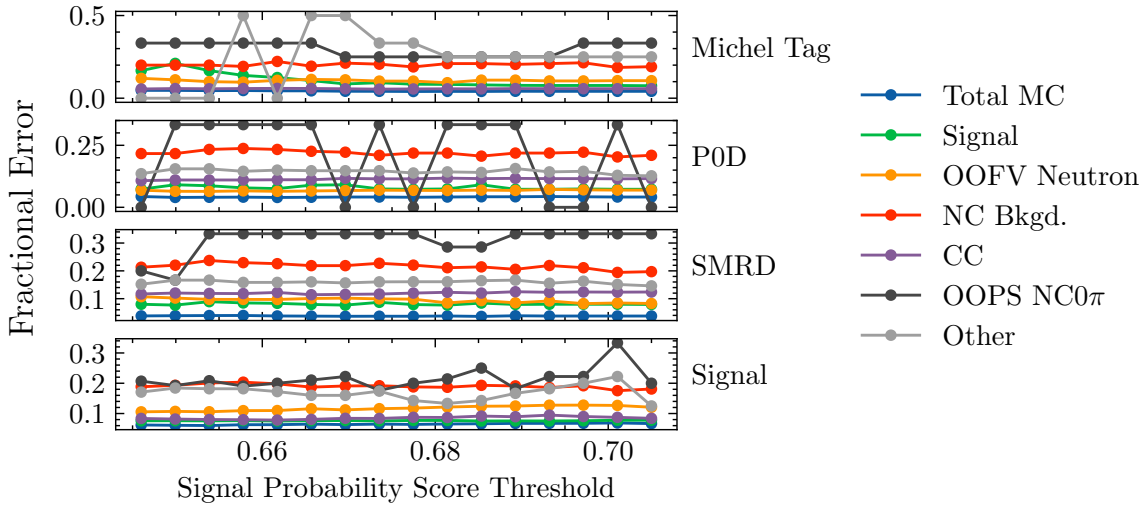

(a) Fractional uncertainties.

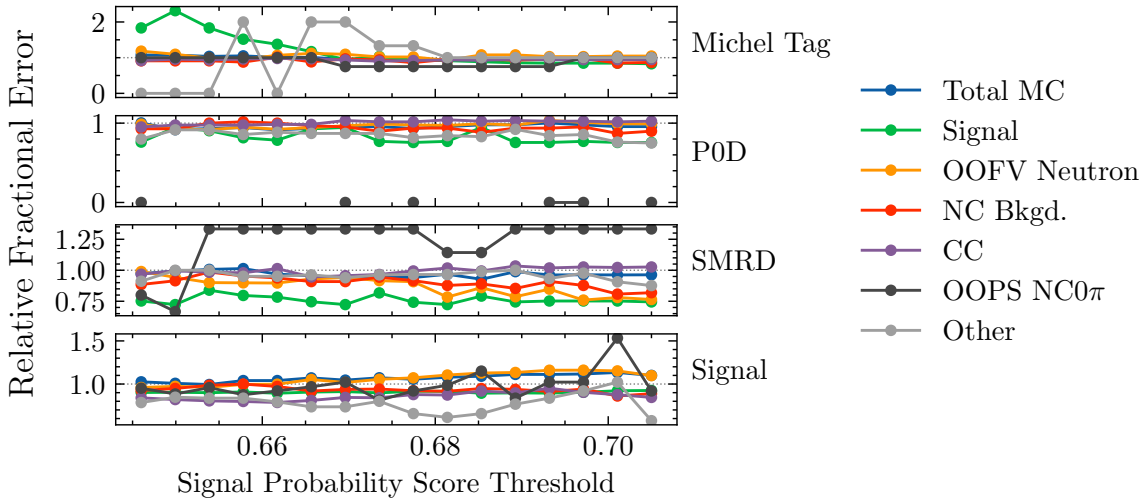

(b) Fractional uncertainties as a fraction of those achieved with the cut-based selection.

Fig. 8.13 Fractional uncertainties on rates of event class contributions to the signal sample. Events in the signal sample are substituted for the nominal MC.

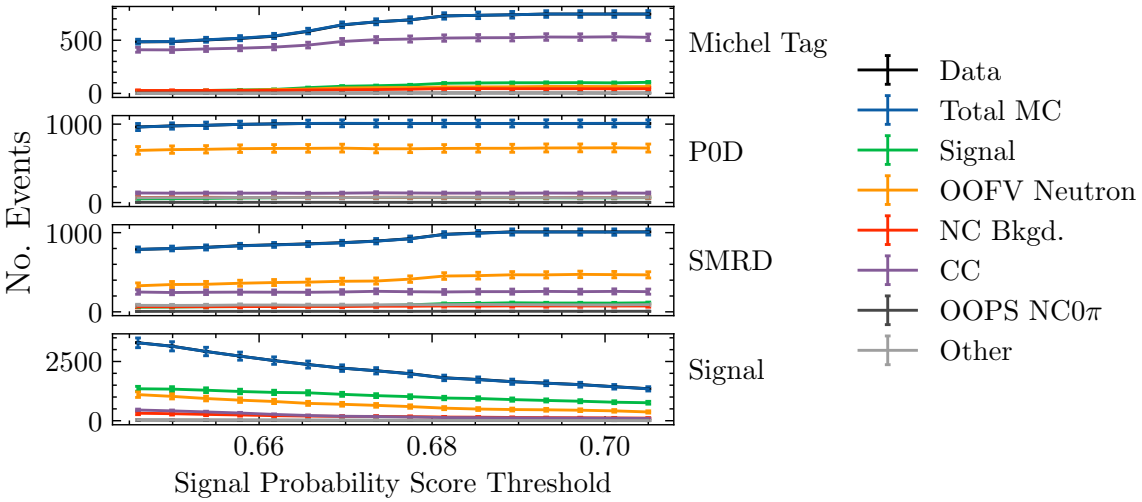

(a) Event rates.

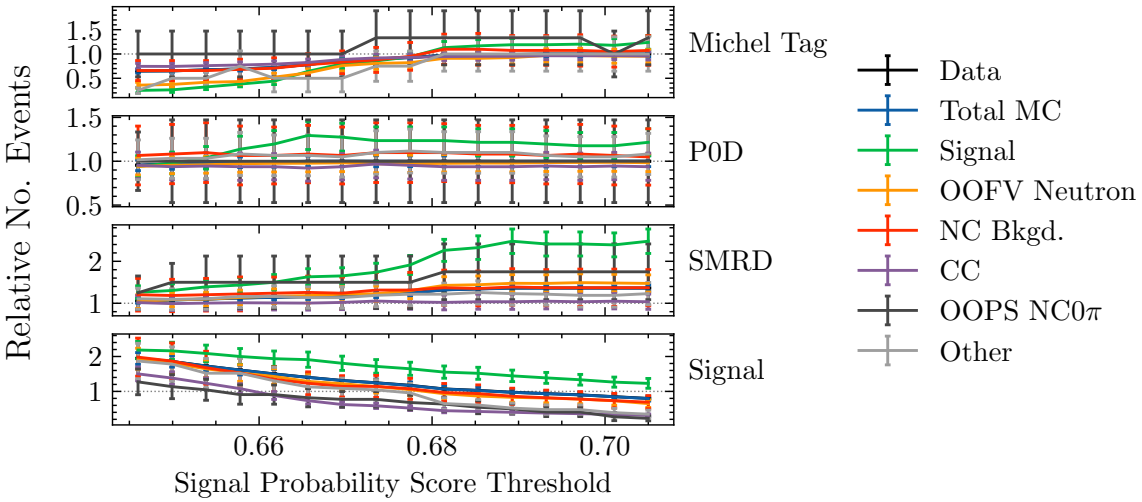

(b) Event rates as a fraction of those achieved with the cut-based selection.

Fig. 8.14 Rates of event class contributions to the signal sample as a function of threshold applied to signal probability score. Events in the signal sample are real detector data.

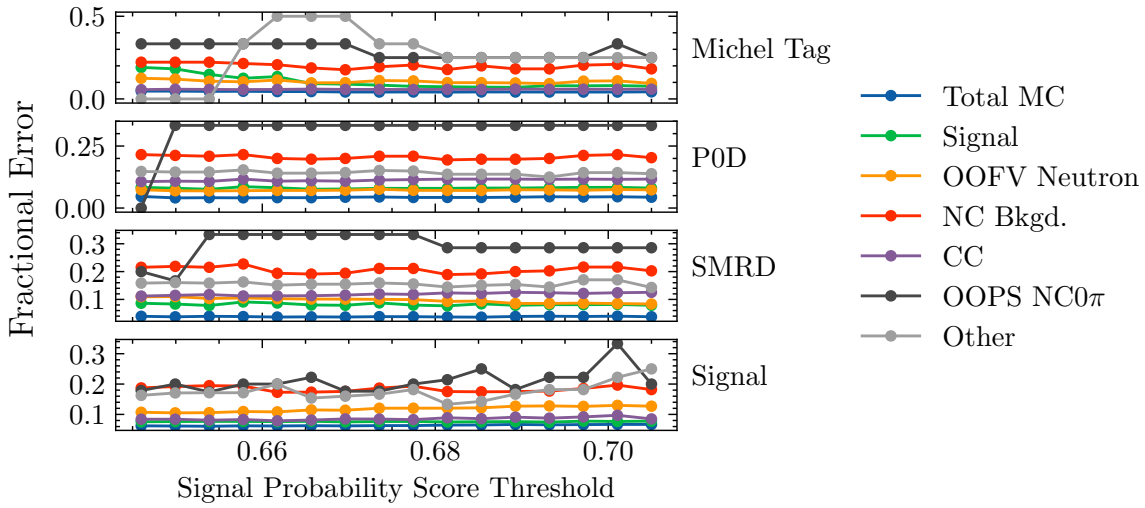

(a) Fractional uncertainties.

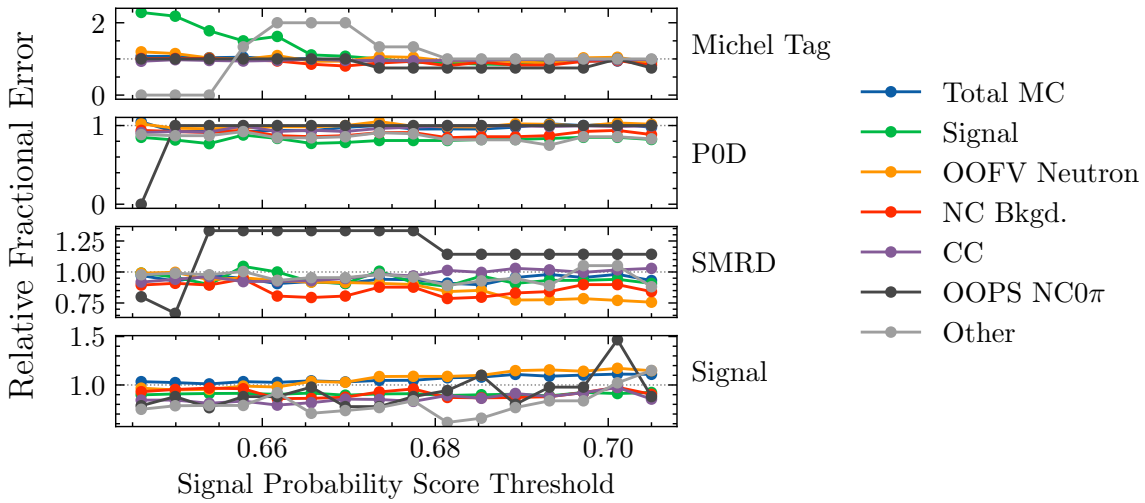

(b) Fractional uncertainties as a fraction of those achieved with the cut-based selection.

Fig. 8.15 Fractional uncertainties on rates of event class contributions to the signal sample as a function of threshold applied to signal probability score. Events in the signal sample are real detector data.

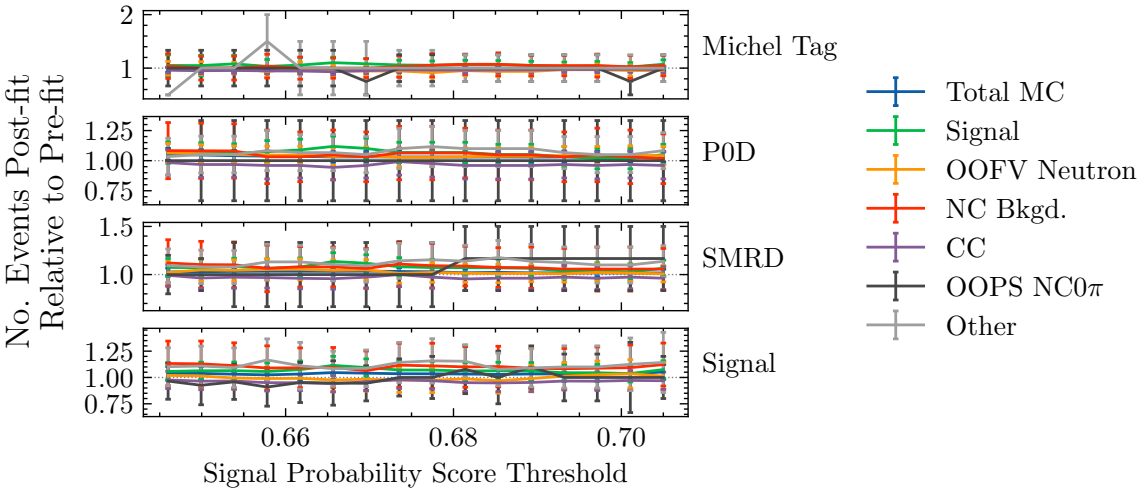

Fig. 8.16 Number of events post-fit relative to pre-fit in the signal sample by event class contribution as a function of signal probability threshold applied. The data in the signal region is real detector data.

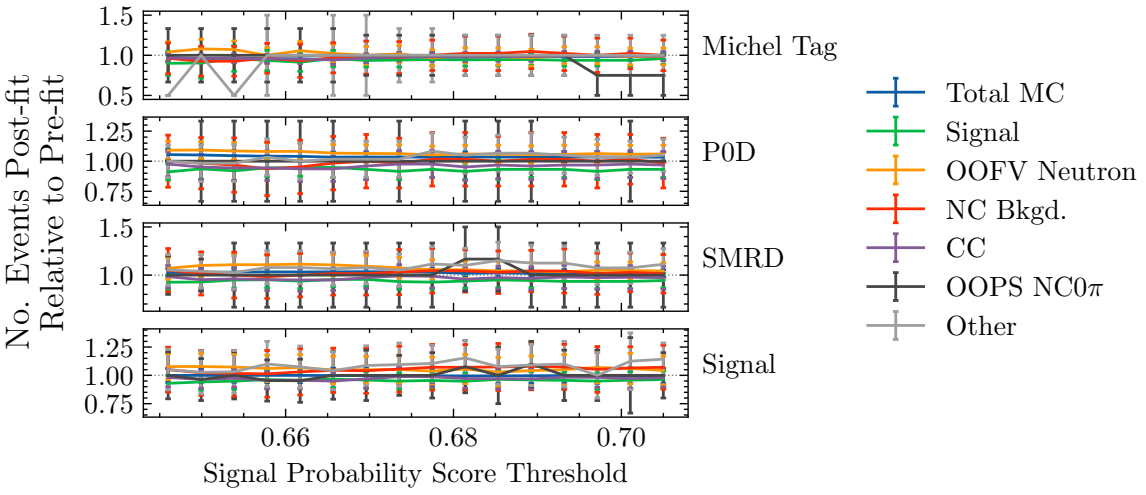

Fig. 8.17 Number of events post-fit relative to pre-fit in the signal sample by event class contribution as a function of signal probability threshold applied. The data in the signal region is substituted for the nominal MC.

the constraint on the OOFV neutron background contribution to the signal sample as derived by the sideband fit.

8.3.1 Timing corrections

The time between a neutrino interaction and the appearance of a proton candidate in FGD1 is probed by the time between the delivery of the proton bunch and the arrival of a proton-like global track in the fiducial volume—this quantity will be called the *relative track time*. There are many monitors deployed along the T2K beamline, of which the five Current Transformer (CTs) monitors are most suited to measure the bunch timing. Though the choice of CT is arbitrary, the pulse times recorded at CT5—the fifth and most downstream CT—have been chosen to represent the bunch delivery time in this study.

Data calibration

In Fig. 8.18 it can be seen that the relative track times in data are subject to the drift of the sub-detector calibrations; in addition, a large jump can be seen during run 2 which corresponds to an increase to the electronic delays. These instabilities are calibrated out at the analysis level by subtracting off the mean value found within graduated periods of wall time—this process is performed separately for runs 2–4 and run 8. In run 8 there is a noticeable loss of timing resolution due to instabilities in the FGD timing markers during that data-taking period. In light of this, the extraction of neutron background rate is attempted only on data taken during runs 2–4.

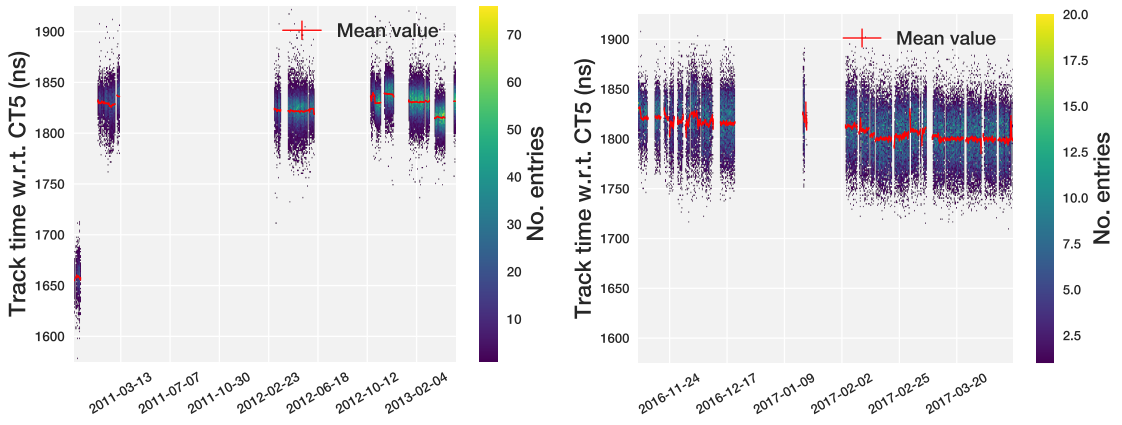

(a) Runs 2–4.

(b) Run 8.

Fig. 8.18 Variation of FGD1 track time relative to the CT5 proton bunch delivery monitor with wall time in data for all proton candidates in the fiducial volume.

Monte-Carlo remodelling

The timing of the neutrino beam is poorly modelled in the nominal near-detector MC—therefore, the relative track times for MC are remodelled on an event-by-event basis at the analysis level using truth information and the shape of the calibrated relative track times in data. The remodelled relative track time for MC is given by

$$t' = t^{reco} - t^{true} + G + z^{true}/c \tag{8.22}$$

where t^{reco} is the reconstructed track time, t^{true} is the track time in truth, G is a smearing term informed by the data shape, z_{true} is the z-axis component of the track position in truth and c is the speed of light.

The smearing term for a given track is a random draw from a probability density function which is, in turn, the result of a fit to the distribution of calibrated relative track times in data. The data shape is noticeably skewed left—a similar shape has been observed by previous analysers [141]. Accordingly, the Double-Gaussian distribution—comprised of two independent one-dimensional Gaussian functions in superposition—was found to describe the data shape sufficiently. The function fitted to data is given by

$$G(t) = \sum_{i=1}^{2} N_i e^{\frac{(t-\mu_i)^2}{2\sigma_i^2}} \tag{8.23}$$

where i indexes the subfunction, N_i is the normalisation, μ_i is the centroid position and σ_i is the width. The fitted parameters are given in Table 8.3.

Table 8.3 Post-fit parameter values of the Double-Gaussian fit to the corrected RTOF distribution for all tracks starting in the fiducial volume in data.

(a) Runs 2–4

Parameter	Value
μ_1	1.482 ± 0.068
μ_2	-3.42 ± 0.18
N_1	2565 ± 32
N_2	662 ± 33
σ_1	8.690 ± 0.080
σ_2	16.36 ± 0.19

(b) Run 8

Parameter	Value
μ_1	19.26 ± 0.41
μ_2	-3.65 ± 0.53
N_1	402 ± 30
N_2	1201 ± 20
σ_1	11.57 ± 0.65
σ_2	23.37 ± 0.22

Resolution effects are double-counted in the MC correction: once by the reconstructed supplied by the nominal MC, and once by the smearing term. To mitigate this, the mean and

the standard deviation of the MC distribution are matched to those of the data. As shown in Fig. 8.19, the relative track times given by data and MC are placed in good agreement by these corrections.

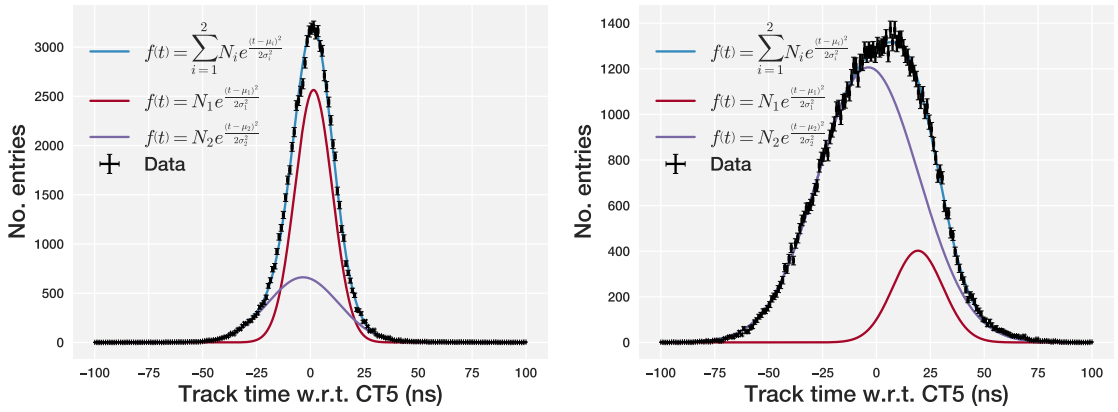

(a) Runs 2–4.

(b) Run 8.

Fig. 8.19 Distribution of calibrated FGD1 track times relative to the CT5 proton bunch delivery monitor for all proton candidates in the fiducial volume. The distributions are shown separately for runs 2–4 (Fig. 8.19a) and run 8 (Fig. 8.19b).

8.3.2 Background rate extraction

The OOFV neutron rates in the signal sample is measured by a fit to the distribution of relative track times in the cut-based signal sample. The fit function is informed by MC, consisting of the shapes of the relative track time distributions for OOFV neutrons and the remainder of the sample; the normalisations of these two shapes are the only free parameters in the fit. This approach was first tested upon the MC signal sample, where an OOFV neutron rate of 1.00 ± 0.17 relative to the true rate was measured. Performed upon the data sample, a ratio of 0.95 ± 0.17 was measured. These fits are depicted in Fig. 8.21. This result is in good agreement with the same quantity found independently by the fit to the sidebands, which is 1.04 ± 0.11.

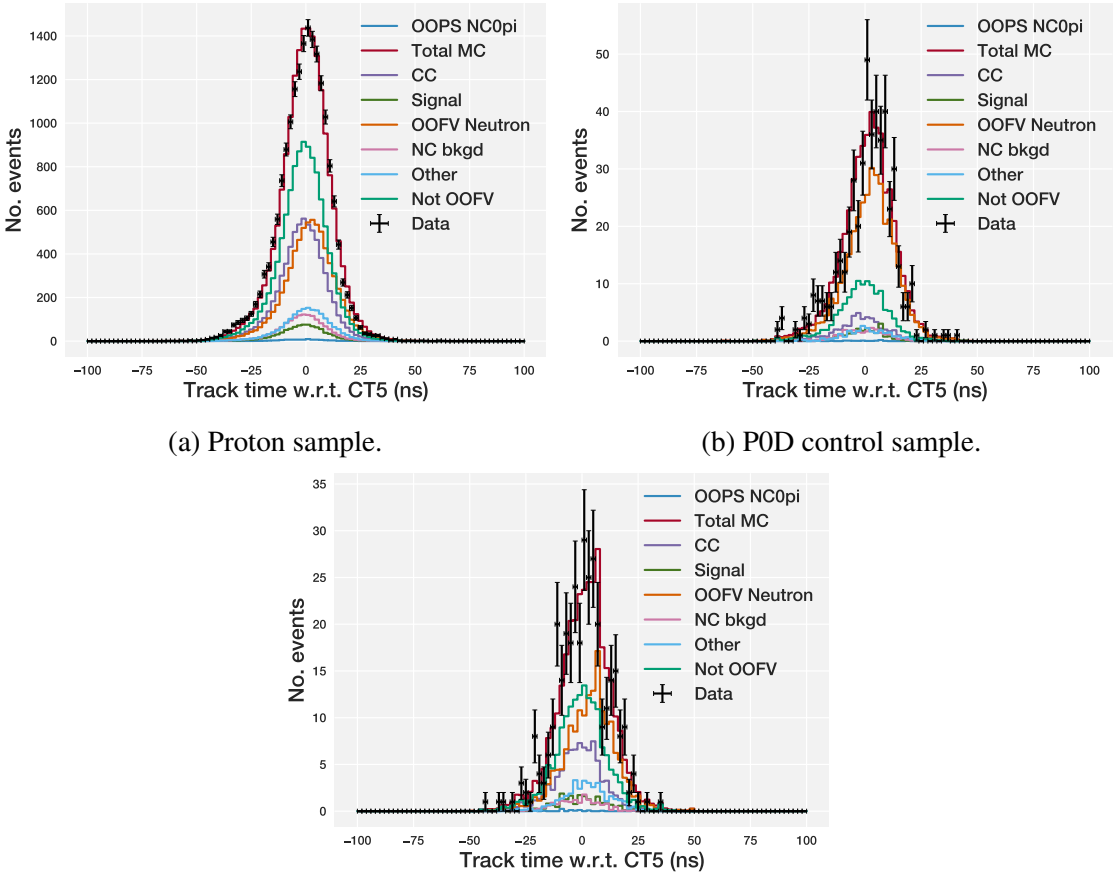

(a) Proton sample.

(b) P0D control sample.

(c) SMRD control sample.

Fig. 8.20 Distribution of calibrated FGD1 track times relative to the CT5 proton bunch delivery monitor for all proton candidates in the fiducial volume. The distributions are shown separately for the proton sample (Fig. 8.20a) and the POD (Fig. 8.20b) and the SMRD (Fig. 8.20c) control samples.

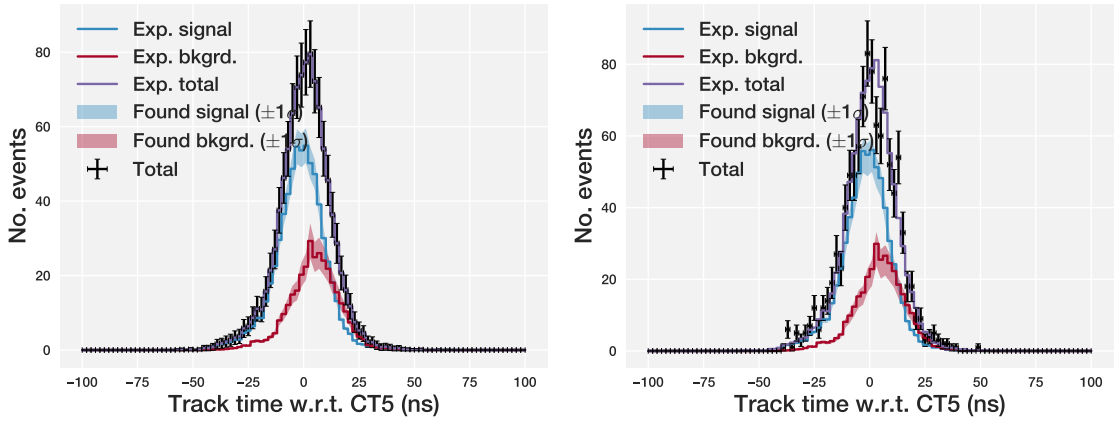

(a) Nominal MC.

(b) Real data.

Fig. 8.21 Extraction of the OOFV neutron background rate in the fiducial volume using FGD track times. The curves labelled *background* are OOFV neutron interactions, the curves labelled *signal* are the remainder. The signal and background template histograms are fit simultaneously to the nominal MC in Fig. 8.21a; in Fig. 8.21b, they are fit to the real data.

8.4 Unfolding Procedure

Any observable quantity is subject to the inefficiencies, thresholds and resolutions of the instrumentation used; to correct for these limitations is the task of *unfolding*. All detector effects—i.e. smearing and efficiencies—are encoded in the *response matrix*. The response matrix is constructed of the joint distribution of simulation data—preferably of high statistics and generated from a variety of interaction models—at the reconstructed and truth-level. Put formally, the response matrix represents the conditional probabilities that an event is reconstructed in bin with index i given that the event belongs to true bin of index j—element ij of the detector matrix T is calculated as

$$T_{ij} = P(i \mid j) = \frac{P(i \cap j)}{P(j)} = \frac{N_{ij}^{selected}}{N_j^{true}} \tag{8.24}$$

where $N_{ij}^{selected}$ is the number of selected true signal events belonging to reconstructed bin i and truth bin j, N_j^{true} is the number of true signal events belonging to truth bin j.

The truth and reconstructed bins of the response matrix are by no means required to be identical; perhaps to the benefit of measurements that suffer from low data statistics, the binning scheme is flexible. However, if discrimination between models is an objective of the measurement, special care must be taken to avoid infecting the response matrix with model dependencies. It is, for this reason, preferable to perform the unfolding with respect

to as many variables that the detector response is dependent on simultaneously in order to minimise the analysis' bias toward a particular interaction model. Care must also be taken to minimise the statistical uncertainty on the elements of the response matrix; the particularly improbable regions in the phase space of the unfolding variables are likely to be starved of MC events and will ultimately contribute to the uncertainty introduced by the unfolding procedure. Ideally, this is diminished by constructing the response matrix with an MC dataset of arbitrarily high statistics and covering a sufficiently wide phase space.

The unfolding is chosen to be performed with respect to the reconstructed kinetic energy and polar angle of the highest momentum proton candidate. Proton kinetic energy is certainly an observable that does particularly depend on the choice of interaction model; later in the analysis, however, is shown that measurements of the differential cross section are dominated by statistical uncertainties—lending little power to discriminate models. Therefore, it is deemed that the model dependency carried by the response matrix is negligible as the analysis currently stands. The binning scheme at both the truth and reconstructed level are identical—25 bins index 5 kinetic energy intervals and 5 polar angle intervals. The kinetic energy binning is the same with which the sideband fit was performed, the boundaries of these bins are

$$[0.0, 181.71, 204.01, 310.56, 427.14, 1000.00]\,\text{MeV}. \tag{8.25}$$

The boundaries of the polar angle bins are

$$[0.00000, 0.09199, 0.13595, 0.17729, 0.22904, 1.00000]\,\pi\text{rad}. \tag{8.26}$$

The number of true signal events is extracted and unfolded simultaneously in a maximum log-likelihood fit of the nominal truth-level Monte-Carlo spectra to the background-subtracted data spectra. The total likelihood to be minimised is given by

$$-2\ln\mathscr{L} = (\vec{N}^{\text{pred.}} - \vec{N}^{\text{obs.}})^T V^{-1} (\vec{N}^{\text{pred.}} - \vec{N}^{\text{obs.}}) \tag{8.27}$$

$$+ p_{E_k} \sum_i \sum_j (c_{ij} - c_{(i+1)j})^2 \tag{8.28}$$

$$+ p_\theta \sum_i \sum_j \left(c_{ij} - c_{i(j+1)}\right)^2 \tag{8.29}$$

where the first term penalises the bin-to-bin discrepancy between the number of smeared truth events $N^{\text{pred.}}$ and the number of background-subtracted data events $N^{\text{obs.}}$, where V is the covariance matrix which functions to propagate through the systematic uncertainties; the two remaining terms are the regularisation penalty terms, where p_{E_k} and p_θ are the regularisation strengths in each dimension. The parameters in this fit are a set of weights c assigned to each truth bin which scale the nominal cross section in that bin. These weights function as corrections to the nominal truth spectra to cover bin-to-bin discrepancies between the truth spectra smeared by the detector matrix and the data. The predicted spectra $\vec{N}^{\text{pred.}}$ is the truth data smeared by the response matrix, calculated as

$$N_{ij}^{\text{pred.}} = \sum_k \sum_l c_{ij} N_{ij}^{\text{truth}} T_{ijkl}, \tag{8.30}$$

where i and j index the reconstructed bins in proton kinetic energy and polar angle, k and j index the truth bins in proton kinetic energy and polar angle, N_{ij}^{truth} is the number of events in truth bins ij and T_{ijkl} is the element of the response matrix describing reconstructed bins ij and truth bins kl.

The regularisation terms exist to address a particular problem with this approach when faced with noisy data, in that there often exists a multitude of viable parameter sets found by the fit—in other words, the problem becomes *ill-posed*. One approach to overcoming this issue of degeneracy is regularisation, whereby, given prior knowledge that cross section values vary smoothly from bin to bin when the bins are sufficiently fine, the spectra weights are also encouraged to vary smoothly. The regularisation is thus implemented as a penalty term in the likelihood function that increases with the difference between weights in adjacent bins. The amount that the penalty term contributes to the total likelihood is determined by a positive constant to be determined known as the *regularisation strength*. At sufficiently low strengths, the problem of degeneracy persists, and adjacent weights are free to become severely anti-correlated; at sufficiently high strengths, the weighted smeared spectra will rather too closely resemble the unweighted smeared truth data rather than the reconstructed spectra. Ideally, the optimal value of the regularisation strength will compromise between these two extremes.

A common approach to determine the ideal regularisation strength is by an *L-curve* analysis [142]. The L-curve is a plot of the total minimised log-likelihood against the corresponding log-likelihood contributed by the regularisation penalty term for a range of applied regularisation strengths. The optimal regularisation strength is said to be the one at the "corner" or "kink" of the curve—the value which maximises the curvature of the L-curve. For each unfolded dimension, there is a corresponding regularisation strength to

be determined; hence, in this analysis, there is p_{E_k} and p_θ that each regularise the fitted weights with respect to the proton's kinetic energy and polar angle. Each strength is found by performing an isolated one-dimensional L-curve analysis—i.e. with the value of the other regularisation strength set at zero. Some typical L-curves are shown in Fig. 8.22.

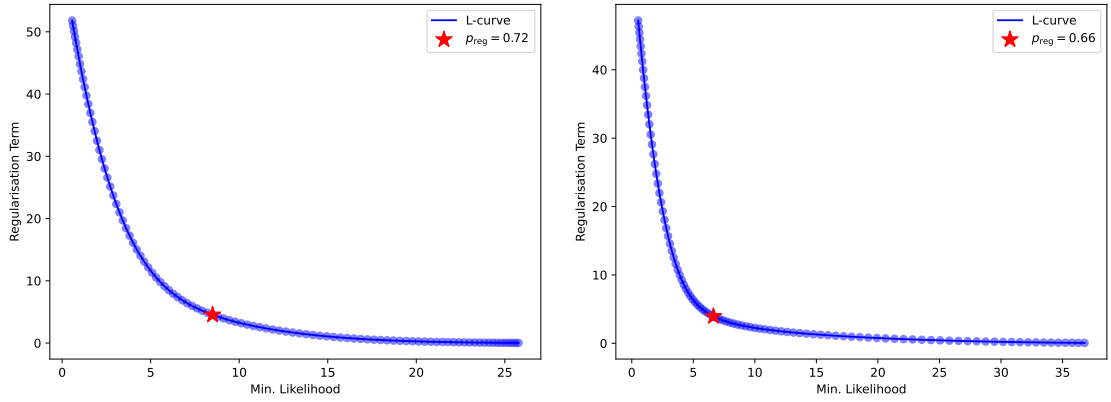

(a) Regularisation to proton kinetic energy.

(b) Regularisation to proton polar angle.

Fig. 8.22 Regularisation likelihood penalty contribution against the total minimised log-likelihood for a range of trial regularisation strengths found for the NC0π unfolding fit to the detector data. Figs. 8.22a and 8.22b show the likelihood penalties contributed by regularisations to proton kinetic energy and proton angle respectively; in each sub-figure, the point of maximum curvature is indicated by the red star.

The best-fit values of the weights c_{ij} are shown for the nominal MC and detector data in Figs. 8.23 and 8.24. The weights are applied to the spectra induced by all true signal events to uncover the unfolded spectra.

The regularisation strengths derived for the blinded analysis can become very large (~ 50). This is because the observed and expected spectra—both derived from the nominal MC—are already fairly consistent prior to the fit, therefore the fitted weights will all be close to unity and the regularisation strength will have to become large in order for the regularisation penalty term to contribute a non-negligible amount to the likelihood. For the unblinded analysis, the observed distributions are subject to significant statistical noise; therefore, the regularisation strengths can be fairly low. There is some reason to believe that the regularisation strengths are being overestimated with the current one-dimensional L-curve method. The unfolding procedure was earlier performed in 10 kinetic energy bins and 10 polar angle bins rather than 5 each, in which the bin values for data were subject to large statistical fluctuations. The unfolded data was seen to be biased to the truth MC, ultimately effecting a bias in the total number of unfolded events. Both regularisation terms contribute to the likelihood simultaneously, therefore the optimisation procedure should

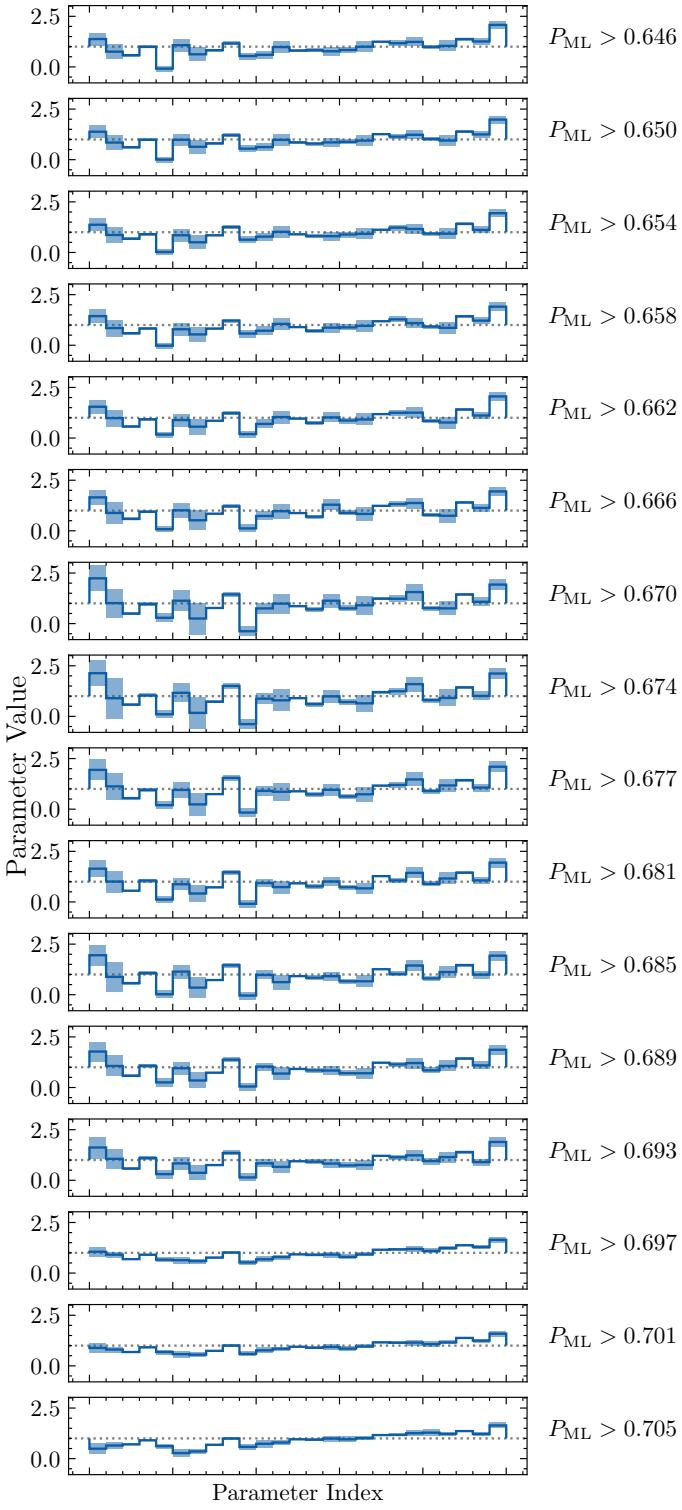

Fig. 8.23 Best-fit values of the weights used to unfold the detector data.

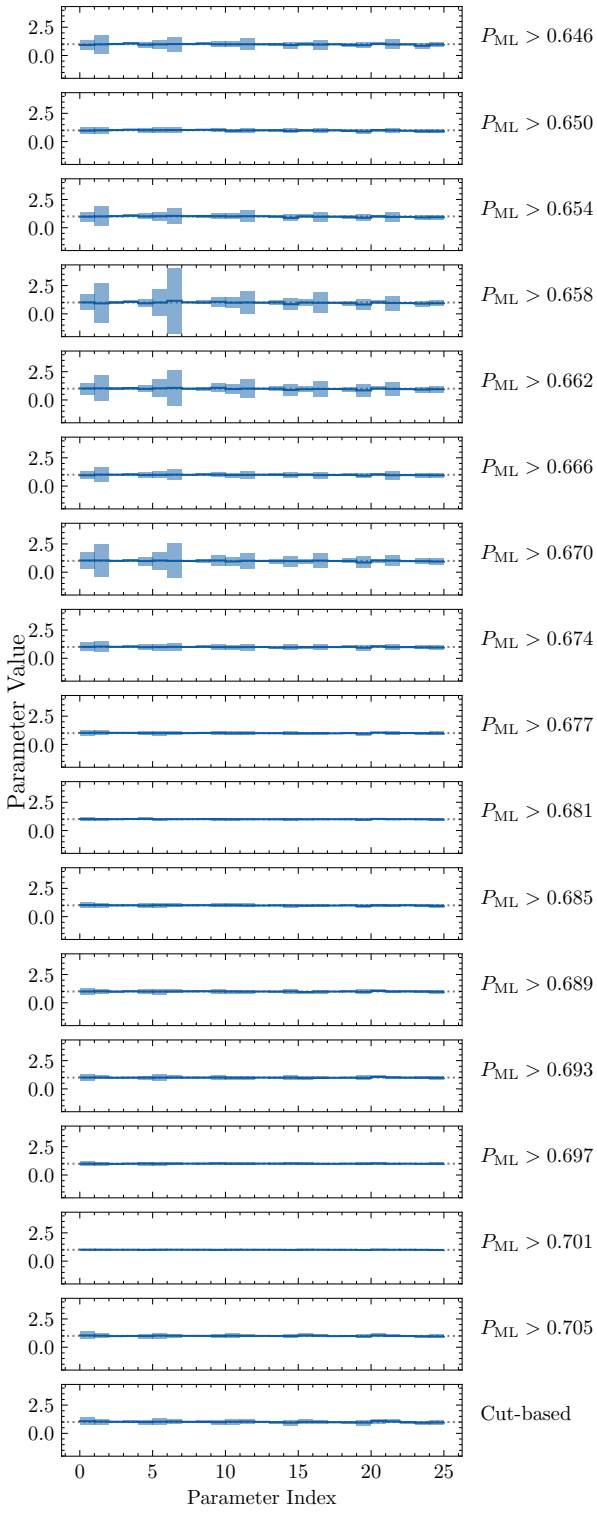

Fig. 8.24 Best-fit values of the weights used to unfold the nominal MC.

optimise both regularisation strengths simultaneously. One obvious approach, for future analysers to consider, is to instead conduct a two-dimensional L-curve analysis, maximising the Gaussian curvature on the minimum likelihood surface with respect to the two penalty terms. Otherwise, other regularisation methods could also be considered.

8.5 Cross Section Extraction

The cross section σ per unit flux per nucleon measured in kinetic energy bin i and polar angle bin j is calculated as

$$\sigma_{ij} = \frac{N_{ij}}{T\phi_{ij}} \tag{8.31}$$

where N_{ij} is the number of efficiency corrected signal events in bin ij, T is the number of nuclei in the target and ϕ_{ij} is the neutrino flux in bin ij. The number of nucleons within the fiducial volume is taken from Ref. [115] as 5.50×10^{29} with an uncertainty of 0.67% [143]—the FGD is composed of 86.1% carbon, 7.35% hydrogen and 3.70% oxygen [115]. The flux values are taken from the fit to the sidebands.

The measured cross sections as averaged over all flux bins are presented for the blinded analysis and unblinded analysis in Figs. 8.25 and 8.26. In the blinded analysis, the cross sections measured in the multivariate analyses and the cut-based analysis are closely aligned. All of these values are consistent with the cross section predicted by the nominal GENIE model to within 1σ uncertainties, thus validating the analysis framework. The fractional uncertainties of the multivariate analyses are all smaller than that of the cut-based analysis. The uncertainty is 14.7% for the cut-based analysis; the fractional uncertainties for the multivariate analyses may be seen in Fig. 8.28a. At maximum, the fractional uncertainty is 14.6%—a 3% reduction compared to the cut-based analysis; at minimum, the fractional uncertainty is 10.7%—a 29% reduction compared to the cut-based analysis. These results align with their expectations provided earlier in Section 8.2. Either way, the precision of the measurement grants a negligible power to discriminate models—therefore, it should be concluded that any bias to any particular model introduced by the unfolding procedure has a negligible effect on the unfolded cross section. When it comes to real detector data, the values found with multivariate signal samples are increased $\sim$10% ($\sim 1\sigma$) compared to those measured of the blinded data. At root, this is caused by an excess of data for the multivariate samples. For the blinded multivariate samples and both the unblinded and blinded cut-based samples, the fit to the sidebands both decreases the signal rate and increases the background rates by less than 5% each. For the unblinded multivariate samples, the fit to the sidebands

corrects for this excess by boosting the signal rate by around ∼5%. The extracted cross sections of the unblinded multivariate samples still, however, lie in good agreement with the value measured with the cut-based signal sample and their corresponding blinded samples. This bias is nonetheless large enough to effect a small tension ($\sim 2\sigma$) with the relatively lower cross section predicted by NuWro when the strange quark contribution to the proton spin Δs is set to $+0.15$. The values achieved with the cut-based signal sample are consistent to ∼1% whether the data is blinded or unblinded. The cross sections measured using the multivariate signal samples are stable with respect to the signal probability threshold applied as expected. The fractional uncertainties for the multivariate analyses may be seen in Fig. 8.27a. For the multivariate samples, the fractional uncertainty is 12.7% at maximum; at minimum, the fractional uncertainty is 10.8%—a 28% reduction relative to that achieved with the cut-based signal sample.

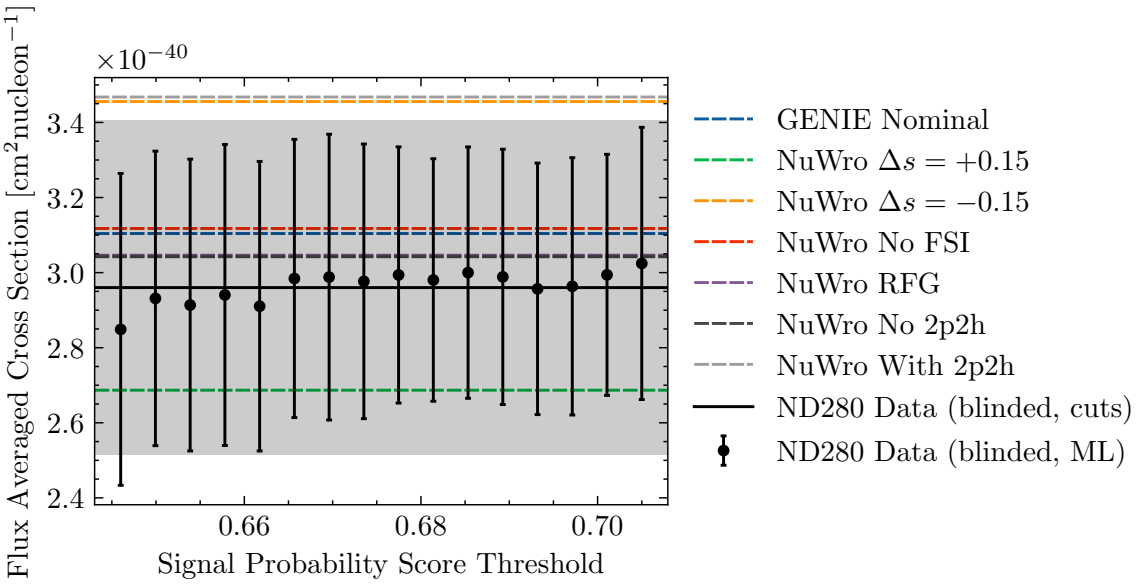

Fig. 8.25 Values of the flux-averaged cross section as a function of signal probability threshold. Events in the signal region are substituted for the nominal MC.

A breakdown of the fractional uncertainties as a function of the signal probability score threshold are presented for the blinded and unblinded data in Figs. 8.27a and 8.28a and compared to those achieved with the cut-based in Figs. 8.27b and 8.28b. The uncertainty contributions were calculated by conducting a high number of pseudo-experiments that randomise the parameters of the cross section calculation in appropriate ways; hence, the uncertainty on the uncertainty contributions is the statistical error due to the number of pseudo-experiments conducted. These contributions are the uncertainty associated with the number of background events in the signal sample, the statistical uncertainty associated

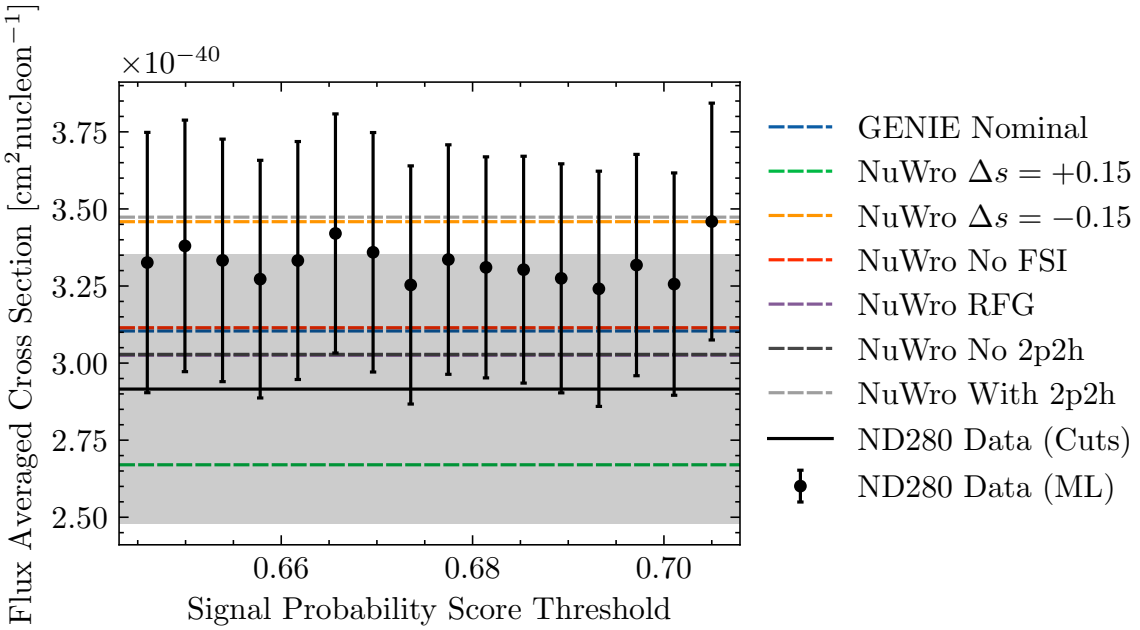

Fig. 8.26 Values of the flux-averaged cross section as a function of signal probability threshold. Events in the signal region are real detector data.

with the number of data events in the signal sample, the uncertainty on the predicted flux, the statistical uncertainty associated with the elements of the forward-folding matrix, and the uncertainty on the number of target nuclei in the FGD. The reduction in uncertainty on the flux-averaged cross section compared to that of the cut-based analysis is driven by a reduction in the uncertainty on the background in the signal sample, which remains the dominant uncertainty in the measurement. As expected, Section 8.2, there exists a minima in the background uncertainty contribution, and hence the total uncertainty, close to the 10^{th} threshold. The statistical error on the selected data events is expected to increase as $\sqrt{p/\varepsilon}$, where p and ε are the purity and efficiency of the signal sample. Deviations from this trend can be attributed to statistical fluctuations having run a limited number of pseudo-experiments. The uncertainty due to the unfolding procedure is driven by the statistical uncertainty of the MC sample used to generate the forward-folding matrix. Since the forward-folding matrix is drawn from true signal events within the signal sample, the uncertainty due to the unfolding decreases with the signal efficiency. Some fluctuations are expected with increasing signal purity, however, due to the treatment of uncertainty of empty elements of the forward-folding matrix. As the phase space with respect to the proton kinetic energy and angle at both the reconstructed and truth levels of events within the signal sample becomes more restrictive with increasing signal purity, some elements of the forward-folding matrix become unfilled. These empty elements can only be assigned an uncertainty of zero; hence, elements that

eventually become unfilled contribute an increasing amount to the total unfolding uncertainty as the statistics of that element diminishes, then suddenly zero when unfilled. Though subdominant, the unfolding uncertainty can easily be decreased with a larger sample of signal MC. The uncertainties on the number of target nuclei and flux, which are subdominant and reside at the 0.7% and 4.8% levels, are invariant to the composition of the signal sample as expected.

The most precise flux-averaged cross section measurements achieved at ND280 are given alongside various interaction model predictions in Table 8.4. A breakdown of the uncertainty contributions to the most precise measurement using the multivariate signal sample is shown in Table 8.5.

Model	$\langle\sigma\rangle_\phi$ [10^{-40} cm^2 nucleon^{-1}]
ND280 Data (Best ML)	3.31 ± 0.36
ND280 Data (Cuts)	2.92 ± 0.44
ND280 Data (Blinded, best ML)	2.99 ± 0.32
ND280 Data (Blinded, cuts)	2.96 ± 0.44
GENIE Nominal	3.10
NuWro $\Delta s = +0.15$	2.67
NuWro $\Delta s = -0.15$	3.46
NuWro No FSI	3.11
NuWro RFG	3.03
NuWro No 2p2h	3.03
NuWro With 2p2h	3.47

Table 8.4 Flux-averaged cross sections as measured at ND280 and predicted by various interaction models.

Source	Fractional Uncertainty [%]
Background	9.27 ± 0.21
Data Statistics	5.71 ± 0.13
Flux	4.94 ± 0.11
Unfolding	2.51 ± 0.06
Num. Target Nuclei	0.66 ± 0.01
Total	10.28 ± 0.23

Table 8.5 Fractional uncertainty contributions to the best measured flux integrated cross section value from detector data. The uncertainty on the fractional uncertainty is a statistical uncertainty due to the number of pseudo-experiments.

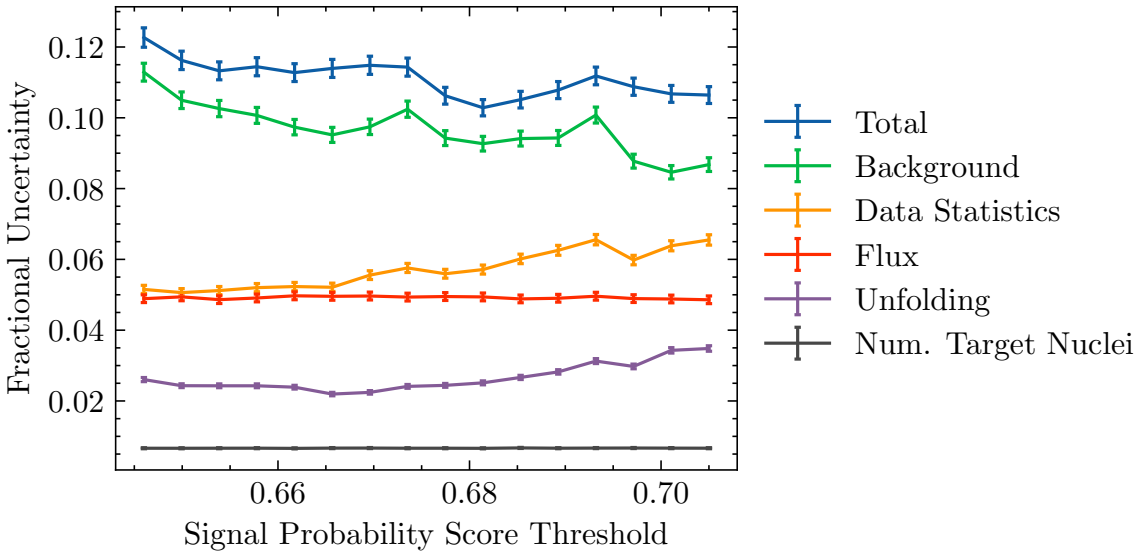

(a) Fractional uncertainties.

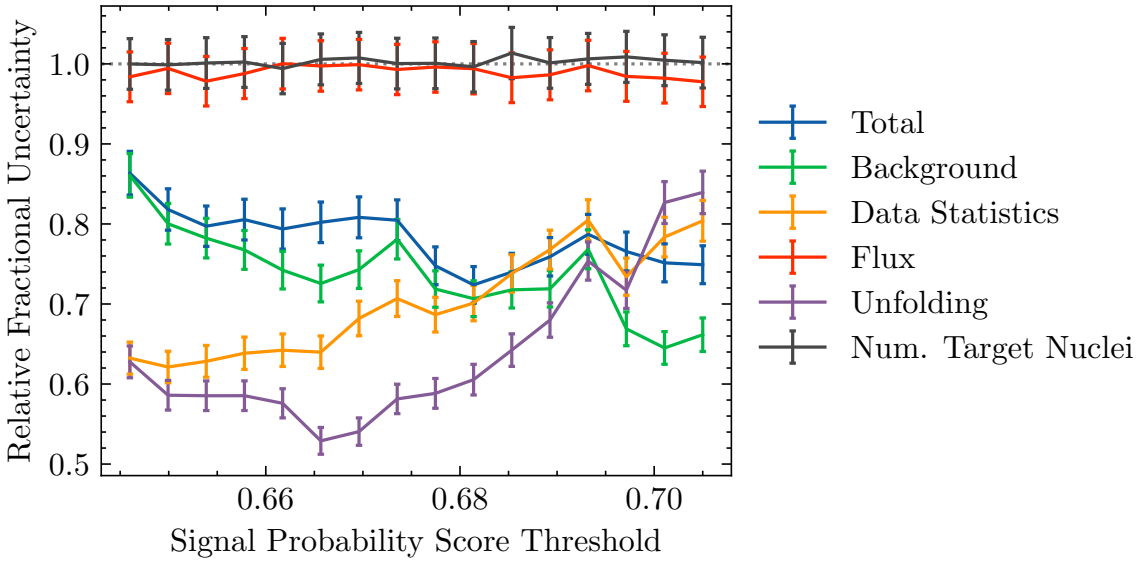

(b) Fractional uncertainties as a fraction of those achieved with the cut-based selection.

Fig. 8.27 Contributions to the fractional error of the flux-averaged cross section as a function of signal probability score threshold. Events in the signal region are real detector data.

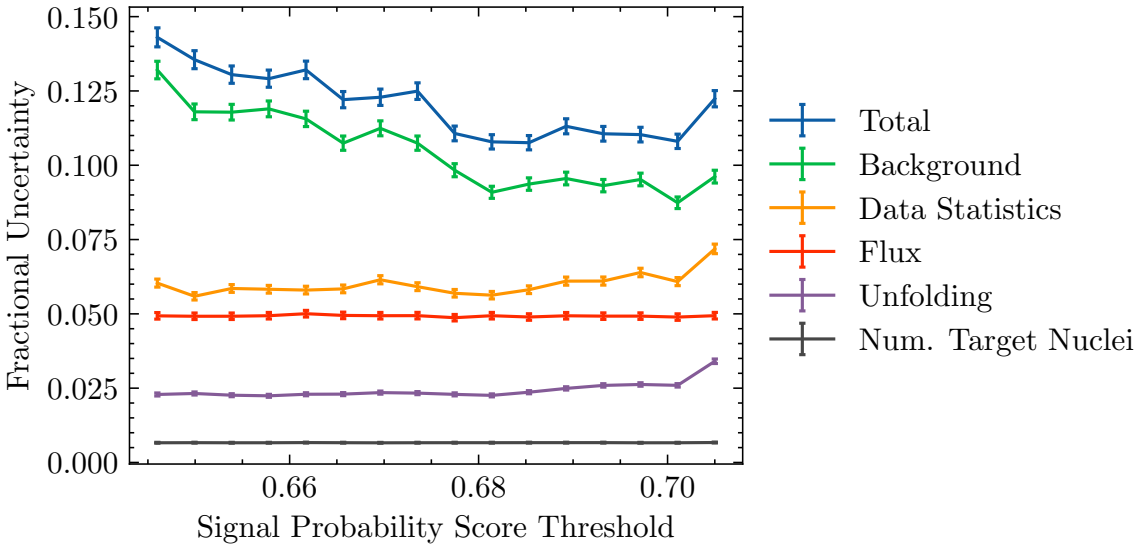

(a) Fractional uncertainties.

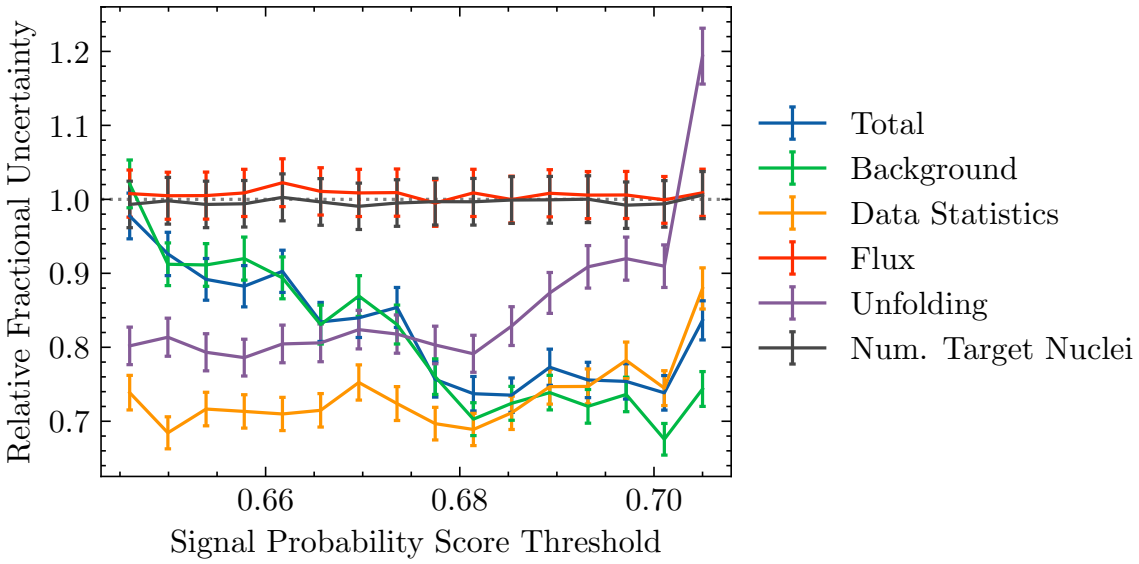

(b) Fractional uncertainties as a fraction of those achieved with the cut-based selection.

Fig. 8.28 Contributions to the fractional error of the flux-averaged cross section of signal probability score threshold. Events in the signal region are substituted for the nominal MC.

The data is unfolded in two dimensions, therefore, in principle, the double-differential cross section may be presented. The measurement is statistically limited, however, hence the single-differential cross section in kinetic energy and polar angle of the primary proton are instead offered. In Figs. 8.29 and 8.30, the differential cross sections distributions corresponding to the most precise flux-averaged values measured of blinded and unblinded data are visualised; in Section 8.5, the values in each bin are tabulated. A phase space requirement according to the detector acceptance is applied, requiring a proton kinetic energy E_k of $>$125 MeV and a proton angle θ of $\cos\theta > 0.4$. Due to the binning scheme, however, the effective phase-space is rather $E_k > 181.71$ MeV and $\cos\theta > 0.752$.

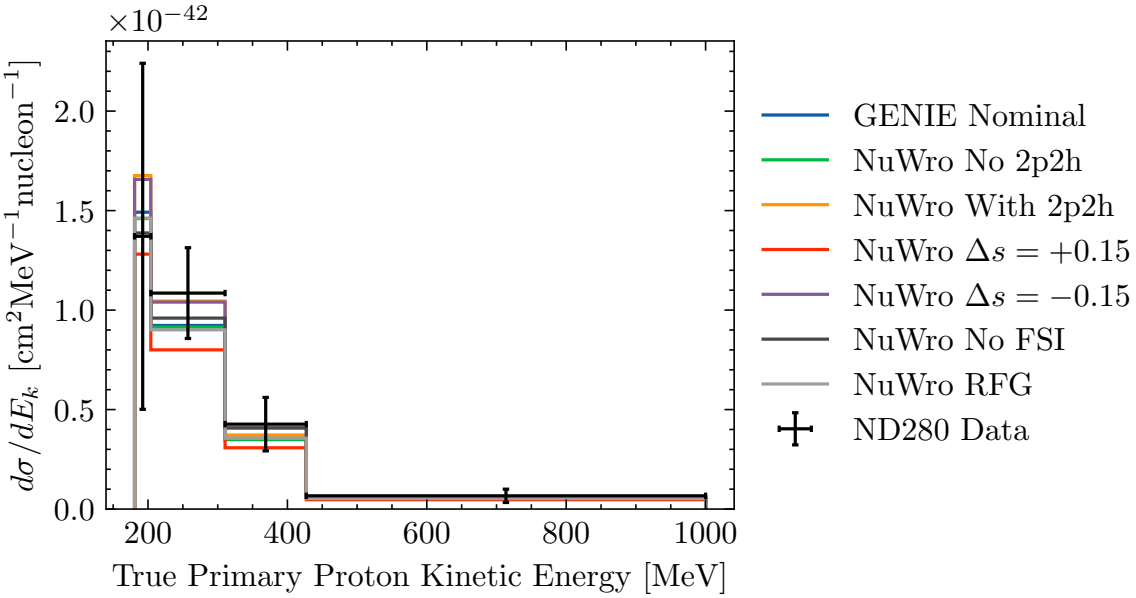

(a) Proton kinetic energy.

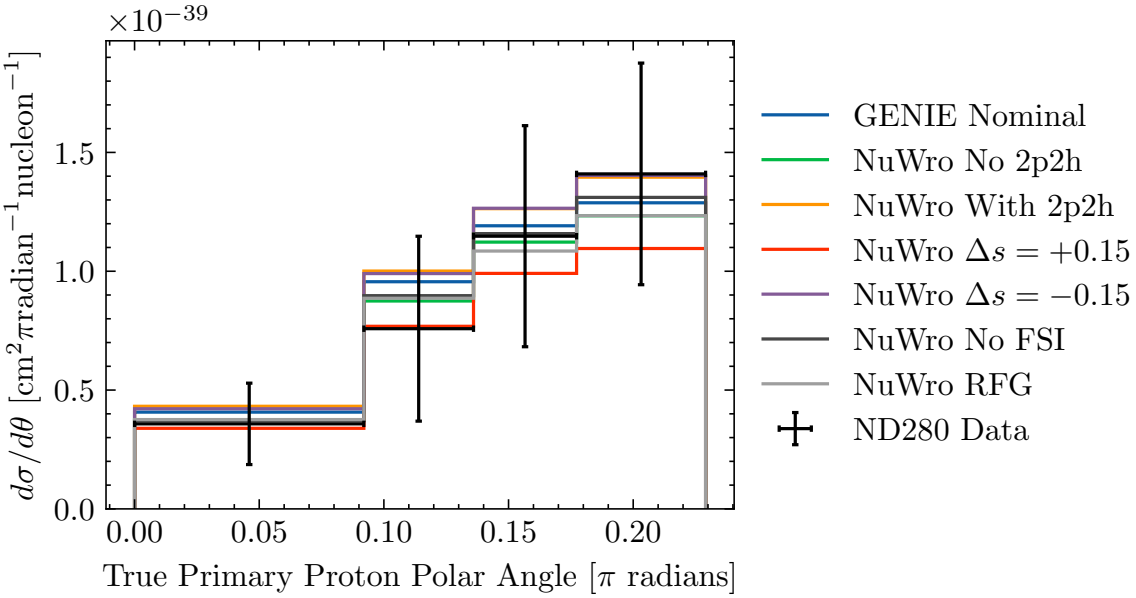

(b) Proton polar angle.

Fig. 8.29 Differential cross sections with respect to the kinetic energy (Fig. 8.30a) and polar angle (Fig. 8.30b) of the primary reconstructed proton. Events in the signal region are real detector data.

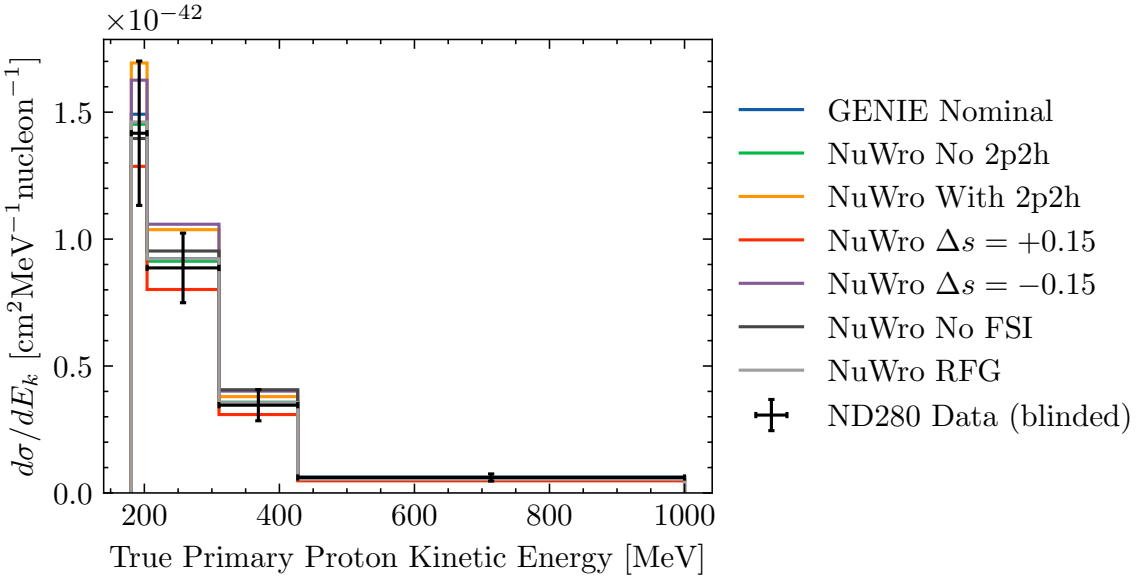

(a) Proton kinetic energy.

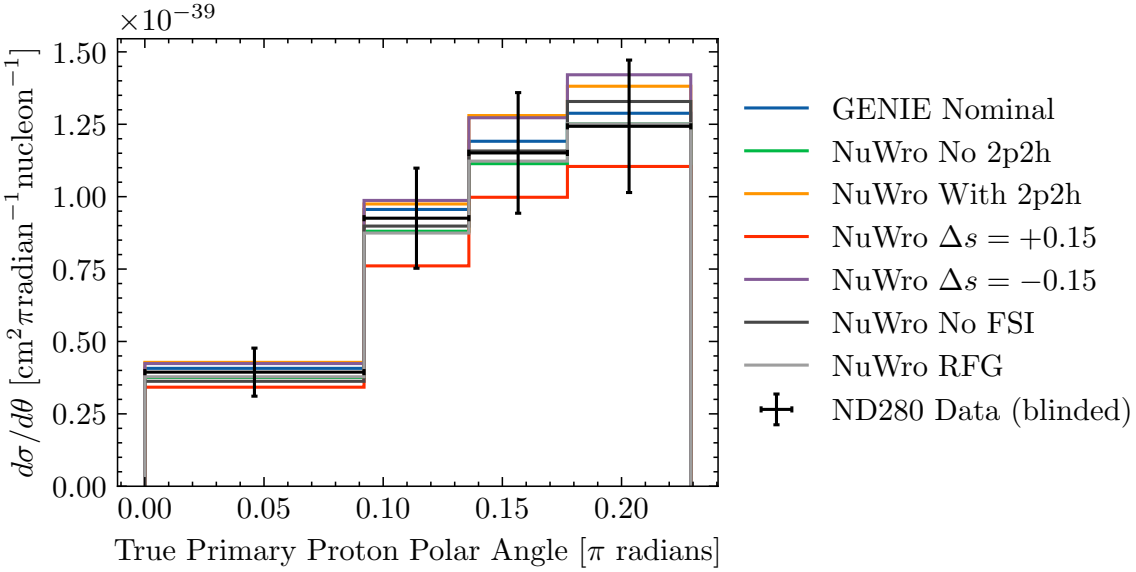

(b) Proton polar angle.

Fig. 8.30 Differential cross sections with respect to the kinetic energy (Fig. 8.30a) and polar angle (Fig. 8.30b) of the primary reconstructed proton. Events in the signal region are substituted for the nominal MC.

Model	$d\sigma/dE_k[10^{-43}\,\mathrm{cm}^2\,\mathrm{MeV}^{-1}\,\mathrm{nucleon}^{-1}]$			
	180.71–204.01 MeV	204.01–310.56 MeV	310.56–427.14 MeV	427.14–1000.0 MeV
ND280 Data	13.71 ± 8.70	10.85 ± 2.28	4.27 ± 1.34	0.67 ± 0.33
ND280 Data (blinded)	14.17 ± 2.84	8.86 ± 1.37	3.46 ± 0.61	0.61 ± 0.14
GENIE Nominal	14.92	9.23	3.57	0.63
NuWro No 2p2h	14.5	9.24	3.54	0.52
NuWro With 2p2h	17.14	10.37	3.8	0.55
NuWro $\Delta s = +0.15$	12.68	8.0	2.97	0.48
NuWro $\Delta s = -0.15$	16.17	10.47	4.03	0.6
NuWro No FSI	13.82	9.64	4.05	0.62
NuWro RFG	14.4	9.18	3.53	0.54

Table 8.6 Differential cross section values with respect to the kinetic energy of the primary reconstructed proton as measured at ND280 and predicted by various interaction models.

Model	$d\sigma/d\theta[10^{-40}\,\mathrm{cm}^2\,\mathrm{rad}^{-1}\,\mathrm{nucleon}^{-1}]$			
	0.0–0.091 969 rad	0.091 969–0.135 95 rad	0.135 95–0.177 29 rad	0.177 29–0.229 04 rad
ND280 Data	3.58 ± 1.71	7.58 ± 3.89	11.47 ± 4.65	14.10 ± 4.66
ND280 Data (blinded)	3.94 ± 0.83	9.26 ± 1.73	11.51 ± 2.08	12.43 ± 2.29
GENIE Nominal	4.07	9.56	11.91	12.88
NuWro No 2p2h	3.81	8.69	10.94	12.57
NuWro With 2p2h	4.27	9.97	12.79	13.92
NuWro $\Delta s = +0.15$	3.5	7.65	9.87	10.88
NuWro $\Delta s = -0.15$	4.29	9.97	12.49	14.1
NuWro No FSI	3.63	8.8	11.72	13.37
NuWro RFG	3.82	8.77	11.26	12.69

Table 8.7 Differential cross section values with respect to the polar angle of the primary reconstructed proton as measured at ND280 and predicted by various interaction models.

CHAPTER 9

CONCLUSIONS

A measurement of the neutrino–proton neutral current elastic interaction cross section at the T2K near-detector has been performed. The flux-averaged cross section within the phase-space $E_k > 125\,\mathrm{MeV}$ and $\cos\theta > 0.4$ was measured as $(3.31 \pm 0.36) \times 10^{-40}\,\mathrm{cm}^2\,\mathrm{nucleon}^{-1}$—in agreement with the value predicted by the nominal configuration of GENIE, which is $3.10 \times 10^{-40}\,\mathrm{cm}^2\,\mathrm{nucleon}^{-1}$. The differential cross sections with respect to the kinetic energy $d\sigma/dE_k$ and polar angle $d\sigma/d\theta$ have also been reported within the phase-space $E_k > 181.71\,\mathrm{MeV}$ and $\cos\theta > 0.752$.

The measurement was enhanced with a high purity sample of signal events selected by a machine-learning model performing binary classification. The classification performances of three algorithms based on decision trees were investigated, of which the AdaBoost algorithm was found to be the best. The model's hyper-parameters were optimised with a randomised search and its generalisation was evaluated with cross-validation. A dimensionality reduction was then performed with a marginal penalty to the selection performance. Compared to a nominal cut-based selection, the AdaBoost has a superior ability to reject charged-current events, the second-highest background present in the cut-based signal sample; its power to reject the highest background—neutrons emerging from outside of the fiducial volume—is on a par with the cut-based selection, however. The flux-averaged cross section as measured with the cut-based signal selection was $(2.92 \pm 0.44) \times 10^{-40}\,\mathrm{cm}^2\,\mathrm{nucleon}^{-1}$; thus, the machine-learning selection improved the precision of the measurement by 28%. The uncertainty on the flux-averaged cross section is dominated by the uncertainty on the background rate in the signal sample, therefore further work on the event selection should specifically aim for increased discrimination power between out-of-fiducial-volume neutrons and the signal.

9.1 Further Work

A few further studies can lead from a comparison of the neutral-current differential cross sections with those of charged-current quasi-elastic scattering. Even the most sophisticated interaction models around today fall short of describing the charged-current interaction accurately; combining these measurements may provide an enhanced power to discriminate between models and further constrain their parameters. Additionally, the ratio of the charged-current and the neutral-current differential cross section is a particularly valuable measurement given its sensitivity to the strange quark's contribution to the proton spin Δs. Before these studies can be undertaken, however, some improvements can be made to the current analysis. The cross section was unfolded with a relatively coarse binning in order to overcome a particular problem with the unfolding fit. When the binning was fine enough such that the data was dominated by statistical noise, the derived regularisation strengths were too strong and the unfolded data were biased to the MC. Hence, an improvement to the method used to derive the regularisation strengths is in order. In principle, all regularisation terms contribute a penalty to the likelihood simultaneously, whereas the current optimisation method considers the terms in isolation. It is conceivable that the L-curve method can be expanded to multiple dimensions—i.e. whether the regularisation strengths corresponding to the point on the likelihood surface at which the Gaussian curvature is maximised provide better estimates of the ideal strengths. Otherwise, there are certainly alternative regularisation methods to consider. Given that the regularisation can be improved, another improvement is to optimise the binning scheme. Now that the optimum signal probability score threshold is known, the differential cross sections may be measured again with a finer and non-arbitrary binning scheme. Additionally, further validations of the analysis framework could be made. The complete analysis was validated by performing the measurement with data substituted for GENIE nominal in the signal region. Further validations could be conducted with fake signal data from the NuWro generator and a set of statistical variations. Additionally, there is currently no control sample dedicated to neutral-current interactions that produce one or more pions—a subdominant background. It is preferable that the rate of these interactions extracted from the fit to the sidebands should be cross-checked.

The current signal selection can be improved in a myriad of ways. First of all, however, to be checked is an approximate 10% ($\sim 1\sigma$) discrepancy between the flux-averaged cross sections measured of all of the multivariate signal samples and the cut-based sample. This discrepancy originates from a $\sim$4% excess of data in the multivariate signal samples, to which the fit to the sidebands responds by increasing the signal rate. In the cut-based selection, the opposite occurs—there is an initial $\sim$4% deficit, which the fit corrects for by, rather,

decreasing the signal rate. A comparison of the data within each of these samples may reveal the cause of this discrepancy. Once this is understood, the classification performance of the model can be enhanced, perhaps the simplest of which is to conduct the training with a larger sample of signal events. Alternative classification models could be investigated as well—for example, other boosting algorithms such as ε-boost or Logit-boost [130], or perhaps deep-learning algorithms, such as neural networks. Additional training variables also could be considered that target specific backgrounds—for example, FGD track timing is known to contain some information about whether the proton was produced by an out-of-fiducial-volume neutron. Alongside improvements to the classification performance, additional measures should be taken to minimise bias to a particular interaction model. The training sample currently consists of events generated by the nominal GENIE configuration; perhaps training and testing samples could be constructed that draw from a variety of interaction models. In such an arrangement, during the optimisation stage, the learner's generalisation would be evaluated against dedicated validation samples for each interaction model in order to assess biases either to the training or testing sample or to a particular interaction model. The current classifier identifies the signal definition, which certainly varies from interaction model to interaction model—an alternative approach would be rather to train the classifier to identify particle types. In this approach, a high statistics training sample would be constructed from particle gun MC, generated by a variety of interaction models. With control over the truth-level kinematic phase-space of the training sample, model dependencies present in the particle kinematics can be prevented from being learned. This approach would likely improve the efficiency with which neutral-current interactions with multiple protons in the final state are identified which is poor at present.

PART III

DEVELOPMENT OF CALIBRATION OPTICS FOR THE HYPER-KAMIOKANDE DETECTOR

CHAPTER 10

INTRODUCTION TO PART III

In this part of the thesis, the development and prototyping of an optical calibration system for the future Hyper-Kamiokande neutrino detector are reported.

Firstly, the general design and physics potentials of Hyper-Kamiokande will be recounted in Chapter 11. Then, the design principles of the calibration hardware will be reported in Chapter 12. The installation of said calibration system within the Super-Kamiokande detector and subsequent testing will be recounted in Chapter 13. Finally, in Chapter 14, conclusions will be drawn and the present status of the hardware will be reported.

Introduction to Part III

CHAPTER 11

THE HYPER-KAMIOKANDE DETECTOR

Hyper-Kamiokande (henceforth HK) is planned as the third in the lineage of Kamiokande-type detectors; by target mass, it will be the largest of the next-generation neutrino detectors by some margin. Like its predecessor, the purpose of HK is twofold: it will serve its own physics programme of topics that include (but are not limited to) proton decay, supernova monitoring and astrophysical neutrinos, and also will assume the role of far detector, thus replacing SK, in an upgraded configuration of T2K known as T2HK.

As expected of a project in its early stages, the technical details given in this chapter are subject to change but may be accepted as accurate at the time of writing [144]. As well, it should be noted that HK, as a project, may eventually span two detectors—the second to be built some years after the first (and possibly at a larger baseline in Korea [145]). In this text, however, only the first detector will be discussed and the name HK will refer to it exclusively.

11.1 Technical Design

HK will lie just 8 km south of SK in a mine tunnelled into an adjacent mountain. The proposed cavern site is exposed to the T2K neutrino beam at the same off-axis angle and baseline as SK. Much of the detector design is the same as SK albeit on an enlarged scale: the target volume will be a cylindrical tank 68 m wide and 71 m high filled with 258 kt (217 kt inner detector volume) of ultra-pure water [43]. In Table 11.1, the specifications of HK are compared to those of Kamiokande and Super-Kamiokande. Various illustrations of HK can be seen in Fig. 11.1. But this upscaling is not without some technical challenges; one of these concerns the water system. In any Cherenkov detector, transparency of the tank medium is key if the Cherenkov light is to travel a significant distance before reaching the photo-sensors. The standard of water quality in SK, which will be matched in HK, is such that

the attenuation length for photons with wavelengths within 400 nm to 500 nm is maintained above 100 m. A faster processing rate of the purification system of $310\,\mathrm{t\,h^{-1}}$ (from $60\,\mathrm{t\,h^{-1}}$ in SK) is necessary to achieve this. As well, steps must be taken to control radon contamination levels in the water, as the natural decay of radon is the dominant source of background at low energies. In SK, the radon concentration is maintained below $1\,\mathrm{mBq\,m^{-3}}$ by vacuum de-gasifiers supplied with radon free air (as the purge gas) at a rate of $20\,\mathrm{N\,m^3\,h^{-1}}$; a faster rate of $50\,\mathrm{N\,m^3\,h^{-1}}$ is required for HK to match this background level. Pending findings from SK, the possibility of doping the water with $GdSO_4$ is also under consideration.

	Detector		
Parameter	K	SK	HK
Depth underground [m]	1000	1000	650
Tank diameter [m]	15.6	39	74
Tank height [m]	16	42	60
Total volume [kt]	4.5	50	258
Fiducial volume [kt]	0.68	22.5	187
Outer detector width [m]	$\sim$1.5	$\sim$2	1 to 2
No. inner detector photo-sensors	948	11 129	40 000
No. outer detector photo-sensors	123	1885	6700
Photo-sensor coverage	20%	40%	40%

Table 11.1 Specifications of the three generations of Kamiokande-type detectors. Adapted from Ref. [144].

Improved 20" PMTs[1] are to cover 20% to 40% of the inner surface of the inner detector, possibly interspersed with multi-PMT modules. These newly-developed PMTs inherit much from those used in SK[2] by virtue of their reliability, though swap the type of dynode from Venetian blind to box-and-line and the photo-cathode for one larger in quantum efficiency and sensitive area. This makes for a tube around twice as efficient in detecting single photons and twice as finer in resolution in measuring their charge, compared to those of SK.

[1] Hamamatsu R12860

[2] Hamamatsu R3600

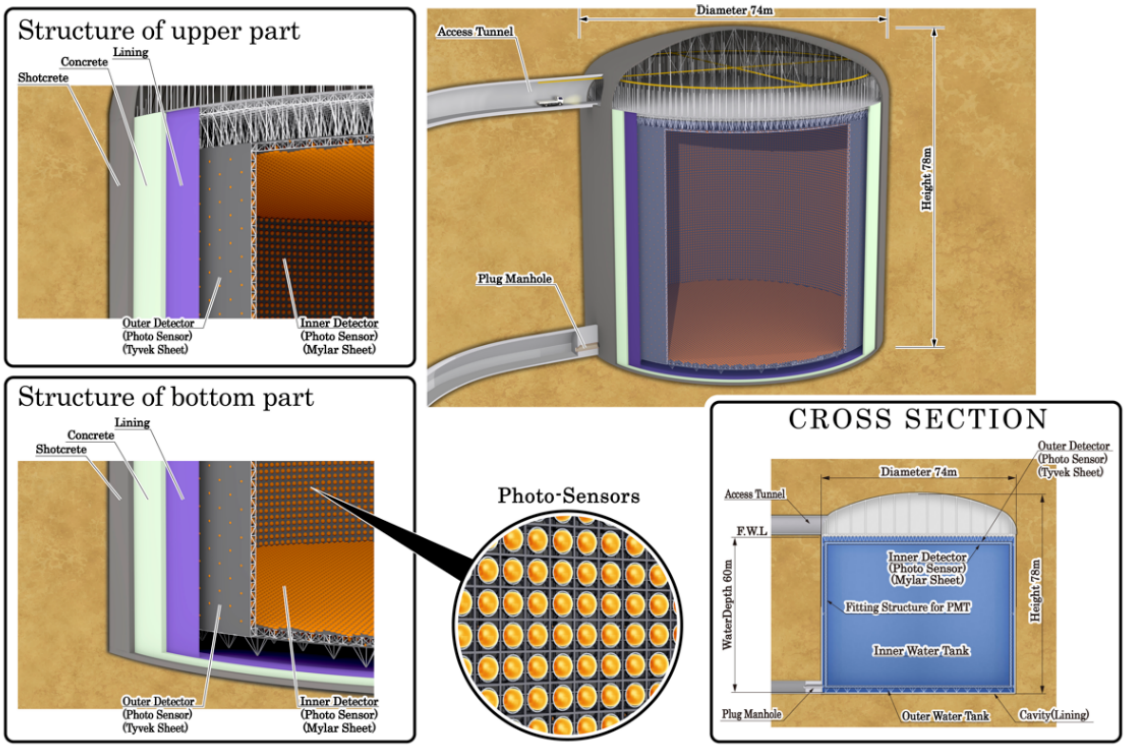

Fig. 11.1 Assorted diagrams of the Hyper-Kamiokande detector. Reproduced from [144].

11.2 Physics Potentials

Of the neutrino oscillation experiments belonging to the next generation, T2HK is perhaps the most qualified to probe CP asymmetry in the neutrino sector. The advantage granted to T2HK in this measurement over ultra-long baseline experiments is the comparatively short baseline at which HK stands. Neutrinos which traverse, and thus oscillate along, very long distances ($> 1000\,\text{km}$) through matter are subject to the Mikheyev–Smirnov–Wolfenstein (MSW) effect, which perturbs the ν_e appearance probability function. Though exposure to this effect grants sensitivity to the neutrino mass orderings, it also introduces a critical uncertainty on mixing parameter measurements when the composition of matter along the baseline is not well understood.

At present, in T2K, the number of ν_e events collected at the far detector limits the confidence level with which the parameter space of δ_{CP} can be probed. Regardless, results from T2K suggest that CP-violation may be maximal ($\delta_{CP} \sim -\pi/2$) [2]. With T2HK, it is projected that with ten years of a beam operating at 1.3 MW δ_{CP} can be measured to better than 23° for all possible values and CP-violation can be observed to a confidence level of more than 3σ (5σ) for 76% (57%) of the phase space of δ_{CP}. In Fig. 11.2, the significance

to which, given normal ordering of the MH, HK can reject CP conservation after ten years of running is shown. Measurement of the atmospheric mixing parameters, Δm^2_{32} and $\sin^2\theta_{23}$, and various interaction cross sections will also benefit from the larger data sample. Over an equivalent running time, DUNE has a marginally weaker sensitivity to δ_{CP}, being projected to establish CP violation to 5σ for 50% of the phase space [146].

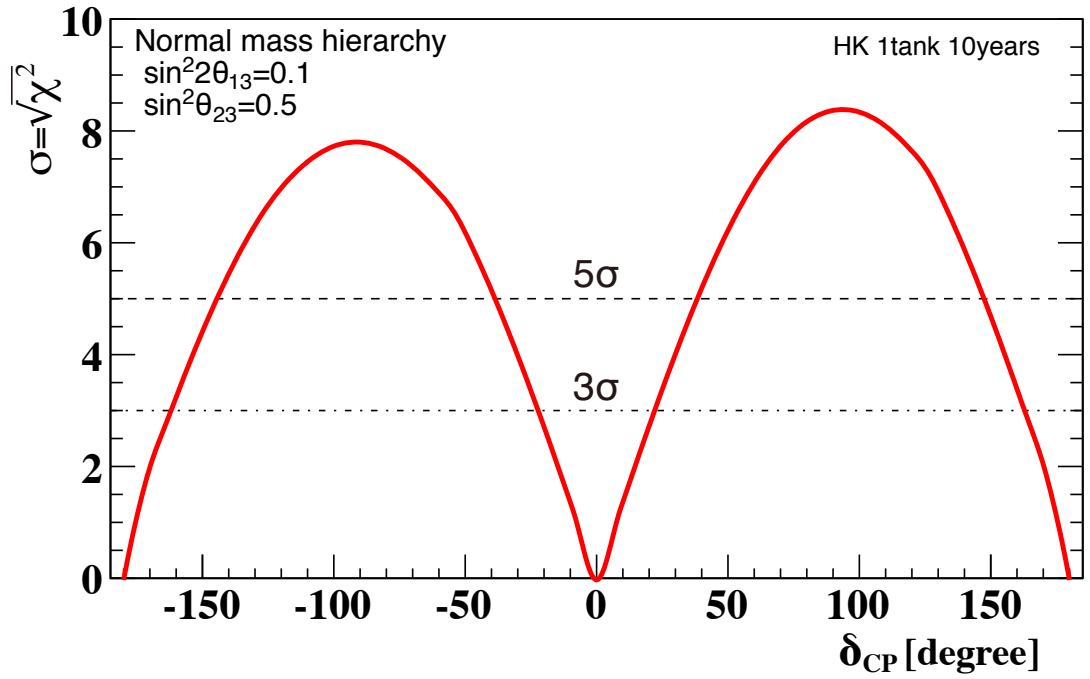

Fig. 11.2 Projected significance with which Hyper-Kamiokande can outrule charge–parity conservation ($\sin\delta_{CP} = 0$) in the neutrino sector given the normal ordering of the neutrino masses. Reproduced from [144].

Water Cherenkov detectors, by virtue of their large fiducial volumes, are naturally sensitive to a variety of nucleon decay modes, and HK, as the largest, can place some of the strictest experimental limits on their lifetimes. In Fig. 11.3, it is seen that HK is generally able to improve on the limits on lifetime set by SK on some of the principal decay modes predicted by various GUTs by approximately an order of magnitude. The upgrades to the photo-sensors as mentioned especially aid HK in these searches, in which atmospheric neutrino interactions are expected to present themselves as a significant source of background. Their presence may however be inferred by the detection of neutrons that are produced either directly or indirectly upon interaction, whereas neutrons are rarely a sign of nucleon decay. Neutrons are already identified in Super-Kamiokande by the detection of a characteristic 2.2 MeV photon emitted on their capture by hydrogen nuclei, though the signal is not strong. With the

improved photo-sensors, the tagging efficiency is expected to greatly increase—from 18% to 73%—for neutron capture events that produce at least ten detector hits.

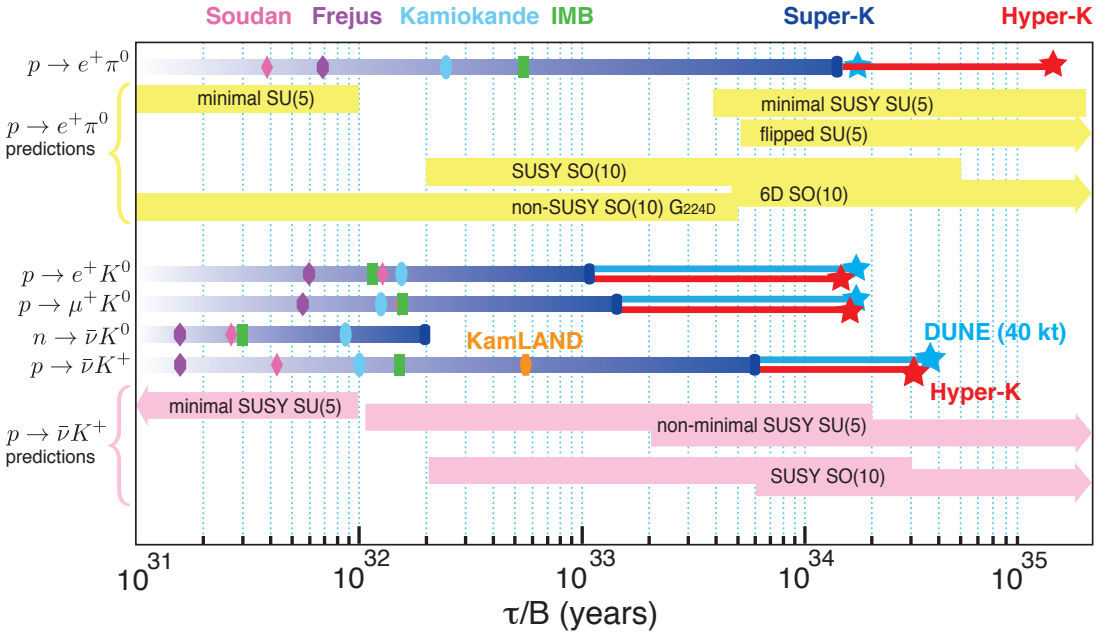

Fig. 11.3 A comparison of lifetime limits at 90% CL set by various historical, present and future experiments on various nucleon decay modes, and their corresponding lifetimes as predicted by various Grand Unified Theories. The projected limits for Hyper-Kamiokande and DUNE are based on 10 years of exposure. Reproduced from [144].

HK's minimum energy threshold of several MeV allows it to study astrophysical neutrinos on an event-by-event basis. The MSW effect, as observed through the behaviour of ^{8}B solar neutrinos, was confirmed at the 3σ level in SK and thus will be measured more precisely in HK. And in doing so, either an existing $\sim 2\sigma$ tension [144] (suspected to be caused by the MSW effect) between the best fit values of Δm^2_{21} as measured by solar and reactor neutrino experiments will be resolved, or, if a real discrepancy, new physics will be hinted. As well, HK is to function as a supernova monitor in the long-term, with sensitivity to collapses up to several Mpc away. It is projected that should a core-collapse supernova occur at the centre of the Milky Way, HK is to register approximately 10^4 neutrinos in a single second. Such an observation would clear up the many questions left by the limited data sample induced by SN1987A.

11.3 The Intermediate Water Cherenkov Detector

Given its size and exposed beam power, HK will accumulate enough ν_e appearance events to diminish statistical uncertainties to such a degree that, for the first time, systematic uncertainties will dominate the total error on oscillation measurements. In T2K, the near detector is vital in reducing the uncertainty on the event rate predicted at the far detector to 7%—but for T2HK this is insufficient. The largest contribution to this error is from uncertainties in nuclear interactions, which in turn originate from the differences in specification between the near and far detectors (their target material, acceptance and exposure to different backgrounds). To combat this the neutrino cross-sections on water are to be measured directly by a second water Cherenkov detector, known as the Intermediate Water Cherenkov detector (henceforth IWCD), placed along the T2HK baseline, not far downstream from the near detector facility.

This new detector will stand 2 km from the neutrino production point—at a position where the beam will appear comparatively point-like to that of the near-detector—to evade occurrences of pile-up. The design of the detector itself is Kamiokande-like—3215 8" photo-sensors will instrument the inner region of a cylindrical tank 8 m high and 10 m wide. These tank dimensions will be sufficient to contain muons with momenta within the region of interest. Perhaps peculiarly, the tank position will not be fixed: the tank is to be suspended on a crane within a 50 m shaft, and lowered or raised as to sample the neutrino spectrum at off-axis angles between 1° and 4°. The motivation for this is threefold. Firstly, as the dependence on the off-axis angle of the neutrino spectrum, these sampled spectra can be linearly combined to recreate practically any arbitrary spectra. Such a technique thus mostly eliminates those uncertainties associated with extrapolating the near detector spectra to the far detector. Secondly, since the proportion of ν_e, $\bar{\nu}_e$ intrinsic to the neutrino beam increases dramatically with off-axis angle, measurement of the ν_e to ν_μ cross section ratio ($\sigma_{\nu_e}/\sigma_{\nu_\mu}$) with good precision becomes feasible with the detector sat at its extremity. And lastly, the LSND and MiniBooNE short-baseline anomalies can be cross-checked robustly by looking for an oscillation signature indicative of a light ($1\,\mathrm{km\,GeV^{-1}}$) sterile neutrino in both the neutrino energy spectra and the off-axis angle.

11.4 Project Timeline

The filling of the tank is planned to begin 2027, and data-taking is expected to start shortly after. As of 2021, excavation of the access tunnel [147] and production of photo-sensors is underway [148]; the latter of which will fulfil its quota in 2026. Excavation of the tank cavern is expected to take three years following the completion of the access tunnel at the end of 2021, after which the tank will be constructed. Upgrades to the beam and near detector, to transform T2K into T2HK, will be performed in parallel as mentioned in Chapter 5.

CHAPTER 12

PRODUCTION OF CALIBRATION OPTICS FOR THE SUPER-KAMIOKANDE AND HYPER-KAMIOKANDE DETECTORS

A variety of techniques are used in SK to calibrate detector properties, such as the photo-sensor responses and the water's optical properties, and to fix particular physics parameters, such as the energy scale, energy resolution and detection efficiencies [149]. These calibration techniques developed for SK are also satisfactory for HK, given their likeness. Hence, HK's calibration systems will not be dissimilar to those of SK. A number of detector calibrations are achieved in SK with automated systems that periodically inject visible light with well-understood properties into the tank. For example, the amount of optical scattering and absorption present in the tank is measured by pulsing collimated laser light of various wavelengths across the tank from various injection points. As well, hanging in the centre of the tank is the diffuser ball—an isotropic light source used to perform relative gain calibrations of the ID PMTs. The operation and maintenance of such systems in a detector as large as HK is foreseen to be particularly labour intensive, given the sheer number of injectors to calibrate the entire detector. Hence, an equivalent automated system for HK is in development by UK groups that integrate the sources required for the photo-sensor and water property calibrations within a single injector module.

In the summer of 2018, Super-Kamiokande was drained for refurbishment[1,2]—the first time in 12 years. The opportunity was seized to install five prototype light injectors along the vertical of the inner detector as the water level was gradually lowered over a period of three months. Soon after, the injectors were connected to the system's dedicated LED pulser; and later, the system was integrated into SK's automated calibration routine. The optics belonging to each injector module were developed in the University of Warwick's HK laboratory. In this chapter, the design and production of these optics will be reported.

12.1 Physics Requirements

The light injection system as a whole consists of a source, a monitor photo-sensor and an array of optics. The source, located at the top of the tank, delivers light simultaneously to the monitor photo-sensor, also at the top of the tank, and the optic within the tank via optical fibres. A number of requirements are demanded of these pieces of hardware to meet the physics goals.

The system of light injectors must be capable of performing a number of different calibrations, these include optical scattering and attenuation of the detector medium, photo-sensor charge uniformity, photo-sensor timing uniformity and the photo-sensor angular responses. These are each accessed by injecting either a diffuse or collimated source into the detector. Hence, each injector module will hold both a collimating and a diffusing optical device. The calibration measurements are described as follows.

The scattering measurement is best performed by injecting a source into the tank of such high intensity that a scattered light signature is observable in the timing distributions of photo-sensors that are not directly exposed to the source. Ideally, the intensity would be as high as possible without saturating the photo-sensors targeted by the source. The scattering coefficient would be extracted by the MC-driven method currently used in SK. A narrow-beam source is preferred for two reasons: a vanishing beam-angle minimises the uncertainty on the scattering position of the direct light and also maximises the number of photo-sensors exposed exclusively to the scattered light, thereby increasing the statistics of the sample. The attenuation measurement is achievable with either a collimated or diffuse source, though the methodology differs for each source type. With a collimated source, the

[1] During which the author took part in various refurbishment activities in the outer detector; these include cleaning the super-structure, replacing dead PMTs, mounting new PMTs to their wavelength-shifting plates and making power connections.

[2] A number of faulty photo-sensors were replaced in the inner detector during this refurbishment period, these can be seen by comparing the hit-maps generated with data taken before and after the refurbishment in Section 13.2 and Section 13.3.1.

relative time-variation of attenuation can be inferred by the variation of charge recorded by the photo-sensors directly exposed to the source. Previously in SK, this measurement was performed by projecting light from an optical fibre from the top of the tank to image a spot with a diameter of $\sim$10 m on the bottom cap. For the diffuse source, the attenuation can be inferred by considering the dependence of the hit rates recorded by each photo-sensor on their distance from the source. In this approach, low-intensity pulses are preferred such that single photo-electron hits are recorded by the sensors. The photo-sensor calibrations are best performed with a diffuse source in order to provoke a response from as many photo-sensors as possible with a single source. Charge uniformity is calibrated by measuring the peak of the single photo-electron distribution for each photo-sensor exposed to the source. The calibration for uniformity in timing response requires short pulses from a source uniform in timing across the angular spectrum. Finally, probing the angular responses of each photo-sensor requires numerous diffuse sources at signal photo-electron intensity, deployed at a diversity of angles with respect to the pointing direction of each photo-sensor.

The number of injectors in the tank is not yet set. There are several considerations that influence this number. Firstly, knowledge of the depth dependence of the water parameters is desired; at a bare minimum, injectors should be located at least five different positions along the vertical of the tank. Secondly, the angular response measurement demands each photo-sensor is exposed to at least two diffusers at different positions. Should a diffuser fail, however, it is essential that no photo-sensor is rendered unreachable to the system; hence, a six-fold redundancy in diffuser coverage for each photo-sensor would be a sensible safeguard against such failures. There are also several practical requirements demanded of the injectors: they must be able to survive decades submerged in water and, potentially, gadolinium doped water, their optical properties must show long-term stability, and they must be able to withstand the high-pressure conditions ($\sim$9.5 bar) at the 70 m depth at the bottom of the tank.

Several properties are also required of the light source for these calibrations[3]. The intensity should be variable over a large dynamic range, from as low as $\sim 10^3$ ppp (photons per pulse) to achieve single photo-electron hits for the photo-sensor calibrations and up to $\sim 10^6$ ppp for the scattering measurement. The pulse width should ideally be small compared to the typical time-of-flight of injected light to a photo-sensor (i.e. less than 10 ns). In the case of the photo-sensor timing calibration, the shorter the pulse the better. Finally, the source must be able to switch between multiple different wavelengths in order to probe the wavelength dependence of scattering and attenuation parameters. Previously, these measurements are performed in SK with 337 nm, 375 nm, 405 nm, 445 nm and 473 nm sources.

[3] For the best account of the electronics Ref. [150] is referred to.

12.2 The Light Injection Module

The current design of the light injection module is shown in Fig. 12.1; photographs can be seen in Fig. 12.2. Each module integrates three light sources: a diffuser, a collimator and an optical fibre, to function as a control source. All water-exposed components are manufactured from 316 stainless steel—a grade recommended for its resilience to gadolinium-induced corrosion. The collimator and diffuser are mounted to the module with grub screws, the control fibre with a threaded FC adapter. The plate onto which the optics are mounted is affixed to an L-bracket—which, in turn, is fixed to the super-structure—with a kinematic mount. Adjusting the bolts on the kinematic mount tunes the pointing direction of the optics in three degrees of freedom. To gauge the pointing direction of the module during alignment of the module during installation there is an additional mounting point on the front plate for a laser pointer.

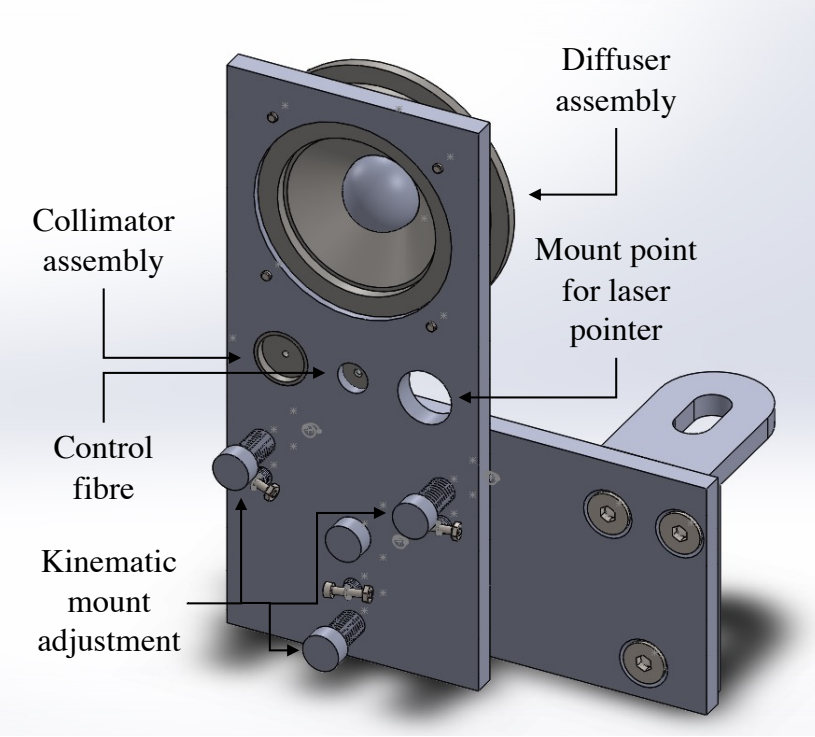

Fig. 12.1 Diagram of the light injector module.

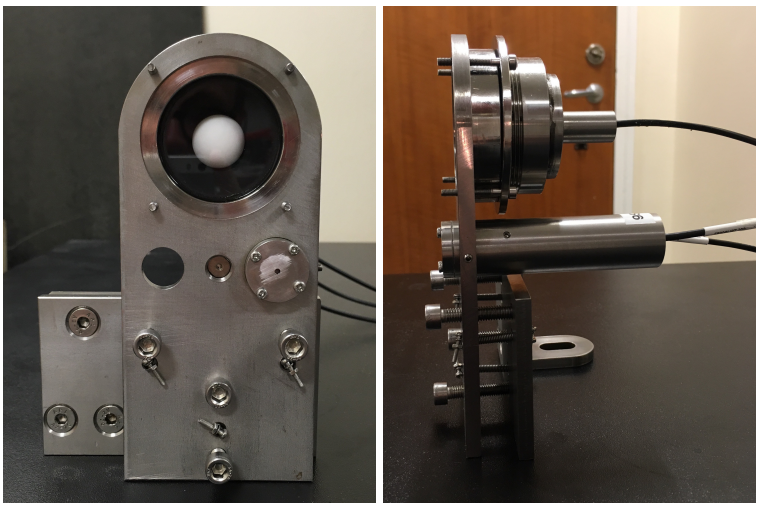

(a) Front view. (b) Side view.

Fig. 12.2 Front (Fig. 12.2a) and side (Fig. 12.2b) views of an assembled injector module.

12.3 The Optical Fibres

The optical fibres coupled to the optics, as well as the control fibre, are 0.5 numerical aperture (NA), 200 μm diameter core step-index type[4]. A 70 m length of fibre connects a port of the light source to each optic in the tank. The optics are assembled with a 1 m length of the same fibre glued in. The source fibre and the injector fibres are ultimately coupled within a water-tight connection box close to the injector. The specifications of this fibre model may be found in Table 12.1.

The attenuation spectrum for this fibre may be seen in Fig. 12.3, the power loss is expected to be around ⅓ as the injected light reaches the optics. In HK, given the longer lengths of fibres required, the fibres are likely to be graded-index type in order to preserve the timing structure of injected pulses.

The light output of the optical fibre type was measured at the collimator's test stand. The test stand's apparatus and the analysis methods are described later in Section 12.5.3. Shown in Fig. 12.4 are measurements of the beam cross section, taken at increasing distances along the beam axis. The half opening angle at full width half-maximum intensity is $(21.6 \pm 1.8)°$. The radial intensity distributions measured of the fibre output at two distances along the beam axis are shown in Fig. 12.5—the intensity profile resembles a Gaussian distribution.

[4] Thorlabs ® FP200URT

Parameter	Value
Wavelength range	300 nm to 1200 nm
Numerical aperture	0.50
Core index (at 589.3 nm)	1.458434
Cladding index (at 589.3 nm)	1.3651
Core diameter	$(200 \pm 5)\,\mu$m
Cladding diameter	$(225 \pm 30)\,\mu$m
Coating diameter	$(500 \pm 30)\,\mu$m
Core material	Pure silica
Cladding material	Hard polymer
Coating material	Tefzel ®

Table 12.1 Specifications of the optical fibres deployed within the Super-Kamiokande light injector system [151].

12.4 The Diffuser

12.4.1 Technical design

The diffuse source is generated by scattering the injected light with a block of compressed polymethyl methacrylate (PMMA) powder. The transmittance of this material is upwards of 90% for wavelengths in excess of 400 nm; at 320 nm the transmission reduces to ∼50%. The shape of the diffuser deployed in SK is a hemisphere of 20 mm radius supported by a cylindrical base of 30 mm radius and 10 mm height. The diffuser material was found to be porous and prone to leach in water—properties neither beneficial for the water quality of the tank nor the longevity of the device. For these reasons, the invention of a water-tight enclosure was a particular focus of the diffuser's development process[5]. The diffuser and enclosure designs used for the SK light injectors are shown in Fig. 12.6. The optical properties of the device are changed by the enclosure. In particular, the intensity of the projected light is uniform to 10% within ±35° field when the diffuser is housed, compared to ±20° for the isolated diffuser. The improvement in performance is prompted by reflections occurring within the *torch*-like interior of the enclosure. The timing delay introduced by the diffuser is uniform to less than 1% within the ±40° field with or without the enclosure.

[5] Ref. [152] is referred to for the best account of these studies.

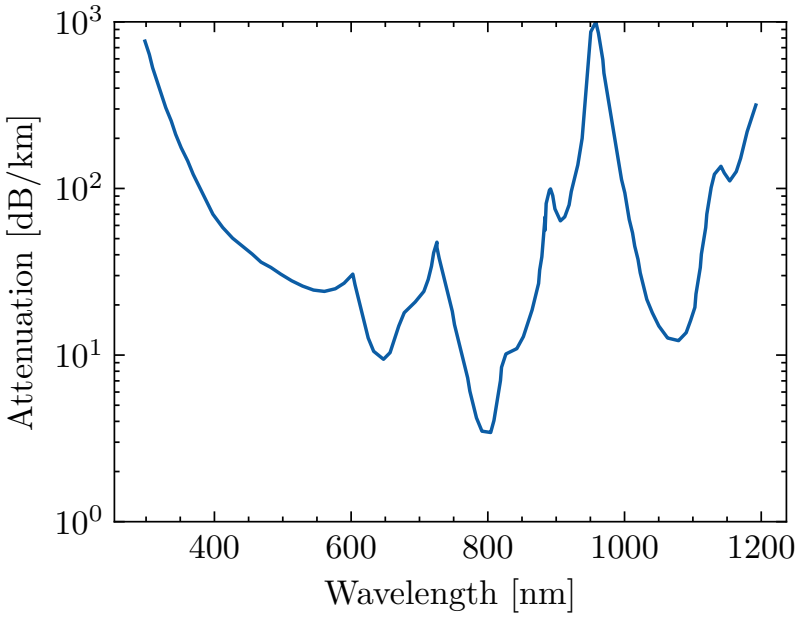

Fig. 12.3 Attenuation levels across the wavelength range of the optical fibres deployed within the Super-Kamiokande light injector system [151].

12.4.2 Test stand

The diffuser characteristics were captured within a dedicated test stand at the University of Warwick. In short, the diffuser to be tested is mounted onto a rotational stage and injected with pulses from a semiconductor laser of either 450 nm or 520 nm wavelength, then an angle-dependent response is recorded of a 2" photo-sensor module as the diffuser is rotated in place. A photograph of the test stand is shown in Fig. 12.7a.

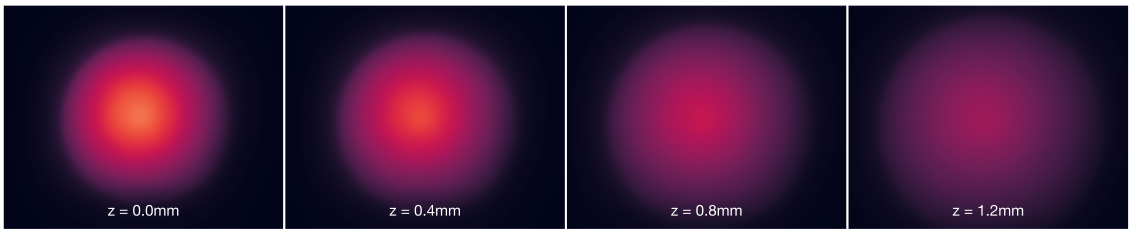

Fig. 12.4 Images of the beam profile emitted by the optical fibre model deployed within the Super-Kamiokande injector system as measured by the collimator test stand system. Each image shown was recorded at a particular distance z from the fibre ferrule.

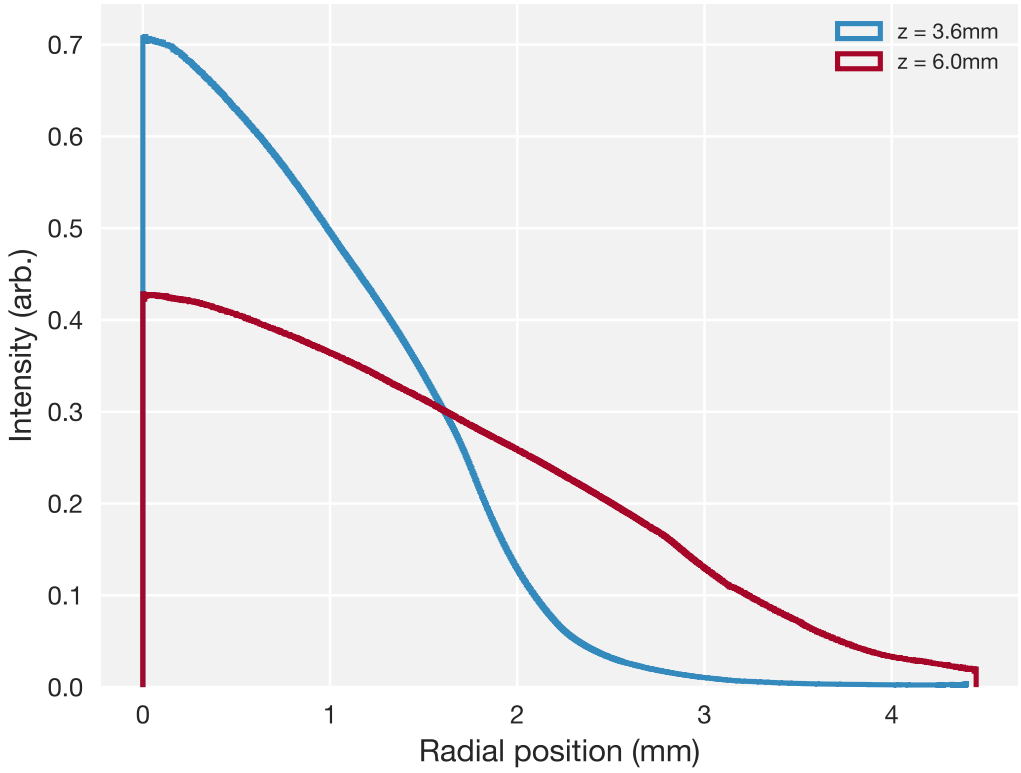

Fig. 12.5 In-air radial intensity distribution of the optical fibre model deployed within the Super-Kamiokande injector system as measured at two longitudinal distances in the laboratory.

12.4.3 Production

Seven diffusers were produced for the barrel injector installation, of which two were designated as spares. After assembly, each unit was required to survive at least 24 hours within a water-filled vessel pressurised to 10 bar. As well, each unit was tested for condensation build-up when chilled—given that SK's water is maintained at 13 °C. Accordingly, each unit was placed at 5 °C for three days, then under 0 °C for two days. No condensation was observed for any unit when visually inspected after their cold exposure. The angular intensity profiles of each unit following these tests are shown in Fig. 12.7b. The current status of each unit is stated in Table 12.2.

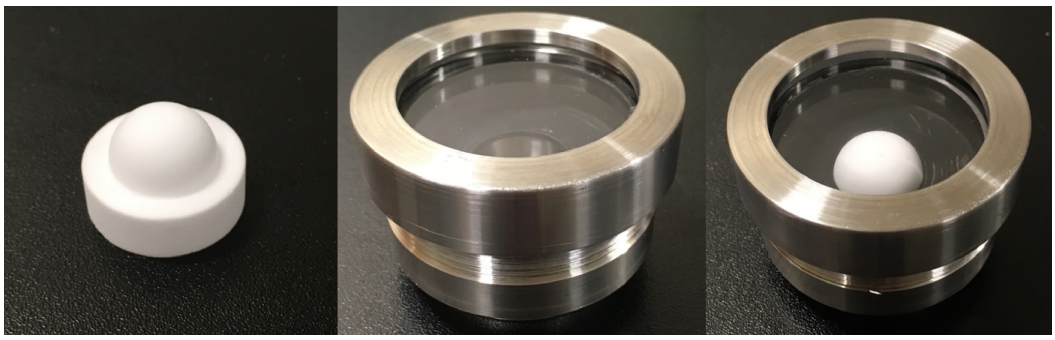

Fig. 12.6 A wide-angle diffuser (left), its enclosure (centre) and the diffuser within the enclosure (right).

Unit Index	Status
1	Installed at B1.
2	Designated as spare, returned to University of Warwick.
3	Installed at B3.
4	Installed at B2.
5	Designated as spare, returned to University of Warwick.
6	Installed at B5.
7	Installed at B4.

Table 12.2 Produced diffusers and their installation point in Super-Kamiokande. Details of the installation points may be found in Table 13.1.

12.5 The Collimator

12.5.1 Technical design

The production design of the collimator for the Super-Kamiokande detector is depicted in Fig. 12.8; the assembly is comprised of three parts: *i)* the lens-mount, *ii)* the sleeve and *iii)* the cap, the function of each will be described. A photograph of these components are shown in Fig. 12.9a.

Collimation of the injected fibre light is provided by a radial Gradient-Index (GRIN) lens of cylindrical shape. According to the distribution and type of material deposited within the lens during the manufacturing process, the refractive index of a radial GRIN lens can be made to vary with distance from the optical axis in a desired manner. The refractive index of the collimator's lens varies parabolically as $n_r = n_0\left(1 - gr^2/2\right)$, where n_r is the refractive index at radial position r, n_0 is the refractive index along the optical axis and g is the gradient constant. Such a property causes incident rays of light to trace sinusoids of period $\frac{2\pi}{g}$ as they

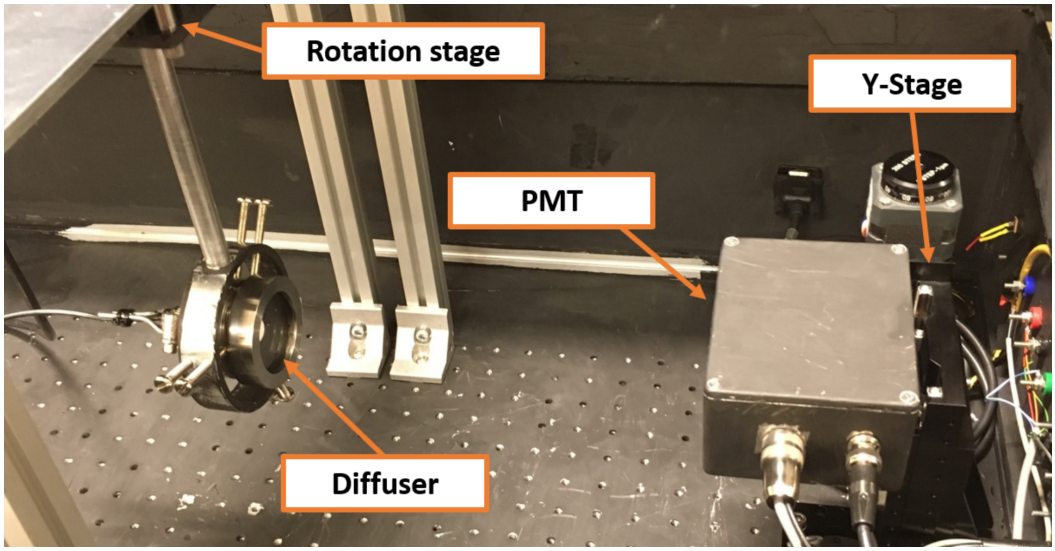

(a) The test stand (reproduced from Ref. [152]).

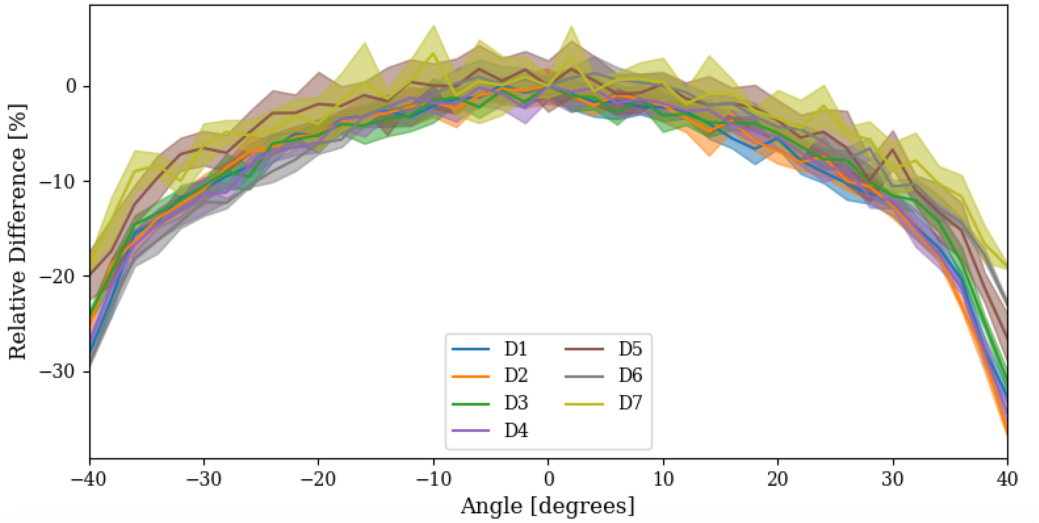

(b) Intensity relative to peak as a function of polar angle from the optical axis of each production diffuser (reproduced from Ref. [152]).

Fig. 12.7 The diffuser test stand belonging to the University of Warwick's HK laboratory (Fig. 12.7a) and angular intensity measurements taken of each production diffuser with said test stand (Fig. 12.7b).

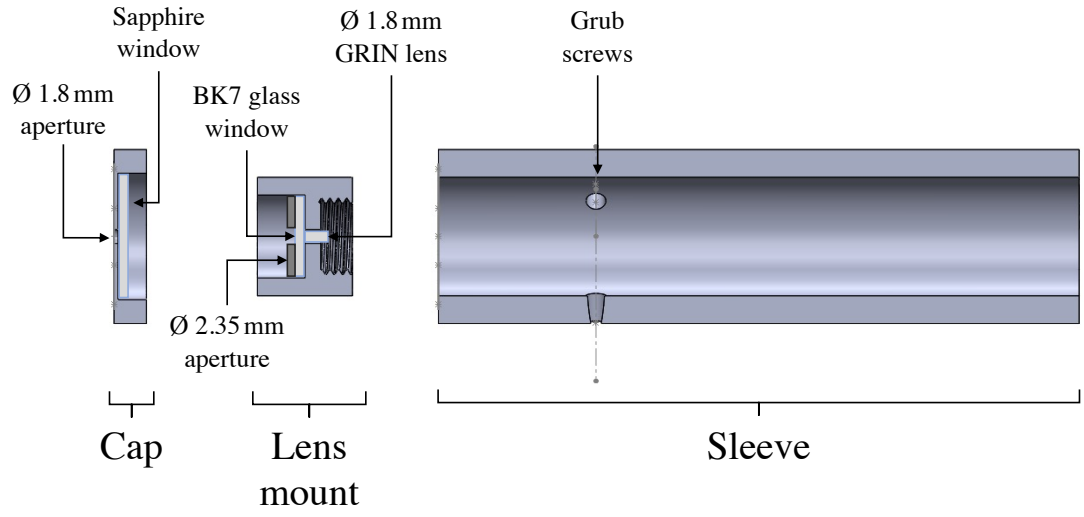

Fig. 12.8 Diagram of the collimator assembly produced for the Super-Kamiokande injectors.

pass from one lens face to the other. In the case that rays are incident at a non-zero angle with respect to the radial axis, the trace is helical. The geometrical length of such a lens is given by $\frac{2\pi P}{g}$, where P is a quantity known as the lens *pitch*—corresponding to the number of oscillation cycles that an incident ray will trace through the lens. In the ideal case where a ray is emitted along the optical axis at one face of the lens and at an angle within its acceptance, the ray direction will be parallel with the optical axis—i.e. achieving perfect collimation—at phases $\frac{\pi n}{2}, n \in \mathbb{Z}$ in the oscillation cycle. Therefore, a focussing GRIN lens is of ¼ pitch, as is the one present in the collimator, and accordingly the lens length is such that incident rays will exit the lens at phase $\pi/2$ in the oscillation cycle. In practice, however, the incident light source with which a GRIN lens is coupled is likely to have a non-negligible diameter relative to that of the lens; this instead produces a light cone with some defined opening angle. The complete specifications of the collimator's GRIN lens[6] is given in Table 12.3.

In the fibre–lens coupled system, the opening angle is set by the beam diameter projected by the fibre at the lens face. The upper limit on the opening angle can be increased either by modifying the specifications of the lens or fibre—increasing the size of the fibre core or decreasing the lens' pitch[7]—or by introducing a separation between the front face of the lens and the fibre ferrule. Coupled with 200 μm fibres, the beam projected by the collimator's lens has a typical half opening angle of $\sim$3.5° in air at full width half-maximum intensity.

[6] Edmund Optics® #64-519

[7] Gradient-index lenses of 0.23 pitch are readily available.

Parameter	Value
Substrate	Aluminosilicate glass embedded with silver ions
Diameter	1.80 mm
Length	4.31 mm
Refractive index at centre	1.629
Gradient constant	0.364 mm
Working distance	0 mm
Numerical aperture	0.52
Wavelength range	400 nm to 1600 nm
Effective focal length	1.69 mm

Table 12.3 Specifications of the collimator's gradient index lens.

The lens and fibre are coupled together in the assembly with the *lens-mount* piece. The lens is held within a 1.81 mm diameter, 2 mm deep cavity in which the lens is restricted from moving downstream by a 10 mm diameter, 3 mm thick BK7 window[8]. The fibre is mounted to the lens-mount via an FC adapter threaded at 3/8"-24. With both the lens and fibre in place, the spring-loaded ferrule of the fibre keeps the lens and fibre in constant contact, thus restricting their relative positions. In addition, a 0.5 mm thick aperture of 2.35 mm inner diameter and 10 mm outer diameter is glued on top of the window. The presence of this aperture aids to block out high-angle light introduced by scattering and internal reflections within the lens and window. The inner diameter is large enough that the direct collimated light emerging from the lens, which has a beam diameter about as large as the lens diameter (1.8 mm) as it exits the window, is well accepted by the aperture. The aperture and inner walls of the lens-mount are accordingly painted black to reduce reflections.

The lens-mount is fixed via three grub screws within a 81 mm long cylindrical *sleeve* of 15 mm inner diameter and 22 mm outer diameter. The apertured *cap* seals the optics when affixed to the end of the sleeve. The inner of the sleeve and the cap are painted matte black to absorb reflections. Inside of the cap is a 15 mm diameter, 2 mm thick sapphire glass (Al_2O_3) window[9]. The primary role of the cap's aperture, however, is not to tune the opening-angle further, as the opening-angle of the beam emitted by the lens–fibre coupled system is already satisfactory. Rather, it is to filter out a high-angle (40°), low-intensity ($\sim$1%), diffuse component that was observed of the light emitted by the lens-mount. The field of the collimated light that is accepted by the aperture is set by the diameter of the aperture and the separation between the lens-mount and aperture: an increasing aperture diameter

[8] Knight Optical®, WQK1000.

[9] Edmund Optical® #43-367.

and decreasing distance between the lens-mount and cap evidently achieves an equivalent beam angle with the inverse configuration within certain bounds. However, the distance between lens-mount should be sufficiently long to allow for high-angle light, initially found close to the optical axis, to propagate far enough to become sufficiently radially separated from the collimated light as it approaches the aperture. Increasing the lens–cap distance, however, amplifies the effects of misalignment of the lens-mount. Hence, a compromise between the aperture size and separation was decided upon. The aperture size was fixed at 1.8 mm, setting a lens–aperture separation of $\sim$15 mm to accept exactly the diameter of the direct light. Unlike the aperture built into the lens–mount piece, however, the cap's aperture accepts a region of the direct, collimated light emitted from the lens. Ideally, this region should be small enough to reject indirect light and large enough to preserve as much of the direct light as possible. However, the persistent problem of misalignment of the beam with respect to the aperture causes the aperture to accept an off-axis portion of the direct light. Hence, if the accepted region is too large, then misalignment will cause a beam with an almond-shaped cross section. The scale of misalignment between each unit was assessed to motivate the field accepted by the aperture. It was found that, in order to select a symmetrical portion of the beam reliably for each collimator, the acceptance of the aperture should be such that a beam with half opening-angle of 1.8° at FWHM intensity is projected. In this configuration, around half of the field of the direct light is rejected by the cap's aperture. The beam diameter projected on the opposite wall of the filled tank becomes 2–3 photo-sensors for the apertured configuration, in comparison to 4–5 for the lens–mount piece in isolation. This certainly falls short of the original target of 5 photo-sensors, however it was decided that good circularity of the beam spot and minimal emission of background light were more favourable characteristics of the collimator for its physics purposes.

With the collimator assembled, the cavity at the back of the sleeve is potted with epoxy resin. As the epoxy cures, a strain relief is implanted to support the fibre. A completed collimator is shown in Fig. 12.9b. As the position of the lens-mount within the sleeve was not known a priori, the length of the sleeve was taken to be longer than required. Due to time constraints, the sleeve was not trimmed once the optimum lens-mount–aperture separations were derived; instead, the large volume for potting was recognised as a good support for the fibre. In spite of this, the length of the assembly was not found to be an issue during deployment.

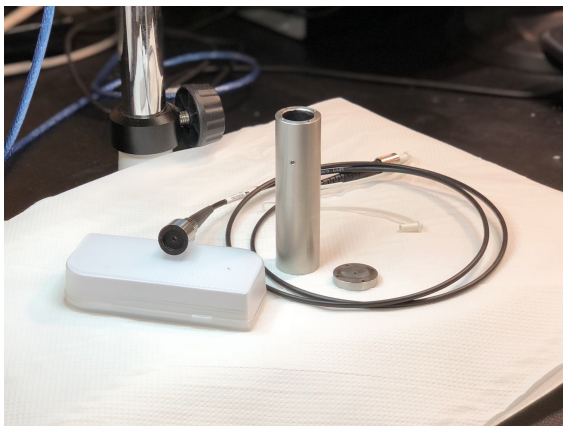

(a) The lens-mount, sleeve and cap.

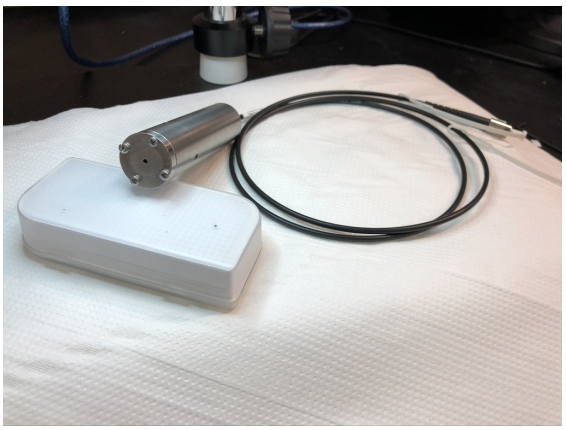

(b) An assembled collimator.

Fig. 12.9 A disassembled (Fig. 12.9a) and an assembled (Fig. 12.9b) collimator.

12.5.2 Simulation

The physics of the lens–fibre coupled system was simulated using two independent simulation packages: *i)* a ray propagation simulation built in-house and *ii)* the COMSOL®[10] simulation software. In the in-house simulation, the propagation of rays through the lens is solved numerically [153]; the COMSOL® simulation conducts a finite element analysis. Some example ray traces through the gradient-index lens as calculated by the former simulation are shown in Fig. 12.10. The simulations agree qualitatively about the resultant beam characteristics when the optics are systematically misaligned.

In both simulations, the fibre light is modelled with a source with a uniform spatial distribution for $r < r_{\text{core}}$, where r_{core} is the fibre's core radius, and a uniform polar angle distribution uniform with respect to $\cos\theta$ for $\theta < \arcsin(NA)$. Translational misalignments that are directionally on-axis project a directionally off-axis beam. As rays begin to take helical paths through the lens, translating the fibre along the $\hat{x}$-axis produces a beam at a non-zero angle with respect to the $\hat{y}$-axis and vice versa. The beam shape and position are largely insensitive to small ($< 10°$) directional misalignments that are positionally on-axis—the lens, in fact, corrects for such misalignments. Introducing a separation between the lens and fibre not only increases the beam size but smears the beam shape.

Simulation studies also found that if internal reflections should occur within the lens, a wide-angle diffuse source is produced. A reflection upon the sides of the lens happens when the incident position and angle of an incident ray exceeds the acceptance of the lens. For a point source, reflections will surely be produced should the numerical aperture of the source exceed that of the lens. Given sufficient positional or direction misalignments, the same effect

[10] Release 5.4, with the Ray Optics Module.

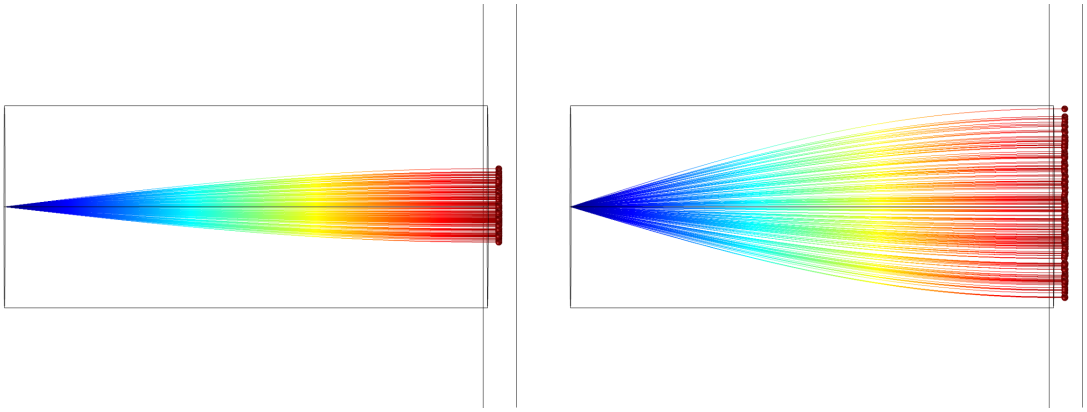

(a) 0.22 NA. (b) 0.5 NA.

Fig. 12.10 Traces of rays through the collimator's GRIN lens from a point-like source with a numerical aperture of 0.22 (Fig. 12.10a) and 0.5 (Fig. 12.10b).

may be observed for a source with a numerical aperture within the acceptance of the lens. With aligned extended sources, whether reflections are produced or not is contingent also on the source size. Simulations ran with the extended source model of the fibres show that rays emerging at extreme radii and angles are likely to be beyond the acceptance of the lens. Such reflections contribute to a diffuse background up to $\sim 40°$ of about 1% intensity relative to that of the direct light. Light emerging from the cladding of the fibre, fibre–lens misalignment and imperfections within the lens evidently increase the rate of internal reflections further. Such a problem is not found in equivalent simulations of a fibre with a lower numerical aperture of 0.22. These predictions were later confirmed in the laboratory.

12.5.3 Test stand

In a word, the test stand injects a collimator with light and scans the cross section of the resultant beam. A diagram and photograph of the test stand are shown in Fig. 12.11. The collimator to be tested is mounted onto a stand and coupled to a blue LED with an optical fibre. Pointing along the direction opposite to the beam axis is a CMOS camera[11] which is mounted to an array of motorised linear stages that may step along and perpendicular to the beam axis. All of the equipment—with the exception of the high-voltage supply that powers the LED—is kept within a light-tight box.

The data-acquisition software used during the development of the SK optics can either step across horizontally or longitudinally to the beam axis in a single run—the data collected

[11] The camera model is Thorlabs®DCC1545M which uses a model Aptina®MT9M001 sensor; the sensor is monochrome and of 8-bit colour depth and of $1/2$" size and 1280 by 1024 pixel resolution.

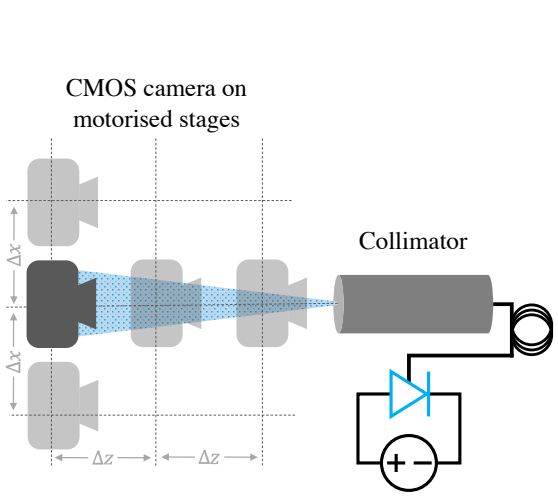

(a) Diagram.

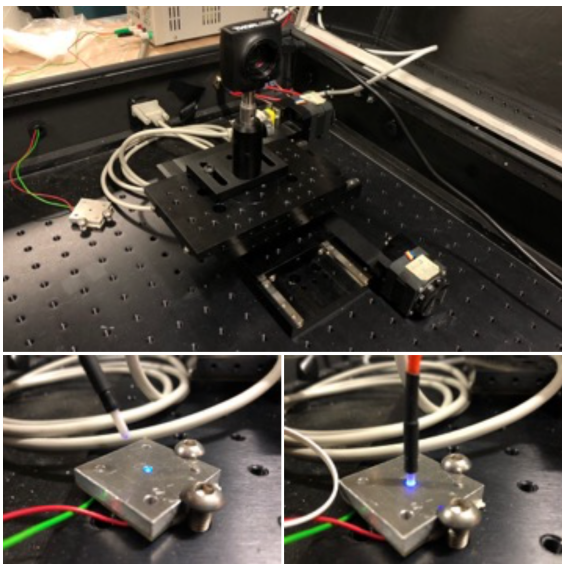

(b) Photograph.

Fig. 12.11 Diagram (Fig. 12.11a) and photograph (Fig. 12.11b) of the collimator test stand.

by each of these two scan types allows a particular assessment of a collimator's quality. The data collected by the longitudinal scans is used for *i)* visual checks of the circularity of and imperfections in the beam at low angles ($\lesssim 5°$), *ii)* checks of linear growth of the beam over the scan range and *iii)* measurement of the effective opening angle of the beam's light cone; the transverse scans are explicitly purposed to check for the presence of high-angle ($\lesssim 40°$) light.

The data is processed and analysed with a dedicated Python package. The images are subject to the following processing chain in order to measure the beam cross sections: *i)* subtraction of the background, the background image is an average of two images collected before and after a run with the light source powered off; *ii)* filtering of the image, a median filter[12] with a circular brush of 10 pixel radius is applied in order to reduce apparent imperfections in the beam that might interfere with the further processing steps; *iii)* binarisation of the image at half-maximum intensity; *iv)* application of the Canny edge-detection algorithm[13] [154], in order to extract the perimeter of a beam cross section; *v)* masking of pixels within six pixels of the image borders, removing bright pixels which tend to appear at the extremities of the sensor that might interfere with the next step; *vi)* fitting of the beam

[12]The median filter is well-recognised for preserving image features.

[13]The Canny edge-detection algorithm has three parameters: the width of the Gaussian filter applied, which was set at 0.1, and the lower and upper bounds of the hysteresis thresholding applied, which were set at 10% and 20% of the maximum pixel value of the image.

edges with a circle, to measure the radius of the beam at half-maximum intensity and the coordinates of the beam centre.

With all cross sections measured, the half opening angle of the beam's light cone at half-maximum intensity is derived by a linear fit to their radii against their corresponding longitudinal camera position. The beam radius grows linearly with distance d as $(d\tan\theta + r_0)$ where r_0 is the beam radius at $d = 0$. The half opening angle is hence $\arctan m$, where m is the gradient derived by the linear fit. The test stand records only the distance displaced by the motorised stage, the distance between the camera and test collimator was unrecorded and free to vary between runs; this is reflected in the variation of the beam radius at the origin between runs shown in Fig. 12.19. Instead, the distance between the camera and the virtual origin of the beam's light cone—the value of d where $r = 0$—is used in order to compare the beam shape at equivalent distances between runs. The distance between the virtual origin of the light cone is calculated as $d - (c/m)$, where c is the intercept of the beam angle fit. The distributions are either calculated from data in the near-field, from images captured at the closest recorded distance to the collimator, or the far-field, from the furthest. Given that the beam-angle varies between collimators, so does the set of distances from the virtual origin sampled in each run. Therefore, the closest set of data points are presented. In spite of this, the shape of the distributions do not inherently vary with the distance from the collimator. A bias in the distributions due to missing data begins to take effect, however, as the beam diameter becomes larger than the sensor. The azimuthal position distributions are particularly prone to this effect and hence are presented in the near-field.

12.5.4 Production

Not including pressure tests, the turn-around time for commissioning a collimator is 2–3 days. The drying of the epoxy and paint are the largest bottlenecks, so certainly production can be conducted in small batches of up to five units. In total, six collimators were produced and shipped to SK, with one reserved for contingency. Upon arrival at SK, the seal of unit 4 appeared to be compromised; hence, the production yield was ⅚.

Before assembly, all components are cleaned with solvent, then various surfaces are painted matte black to absorb reflections: the inside of the tube, the inner-face of the cap and all surfaces on the lens mount facing the optical volume in the tube. After drying, the Schott glass window and the 2.35 mm aperture can be glued into the lens-mount, and the Sapphire glass window can be glued into the cap. After those have dried, a clean lens is mounted into the holder and the optical performance of the lens–fibre coupled system is measured. If the performance is satisfactory, then the fibre can be glued to the lens-mount assembly. A photograph of the assembly process for some of the lens-mounts are given in Fig. 12.12a.

Next, alignment of the complete assembly is attempted. If good alignment is possible, the separation between the lens-mount and the cap is tuned for the target half opening angle in air. This is achieved by interpolating between measurements made of the opening angle at two lens–cap separations. These separations are given in Table 12.4. This configuration is validated by a final opening angle measurement before glueing the cap and potting the cavity on the fibre side of the assembly with epoxy. A stain-relief is fitted to the fibre as the epoxy cures. A photograph of the newly assembled collimator is given in Fig. 12.12b. After drying, an opening angle measurement is performed; then after at least 24 hours of pressure-testing, the opening angle measurement is performed for a final time. If the optical performance is satisfactory and there is no sign of leakage, then the collimator is eligible for installation.

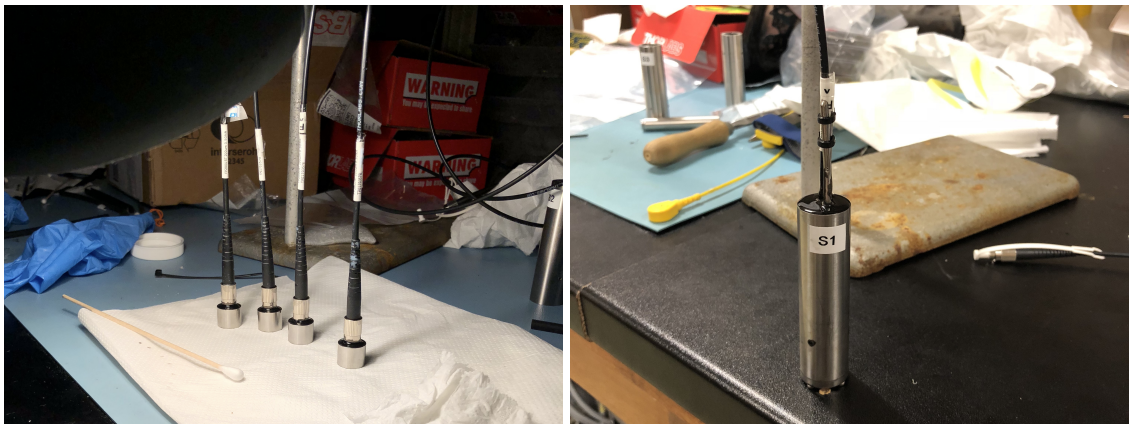

(a) Lens mounts. (b) A potted collimator assembly.

Fig. 12.12 Photographs taken of two stages in the production of the collimators for the Super-Kamiokande light injection system. Fig. 12.12a shows four lens-mounts after glueing in the fibre and Fig. 12.12b shows a newly-completed collimator after potting.

12.5.5 Laboratory measurements

The opening-angle measurements of the six collimators produced are given in Table 12.4.

As mentioned, the lens-mounts were measured after being assembled[14]—some images from these scans are shown in Fig. 12.13. Newton's rings come into focus as the distance along the beam axis increases, likely this arises from the point of contact between the lens face and the glass window. As seen in Fig. 12.16, this effect is particularly noticeable in the extremities of the radial and polar angle distributions of unit 5; the other units are consistent. The half opening angles at full width half-maximum are given in Table 12.4 and Fig. 12.14—the mean of the batch is $(3.46 \pm 0.40)°$. As the lens-mount assembly does nothing other than

[14] With the exception of unit 6, which was not measured.

Unit index	Lens–cap separation [mm]	Lens-mount opening-angle [°]	Assembled opening-angle [°]	Current status
1	18.36	3.42 ± 0.05	1.75 ± 0.03	Installed at B4.
2	17.52	3.47 ± 0.05	1.61 ± 0.04	Installed at B1.
3	17.53	3.45 ± 0.05	1.67 ± 0.05	Installed at B3.
4	17.46	3.46 ± 0.05	1.77 ± 0.04	Seal compromised, returned to laboratory.
5	17.37	3.49 ± 0.05	1.56 ± 0.03	Installed at B2.
6	16.09	—	1.56 ± 0.03	Installed at B5.

Table 12.4 Details of the collimator units produced for the Super-Kamiokande injector installation. Details of the installation points may be found in Table 13.1.

couple the lens and fibre, this small variation of angles can be assumed to be a product of the manufacturing tolerances of the core diameter of the coupled fibre and the length of the lens. As demonstrated in Fig. 12.16, the spots imaged by all units are sufficiently circular.

The collimators were measured before and after at least 24 hours at 10 bar within the pressure vessel; no changes in performance were observed, thus only the data collected after the pressure test will be presented. Some images of these scans can be seen Fig. 12.18. The increase in the irregularity of the beam projected by some units compared to those of their lens-mounts should be attributed to the quality of the cap's aperture and the alignment of the lens-mount's beam with respect to the aperture. It should be acknowledged, however, that these irregularities are unresolvable with SK's spatial resolution when the beam is projected across the diameter of the tank. For such a small beam diameter, the shape of the hit pattern recorded by the detector is dominated by spatial aliasing; so rather the resulting hit pattern is driven mostly by the beam position relative to the photo-sensors. Their deviations from a circular beam are reflected in the azimuthal angle distributions given in Fig. 12.22. Unit 2 and unit 5 are perhaps the poorest in this respect. The relative noisiness observed of unit 5 is likely due to the outside face of the window in the cap having not been cleaned properly before measurement. Due to the small size of the aperture, it is an arduous task to attempt to rid the outside face of the window of dust or smudges after being adhered to the cap. Likely this is of no issue, however, as this area would be cleaned when submerged within the tank. The half opening angles at full width half-maximum intensity of the fully assembled collimators are given in Fig. 12.19—the average of the batch is $(1.65 \pm 0.12)°$. The variation of opening angles can be attributed to the deviation of lens-mount–cap separations from their value necessary to achieve 1.8°. Likely their separations were set with a good degree of precision, rather this is more likely this variation is driven by the uncertainty on their

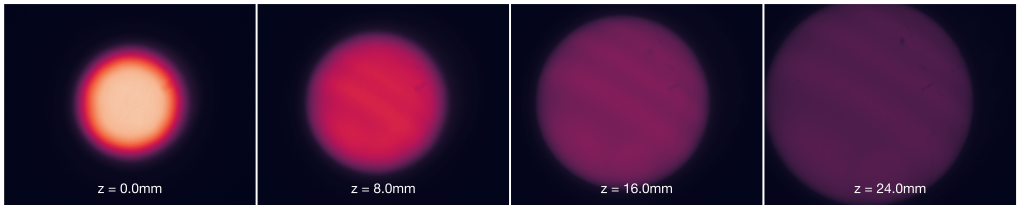

(a) Unit 1 (B4).

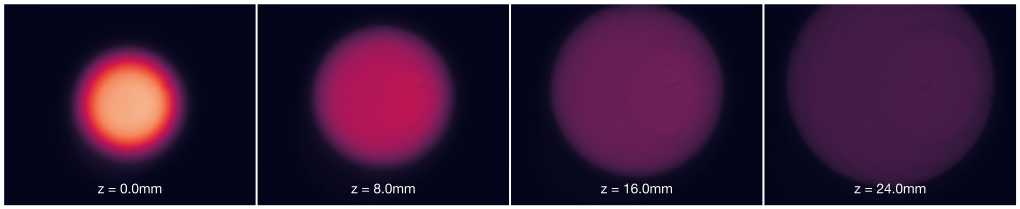

(b) Unit 2 (B1).

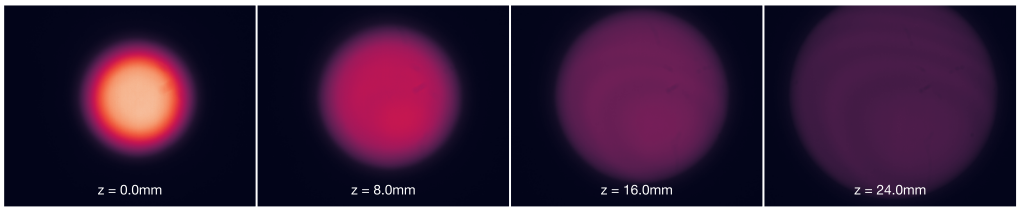

(c) Unit 3 (B3).

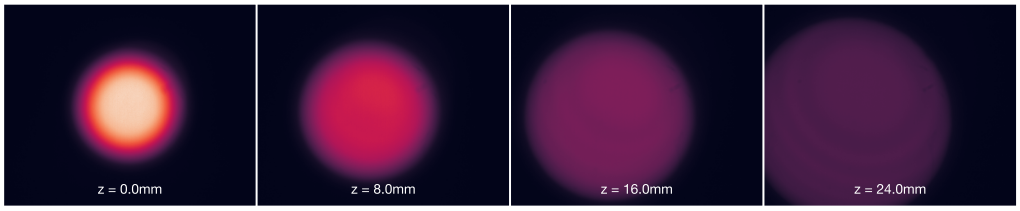

(d) Unit 4.

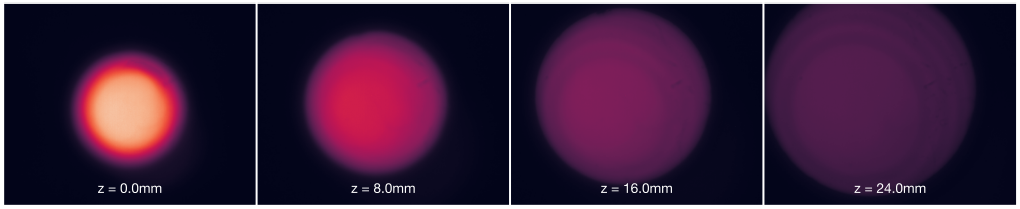

(e) Unit 5 (B2).

Fig. 12.13 Background-subtract beam images taken at increasing beam distances of the lens-mounts as measured in the laboratory test stand.

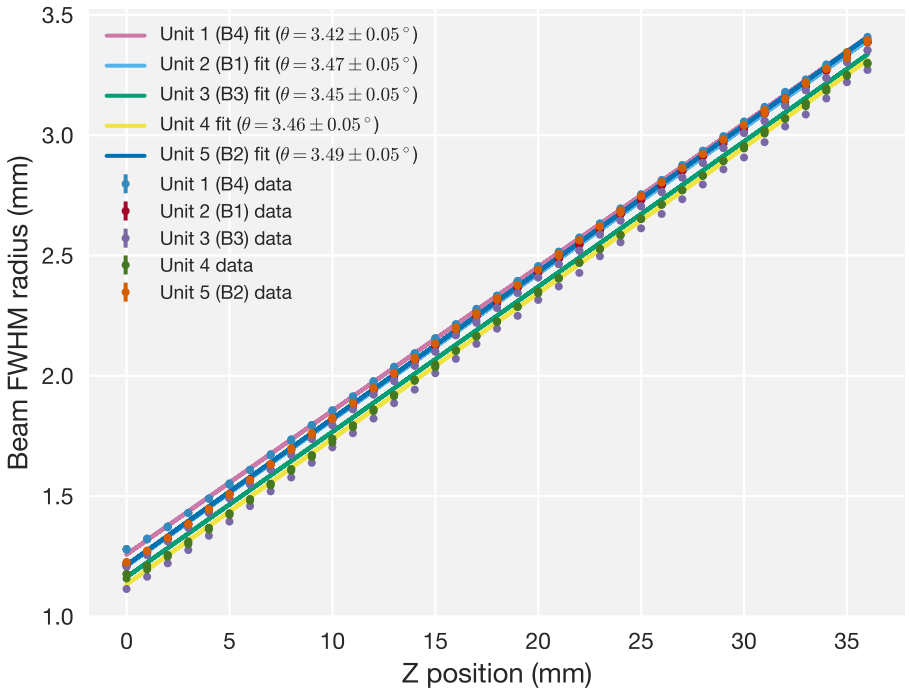

Fig. 12.14 Variation of beam widths with longitudinal distance of lens-mounts of the collimator assembly measured in the laboratory test stand.

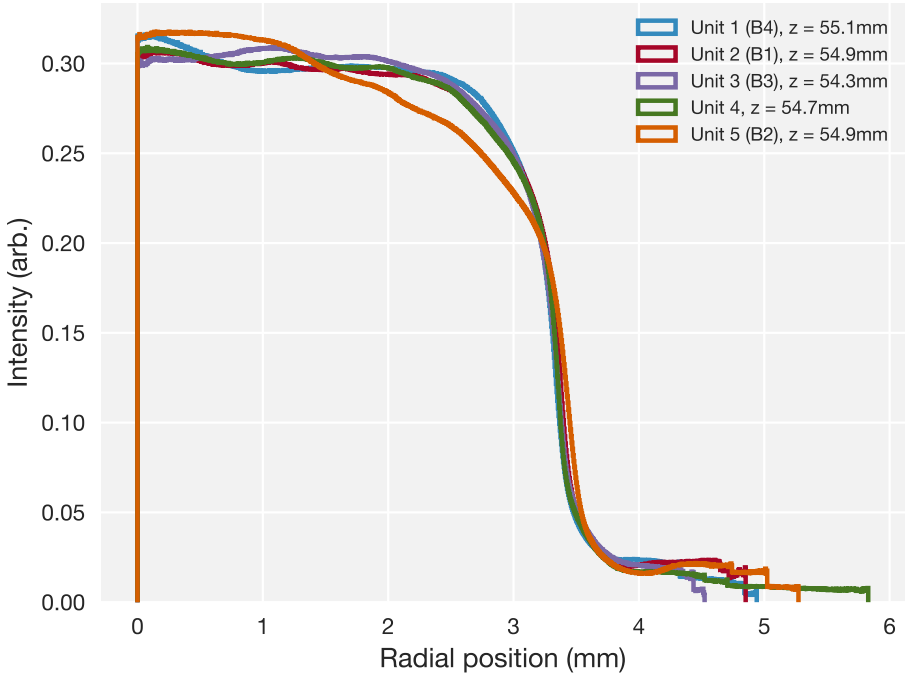

Fig. 12.15 Radial intensity distributions in the far-field of lens-mounts of the collimator assembly as measured in the laboratory test stand.

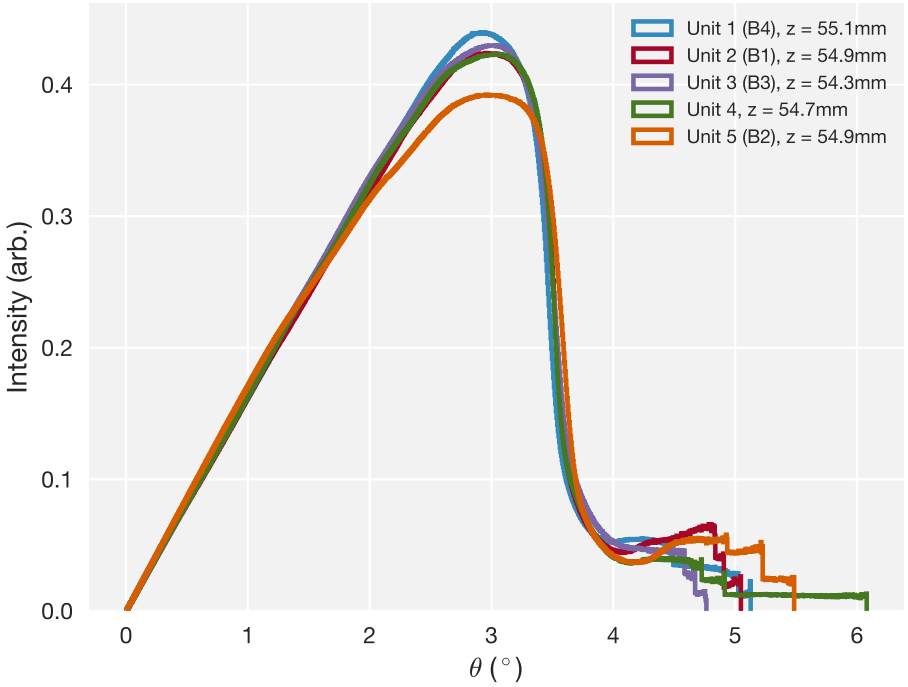

Fig. 12.16 Polar angle intensity distributions in the far-field of lens-mounts of the collimator assembly measured in the laboratory test stand.

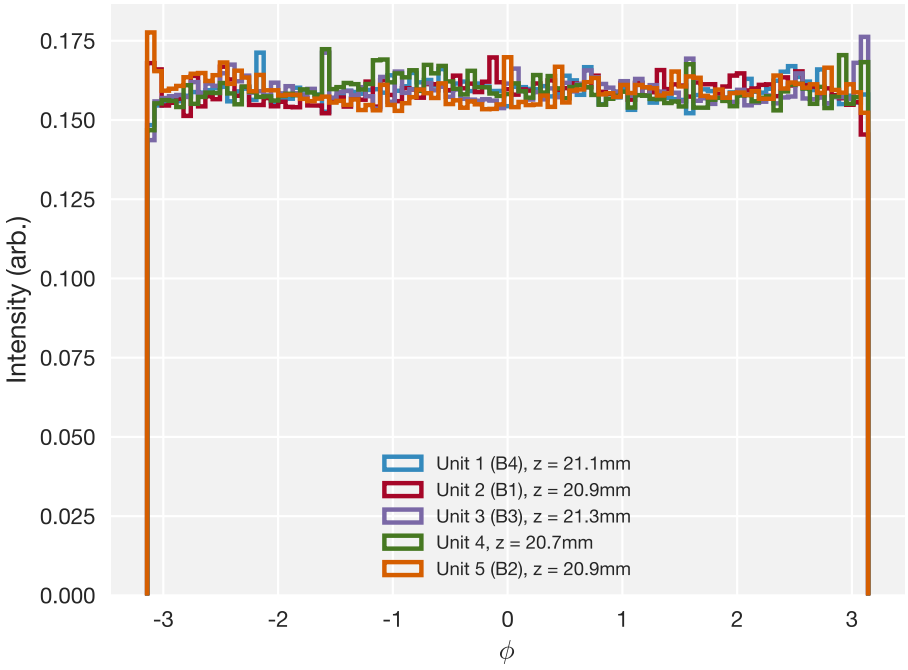

Fig. 12.17 Azimuthal angle intensity distributions in the near-field of lens-mounts of the collimator assembly as measured in the laboratory test stand.

calculated target separations. As mentioned, only two data points were used to calculate their target separations—in retrospect, the calculated separation should have then been validated with a measurement; however, due to time constraints, this was overlooked. The far-field polar angle distributions are given in Fig. 12.21—reflecting the opening angle measurements, it can be seen the spread is tighter for the units installed at B2 and B5 the remainder of the now installed units show consistent shapes. The opening angle of unit 4 is the largest of the batch, but perhaps the tail of its distribution is more prominent than expected given this opening angle. The far-field radial distributions as averaged over all azimuthal angles are given in Fig. 12.20. Units installed at B1 through to B4 achieve uniformity in intensity in the radial region close to the beam axis; the same cannot be said for B5. The relatively tighter spread for units installed at B2 and B5 may also be observed in the radial distribution. All this said, the production can certainly be considered successful, especially given the development time scale.

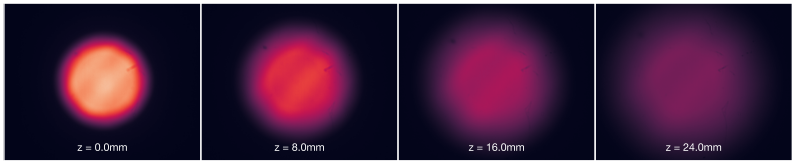

(a) Unit 1 (B4).

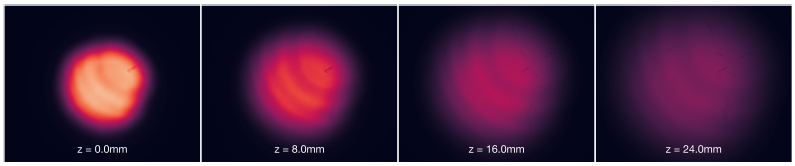

(b) Unit 2 (B1).

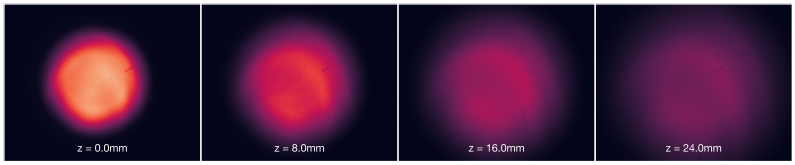

(c) Unit 3 (B3).

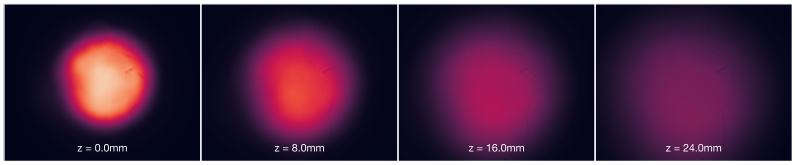

(d) Unit 4.

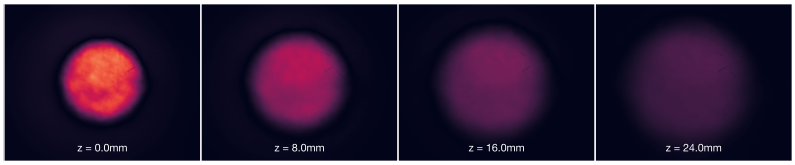

(e) Unit 5 (B2).

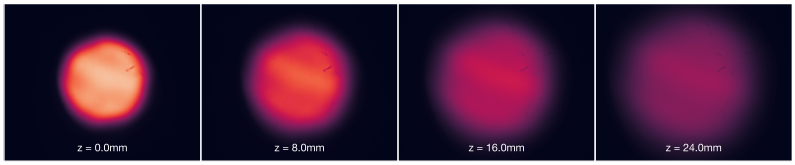

(f) Unit 6 (B5).

Fig. 12.18 Beam profile images taken of each collimator produced for installation in Super-Kamiokande as measured at the collimator test stand. In each subfigure, the images are displayed in order of the beam distance z from the collimator at which they were recorded.

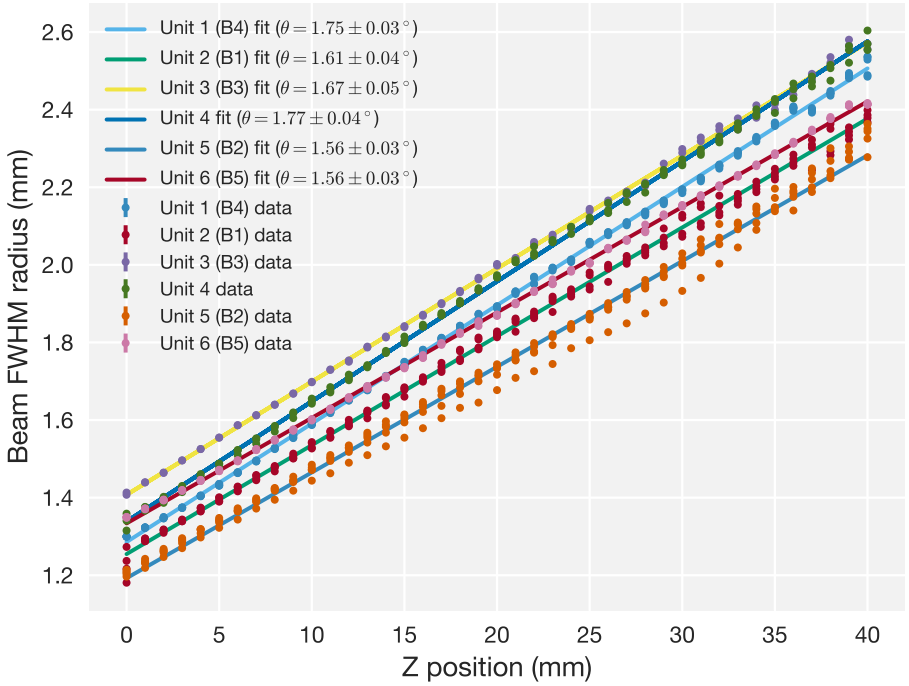

Fig. 12.19 Variation of beam widths with longitudinal distance of all collimator optics installed inside of Super-Kamiokande as measured in the laboratory test stand.

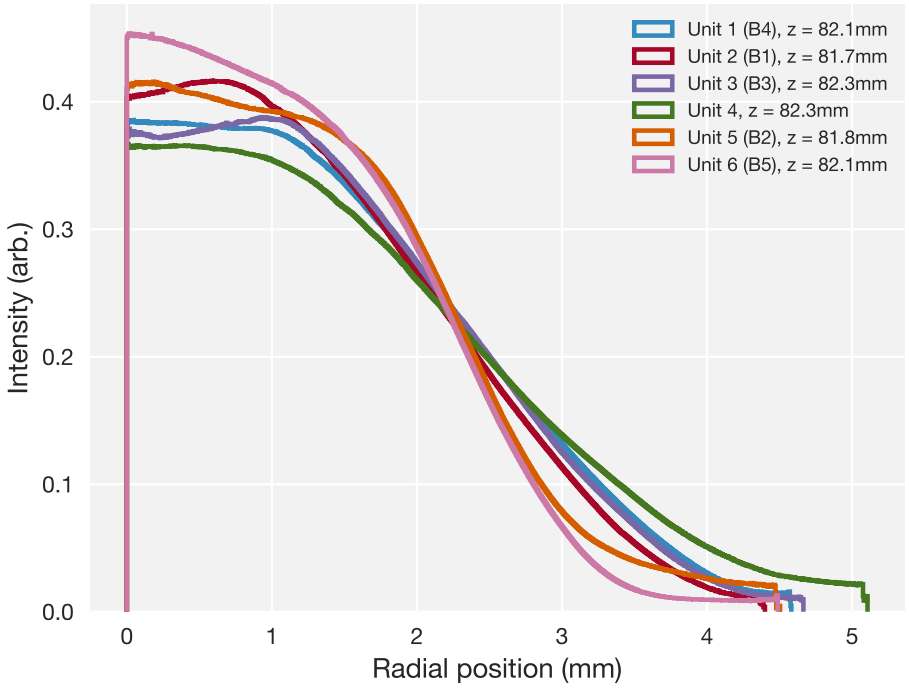

Fig. 12.20 Radial intensity distributions in the far-field of all collimator optics installed inside of Super-Kamiokande as measured in the laboratory test stand.

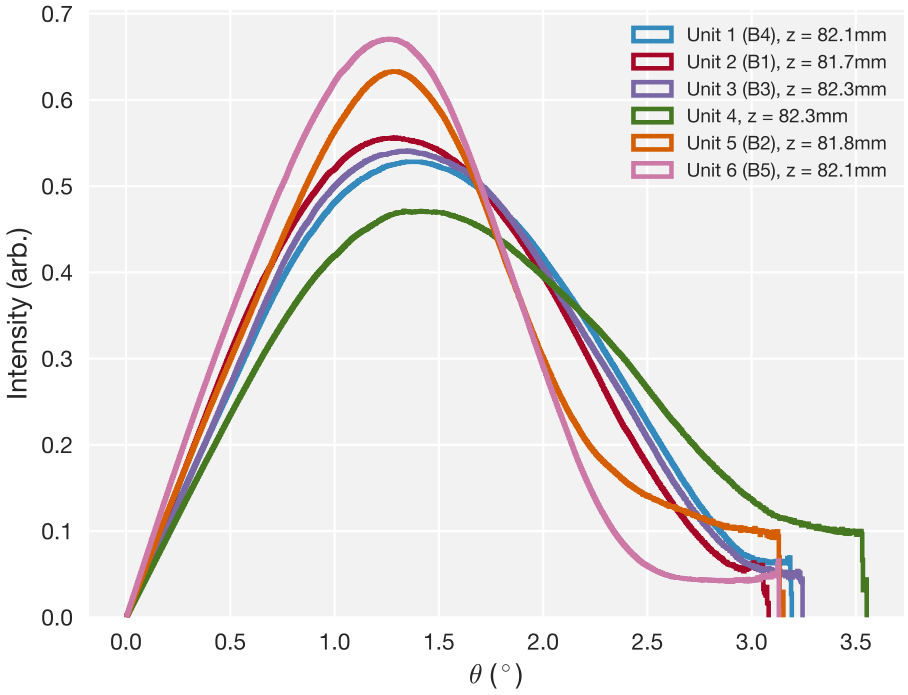

Fig. 12.21 Polar angle intensity distributions in the far-field of all collimator optics installed inside of Super-Kamiokande as measured in the laboratory test stand.

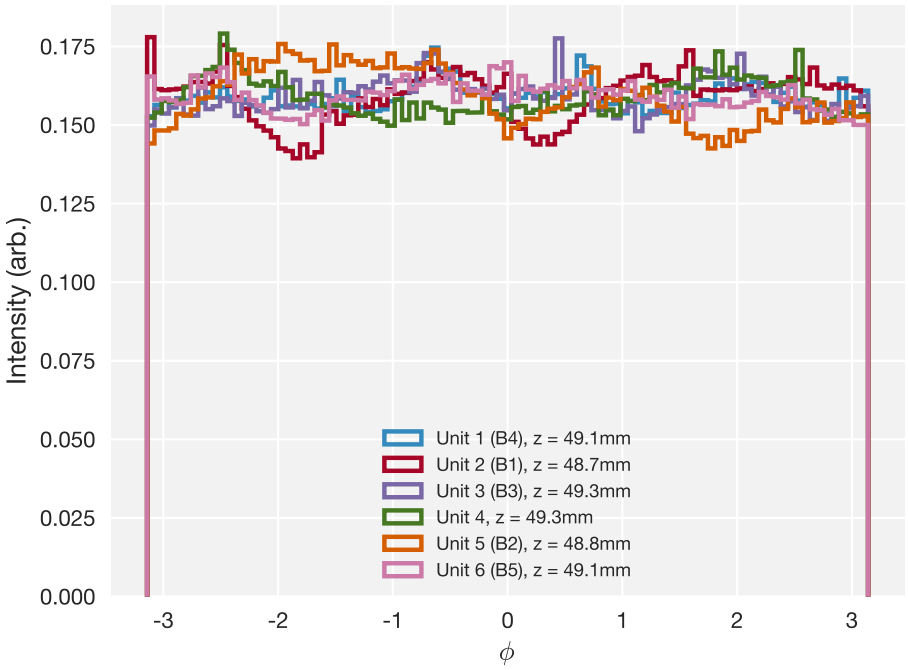

Fig. 12.22 Azimuthal angle intensity distributions in the near-field of all collimator optics installed inside of Super-Kamiokande as measured in the laboratory test stand.

CHAPTER 13

DEPLOYMENT OF CALIBRATION OPTICS WITHIN THE SUPER-KAMIOKANDE DETECTOR

The calibration optics described in the preceding chapter were installed within the Super-Kamiokande detector during the summer of 2018. In this chapter, in-situ measurements of the signal from each optic are reported. After the installation, in 2019, signal checks were conducted on two separate occasions. Those data are analysed in Section 13.3. Before all of that, however, a preliminary deployment of the optics was carried out in Super-Kamiokande in the January of 2018. The key findings of that deployment will be covered in Section 13.2. But first, the treatment of the detector data will be described.

13.1 Calibration Data

Each calibration run corresponds to one optic at one injection point being injected with a fixed wavelength, pulse length and intensity. The positions and targets of all available injection points in the inner detector are listed in Table 13.1 and visualised in Fig. 13.1. The treatment of calibration data is the same for all runs; the event selection and corrections applied will be described.

	Injector (m)			Target (m)		
	$\hat{x}$	$\hat{y}$	$\hat{z}$	$\hat{x}$	$\hat{y}$	$\hat{z}$
Old Top	-0.4	7.0	18.1	-0.4	7.0	-18.1
New Top	-0.7	-7.8	18.0	-0.2	-6.9	-18.1
B1	14.9	7.7	13.0	-14.7	-8.3	13.1
B2	14.9	7.7	6.7	-14.5	-8.6	6.7
B3	14.9	7.7	-1.1	-14.9	-7.9	-1.7
B4	14.9	7.7	-6.8	-14.6	-8.5	-6.4
B5	14.9	7.7	-13.1	-14.3	-9.0	-13.0
Bottom	-0.7	7.8	-18.0	-0.7	7.8	18.0

Table 13.1 Coordinates of injection points and corresponding target positions within the Super-Kamiokande inner detector.

13.1.1 Event selection

The selected hits are those *in-time* and collected on the light injection trigger. In-time hits are defined as detector hits within a window placed on the time-of-flight subtracted hit times. If the run was taken with a monitor photo-sensor, then the detector hit times are made relative to the hit timing reported by the monitor in the same event; otherwise, the hit time is simply the photo-sensor timing. The expected time-of-flight of a photon emitted at the injection point and collected at photo-sensor i is calculated as

$$\Delta t_i = \frac{n}{c_0}|\vec{r}_i - \vec{r}_0| \tag{13.1}$$

where n is the refractive index of water (1.333), c_0 is the speed of light in vacuum, $\vec{r}_i$ is the position of photo-sensor i and $\vec{r}_0$ is the position of the injector. The timing window was adjusted by hand for each run in order to select as much direct light and as little scattered light as possible in order to assess the performance of the optics.

A relatively large number of hits can be found close to the injector for some optics, this is either due to indirect light that reflects from the collimator tube or back-scattering. Hence, photo-sensors within a $4\,\mathrm{m} \times 4\,\mathrm{m} \times 4\,\mathrm{m}$ cube centred on the injector are removed from the analysis, so as to not saturate the hit maps.

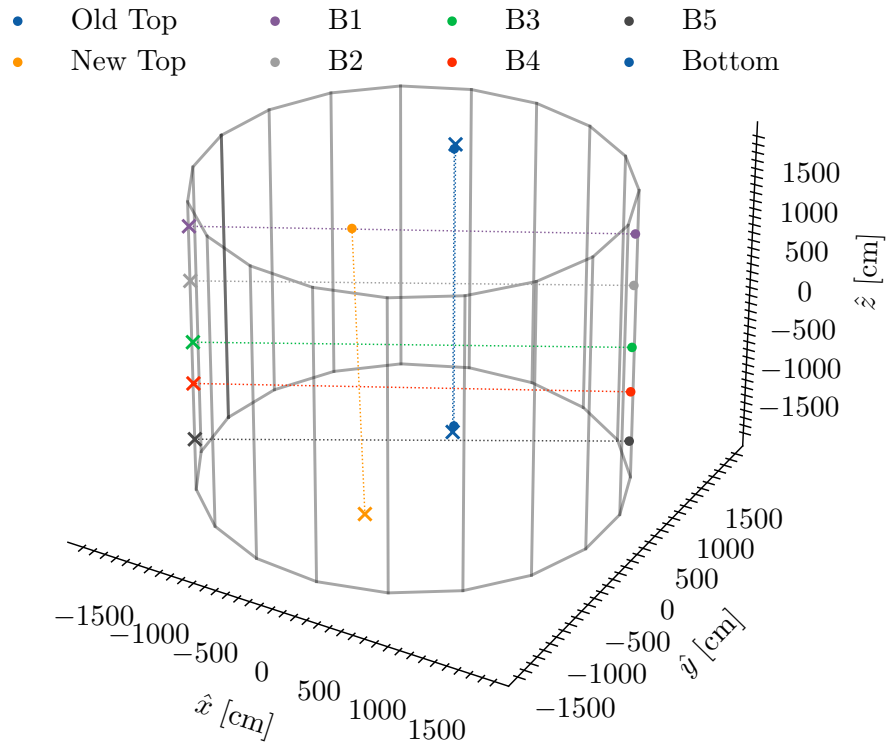

Fig. 13.1 Locations and pointing directions of injection points in the Super-Kamiokande inner detector. Circles mark the injector positions and crosses mark the corresponding target position.

13.1.2 Corrections

A number of corrections are applied to the hit maps to account either for geometrical or detector effects. Each correction j is implemented as weight w_i^j applied to the number of hits collected at each photo-sensor i, whereby the total correction is calculated as

$$w_i = \prod_j w_i^j. \tag{13.2}$$

The applied corrections are as follows:

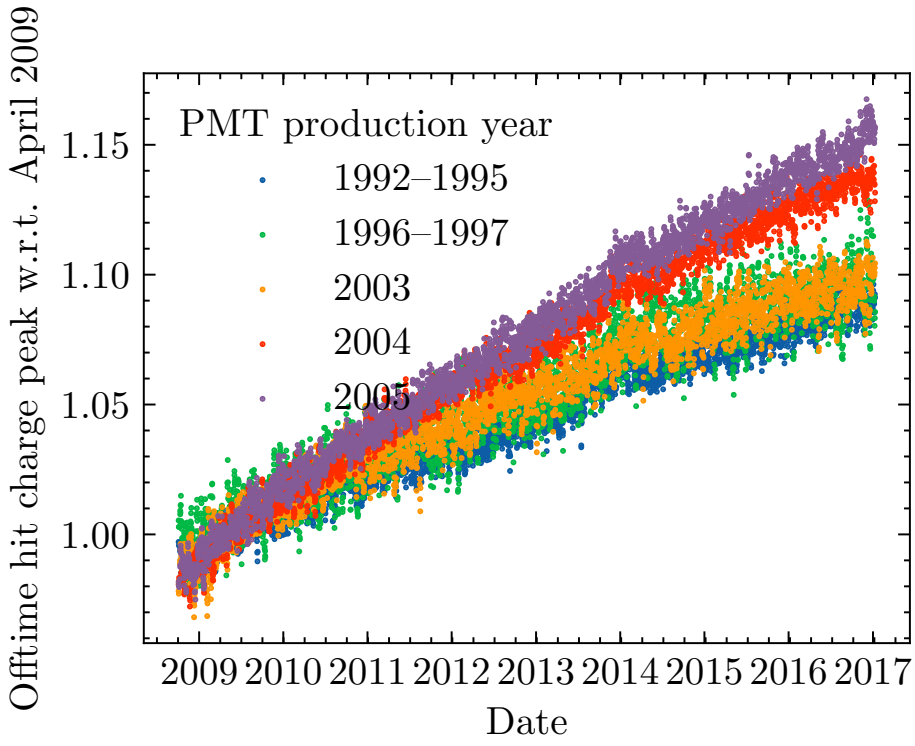

Fig. 13.2 Time variation of gains of Super-Kamiokande inner detector photo-sensors by their production year.

Gain

As defective PMTs have been replaced over the years, a variety of older and newer PMTs are installed in the tank. Given that the gains of the photo-sensors are observed to increase around 2% per year, a standard correction is applied to account for the relative differences in gain between the PMTs.

The relative increase in gain since April 2009 is taken as the increase in the off-time hit-charge peak. This quantity is recorded as a function of SK run number for five different batches of PMTs, these five batches correspond to PMTs manufactured between 1992–1995 and 1996–1997 and in 2003, 2004 and 2005. The time-variation of gains by photo-sensor batch is shown in Fig. 13.2. These gains are used to calculate a correction to the number of hits recorded at each photo-sensor. The number of hits recorded by a photo-sensor belonging to batch i at a particular date is corrected using the function

$$w_i = \frac{1}{1 + C\left(\frac{G_i}{G_i'} - 1\right)} \tag{13.3}$$

where G_i is the gain correction for PMT batch i at that date and G_i' is the gain correction for batch i at a reference date of October 2008. C is a factor—equal to 0.226—describing the fractional increase (22.6%) in the number of hits expected of a photo-sensor whose gain increases by 100%. The minority of PMTs for which the production year is unknown ($\sim$500) are not corrected. The gains at the time of these calibration runs—January 2018, February 2019 and July 2019—were not yet tabled and so were calculated from linear extrapolations of the previous gains. The scale of this correction is around 5% at maximum at these dates.

Solid angle

The solid angle subtended by a PMT from a particular position in the detector can be approximated as

$$w_i = \frac{2\left(a^2 + d_i^2\right)}{a^2} \tag{13.4}$$

where a is the radius of the photo-sensor and d is the distance from the injector to photo-sensor with index i [155].

Angular acceptance

The photo-sensors appear at a diversity of orientations with respect to the source. Given the photo-sensors are hemispherical, a correction to account for the visible area of each photo-sensor can be approximated as

$$w_i = \frac{1}{1 - \frac{|\theta_i|}{\pi}} \tag{13.5}$$

where θ_i is the angle subtending the normal of photo-sensor i and the direction of the injector.

Attenuation

The combined effects of optical scattering and absorption within the detector medium cause the intensity of light to dim according to the path length traversed through the tank. The correction is calculated as

$$w_i = e^{d_i/\lambda} \tag{13.6}$$

where d_i is the distance from photo-sensor i to the injector and λ is equal to 103.4 m—the average attenuation length measured of 450 nm light during SK-IV.

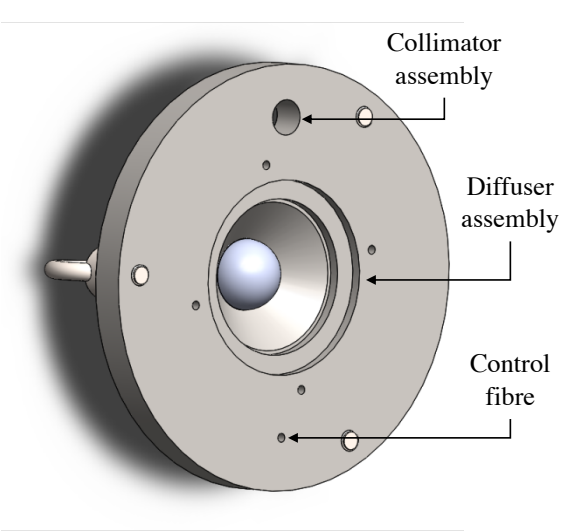

(a) The optical plate.

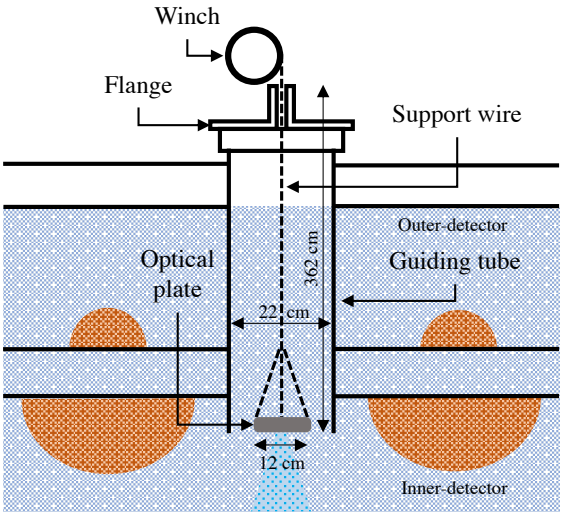

(b) Deployment of optical plate through the top calibration (not to scale).

Fig. 13.3 The optical plate (Fig. 13.3a) and its lowering through the *old top* calibration port (Fig. 13.3a) during the preliminary light injection deployment in Super-Kamiokande.

13.2 Preliminary Deployment

An in-situ test of the optics was conducted in January of 2019. As depicted in Fig. 13.3b, a diffuser, collimator and control fibre were fixed to a dedicated mount and lowered on a winch into the inner detector through the *old top* calibration port. The optical plate sat flush with the bottom of the port's guiding tube, targeting an off-centre position on the bottom of the tank.

13.2.1 Optics

Early designs of the optics were used. The collimator consisted of a gradient-index lens of the same specification as previously mentioned upstream of a 40 mm long, 1.70 mm diameter tube. The design of the assembly can be seen in Fig. 13.5. The 2.5 mm ferrule of the optical fibre was glued directly to the assembly. It was later found that the tube had little to no collimation power. Simulations of this assembly showed that the tube was neither long nor narrow enough to interfere with the beam when the optics are perfectly aligned. Rather, the tube promotes reflections upon one side of the tube walls if the beam projected by the lens is not aligned along the tube axis. Such an occurrence introduces a shape asymmetry to the resultant spot. The collimation tube was removed in later design iterations in light of these findings.

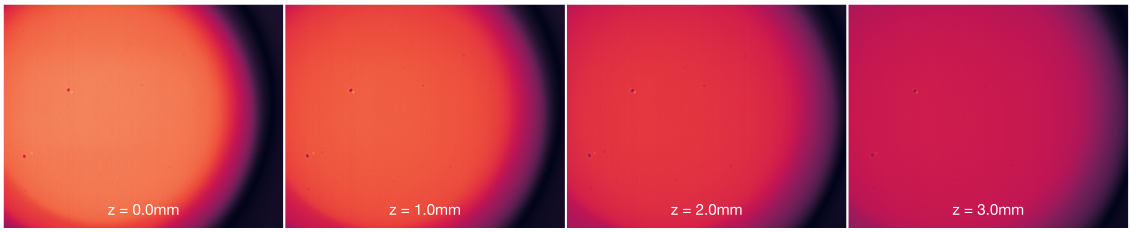

Fig. 13.4 Laboratory scan of the optical fibre model.

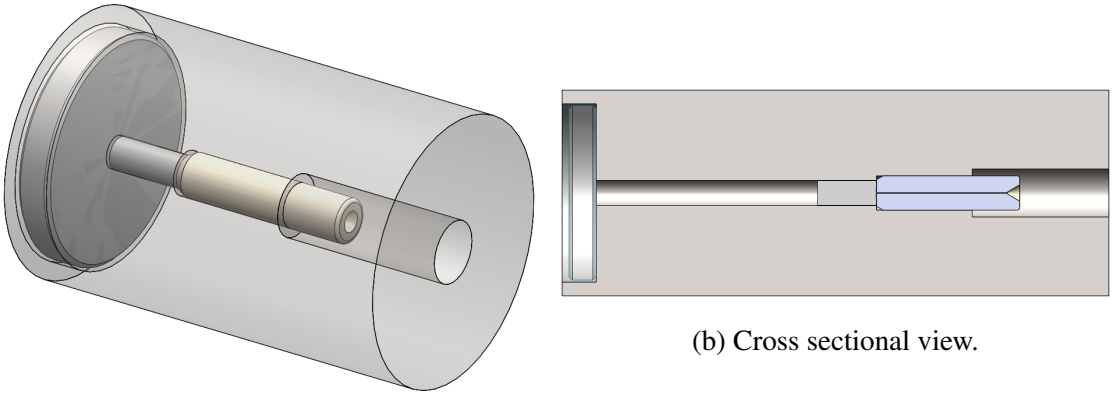

(a) Side view.

(b) Cross sectional view.

Fig. 13.5 Three-dimensional (Fig. 13.5a) and cross sectional (Fig. 13.5b) views of the collimator tested during the preliminary Super-Kamiokande deployment.

The light source was a system carrying a number of 435 nm wavelength LEDs. Coupled to each LED was a 1 mm diameter core fibre that split the light three ways: one to the optic, one to a 2" monitor PMT and one to an on-board photo-diode monitor which was yet to be installed at that time. The fibres coupled to the optics were step-index type with a 0.22 NA and a 200 μm diameter core[1]. A scan of the light output of this fibre model can be seen in Fig. 13.4. The pulse duration for each optic was set by hand according to the intensity measured within the detector—accordingly, the diffuser was assigned comparatively longer pulse lengths than the collimator and the control fibre owing to its wide beam spread. Each run accumulated $\sim 10^5$ inner detector events which were induced at a rate of 100 Hz.

[1] Thorlabs® FG200UEA.

13.2.2 Results

As seen in Fig. 13.6a, the pattern imaged by the collimator in initial runs has some degree of asymmetry, though the beam width is about the size expected and the position is sufficiently on target. The initial asymmetry is caused by poor alignment of the optics as previously described in Section 12.5. In a run taken on the following day, however, both the size, shape and position of the pattern had changed drastically—becoming about three times as large. The pattern in this later run is shown in Fig. 13.6b. It was concluded that the position of the fibre tip had likely drifted with respect to the lens due to insufficient coupling of the fibre to the assembly. The cable was either disturbed directly during experimental activities conducted at the top of the tank or by a spontaneous release of tension in the fibre. The misalignment of the optics then become so severe that multiple reflections of the beam projected by the lens upon the walls of the collimation tube were introduced, thus increasing the beam size and asymmetry. Another possibility, however, is that over time, as the air that had been trapped within the barrel after submerging gradually dissolved, water had made its way into the collimator and disturbed the alignment of the optics. Either way, this observation prompted revisions of several aspects of the collimator design to protect against misalignment and promote stability of the internal optics. First of all, the fibre was mechanically affixed to the holder by an FC connector to restrict its movement and improve reproducibility. Secondly, the movement of the lens was further restricted by manufacturing the assembly to a tighter tolerance. Thirdly, the entire assembly was made water-tight by introducing a glass window downstream of the optics. Lastly, the long and narrow barrel was removed as previously mentioned. The stability of the permanently installed optics demonstrates that these were crucial improvements to the design.

No abnormalities were observed in the diffuser pattern (Fig. 13.8); however, the intensity was still relatively low despite being injected with longer pulses. This informed future designs to minimise light loss due to coupling inefficiencies and attenuation in the fibre. As seen in Fig. 13.7, the control fibre fulfilled its purpose as a control source. The optics remained submerged in the water until refurbishment work began and displayed no obvious signs of fault upon their retrieval.

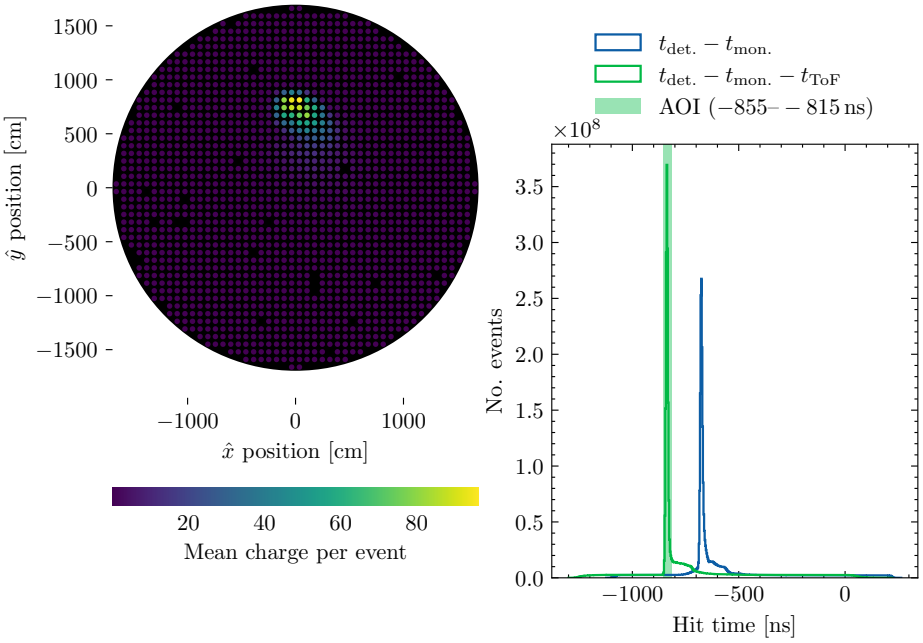

(a) In earlier run 77490.

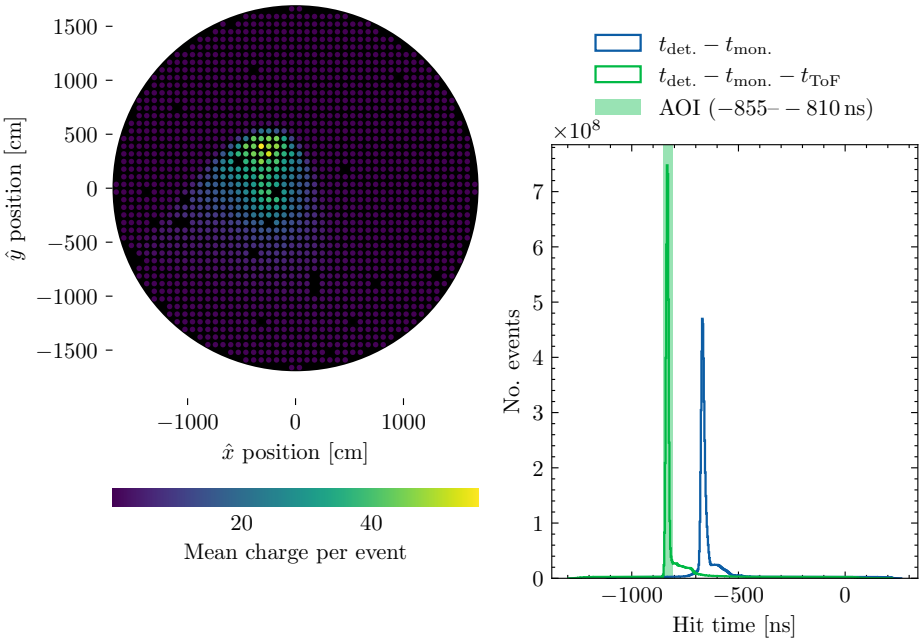

(b) In later run 77496.

Fig. 13.6 Charge hit maps on the bottom cap of the inner detector and corresponding timing distributions for two runs taken with the collimator during the preliminary deployment in January of 2018.

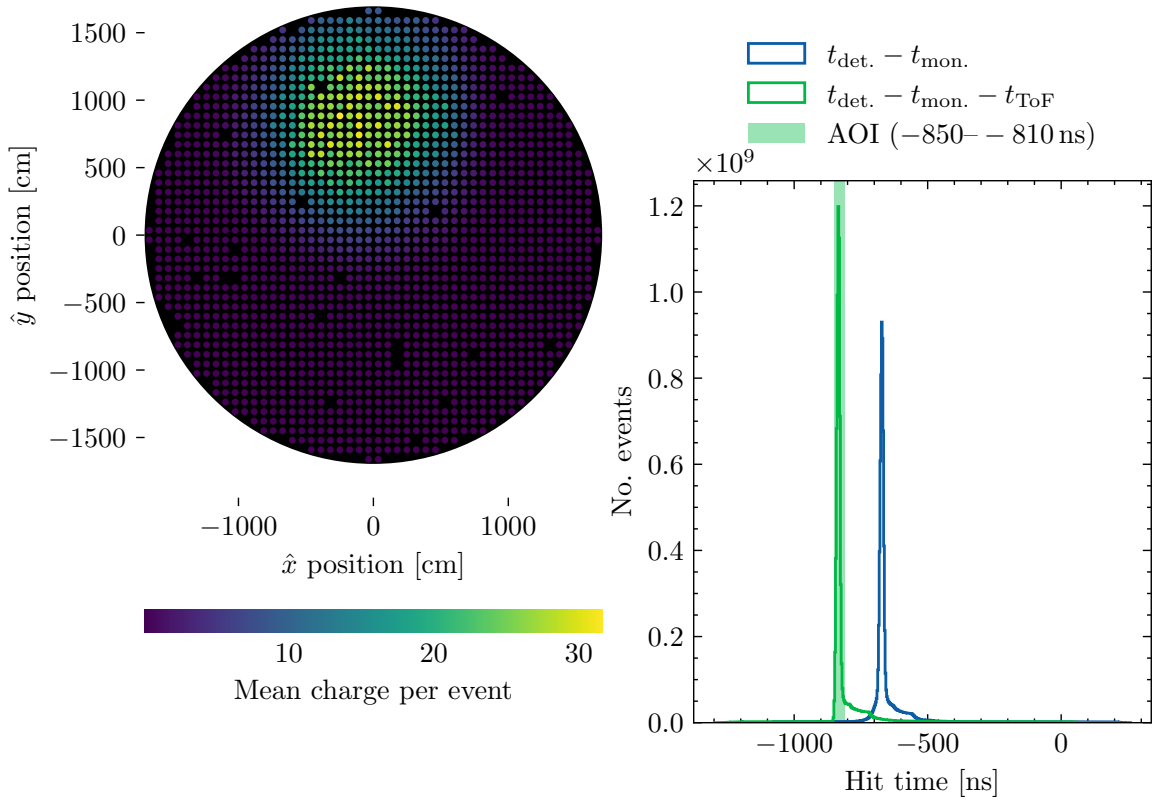

Fig. 13.7 Charge hit-map on the bottom cap of the inner detector and corresponding timing distribution for control fibre run 77486 taken during the preliminary deployment in January of 2018.

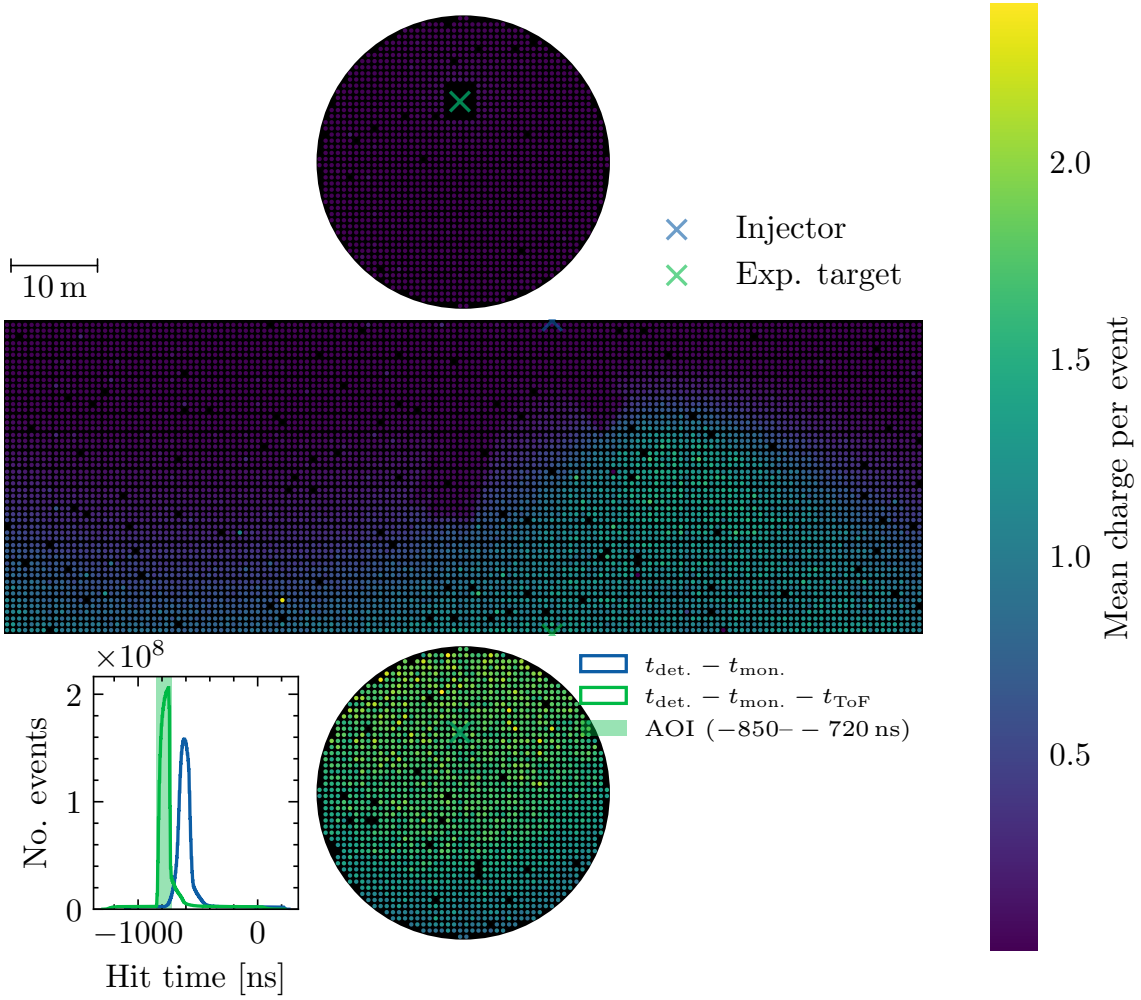

Fig. 13.8 Charge hit-map and corresponding timing distribution for diffuser run 77500 taken during the preliminary deployment in January of 2018.

13.3 Deployment of Barrel Injectors

An injector of the design described previously in Chapter 12 was installed at each of the injector positions B1, B2, B3, B4 and B5. The fibres that deliver light to the optics and the control fibres in the injectors were changed for the barrel injectors[2]. The specifications of the fibres used in the preliminary deployment and the barrel injectors differ by their numerical aperture—all other parameters are the same. A model of fibre with 0.50 NA[3]—compared to 0.22 of those used in the test deployment—was chosen to reduce bend-induced losses.

13.3.1 Results

There exist two datasets of optic tests for the barrel injectors: one conducted in February 2019, with 405 nm light injected from the laser diode pulser borrowed from the existing calibration system; and another conducted in the following July, using light from the dedicated 435 nm LED pulser.

The beam of the control fibre at B3 became very large at some point between February and July. All other optics show stability in shape between those two dates; the stability of the collimators is to be particularly celebrated given the happenings of the previous deployment. As seen in Table 13.2, the optics also display stability with respect to their target positions. The beam positions projected by the optics were measured by fitting a two-dimensional Gaussian function to the observed hit pattern on the region of the detector barrel directly opposite to each injector. The hit patterns induced by the control fibres at B1 and B5 in both datasets and at B3 in the July dataset were not measured due to their asymmetry. The deviations of these measured positions from their expected positions are given in Table 13.2. With the exception of injector B3, which is accurate to within 2°, the optics consistently target a position several photo-sensors clockwise of the expected target position. It well could be that the target positions were not marked correctly within the tank, however the reliability of the method of gauging the pointing direction of the injector is to be criticised as well. Likely a more reliable approach would be, rather, to instead inject the collimator and control fibre with light directly during the alignment procedure. The pointing direction of the diffuser will likely be unobservable to the naked eye owing to the large spread, and so should be inferred from the other optics. That said, these deviations are likely small enough to have a marginal impact on the physics measurements. Given these observations, the installation procedure itself can certainly be said to be successful.

[2] For the best account of the installation procedure, Ref. [150] is referred to.

[3] Thorlabs®FP200URT.

Injector	Optic	Run date	Target position (m)		
			$\hat{x}$	$\hat{y}$	$\hat{z}$
B1	Collimator	Feb. 2019	-15.96 ± 0.03	-5.64 ± 0.07	12.14 ± 0.05
		July 2019	-15.98 ± 0.02	-5.66 ± 0.08	12.13 ± 0.07
	Diffuser	Feb. 2019	-15.89 ± 0.30	-5.93 ± 0.81	10.08 ± 1.92
		July 2019	-16.10 ± 0.12	-5.32 ± 0.36	9.70 ± 0.83
B2	Collimator	Feb. 2019	-15.72 ± 0.03	-6.37 ± 0.07	7.51 ± 0.09
		July 2019	-15.91 ± 0.05	-5.88 ± 0.13	7.16 ± 0.20
	Control fibre	Feb. 2019	-15.96 ± 0.03	-5.72 ± 0.08	6.73 ± 0.09
		July 2019	-15.99 ± 0.03	-5.65 ± 0.07	6.98 ± 0.08
	Diffuser	Feb. 2019	-15.82 ± 0.30	-6.11 ± 0.77	6.11 ± 1.64
		July 2019	-15.94 ± 0.11	-5.77 ± 0.30	6.02 ± 0.61
B3	Collimator	Feb. 2019	-15.37 ± 0.03	-7.16 ± 0.07	-1.85 ± 0.08
		July 2019	-15.29 ± 0.04	-7.34 ± 0.08	-1.73 ± 0.08
	Control fibre	Feb. 2019	-15.45 ± 0.03	-6.99 ± 0.08	-1.71 ± 0.09
	Diffuser	Feb. 2019	-15.32 ± 0.44	-7.27 ± 0.92	-1.76 ± 1.95
		July 2019	-15.48 ± 0.16	-6.91 ± 0.36	-1.37 ± 0.69
B4	Collimator	Feb. 2019	-14.95 ± 0.08	-7.99 ± 0.15	-7.61 ± 0.17
		July 2019	-15.03 ± 0.03	-7.85 ± 0.06	-7.45 ± 0.04
	Control fibre	Feb. 2019	-15.05 ± 0.04	-7.81 ± 0.08	-7.08 ± 0.09
		July 2019	-15.06 ± 0.04	-7.79 ± 0.08	-7.10 ± 0.09
	Diffuser	Feb. 2019	-15.25 ± 0.42	-7.42 ± 0.87	-7.26 ± 2.06
		July 2019	-15.69 ± 0.20	-6.42 ± 0.48	-7.31 ± 1.15
B5	Collimator	Feb. 2019	-15.13 ± 0.05	-7.65 ± 0.10	-14.07 ± 0.12
		July 2019	-15.12 ± 0.08	-7.67 ± 0.16	-13.79 ± 0.25
	Diffuser	Feb. 2019	-15.17 ± 0.49	-7.58 ± 0.99	-10.93 ± 2.41
		July 2019	-15.47 ± 0.24	-6.93 ± 0.55	-11.03 ± 1.35

Table 13.2 Measured target positions of the beams projected by each optic of each barrel position light injector in Super-Kamiokande.

Injector	Deviation from Target (°)		
	Collimator	Control fibre	Diffuser
B1	5.2 ± 0.2	—	7.4 ± 3.8
B2	4.9 ± 0.3	5.6 ± 0.2	5.2 ± 2.5
B3	1.3 ± 0.2	1.7 ± 0.1	1.6 ± 3.5
B4	2.3 ± 0.4	1.9 ± 0.3	3.5 ± 4.1
B5	3.1 ± 0.5	—	5.0 ± 5.0

Table 13.3 Angular deviations of measured target positions from their expected target positions of the beams projected by each optic of each barrel injector in Super-Kamiokande.

The expected FWHM beam diameter projected by the collimators is close to two photosensor spacings; a qualitative evaluation of the hit-maps shown in Figs. 13.9 to 13.13 show that the sizes are about as expected. Given these sizes, the spatial resolution of the detector impedes a quantitative evaluation of their shapes. However, the hit patterns induced by collimators B1 through to B4 display a satisfactory degree of uniformity; at B5, however, the pattern is unusual. By assessing the noise level with respect to the signal peak in the timing distributions, the transmissibility of the optic can be judged. Again, collimators at B1 through to B4—albeit with B3 somewhat less bright than the others—induce a sufficiently strong detector signal. As seen in Figs. 12.19 and 12.21, laboratory measurements of B5 indicate that the opening-angle is lower than the rest of the collimators bar one, and the tails of the intensity distribution are the shortest. There is nothing in the azimuthal distributions (Fig. 12.22), however, that are indicative of the irregular shape observed. It can reasonably be assumed that the low intensity and irregular hit pattern of this unit are interrelated and indicative of a fault introduced sometime between shipment and the first in-situ measurements. Measurements of the optical loss of each collimator were beyond the capability of the test stand at the time; certainly, this should be implemented for the production of the HK collimators. Noticeable in the timing distributions in the collimators of the July dataset (Figs. 13.9 to 13.13) is a double peak structure separated by $\sim$80 ns—corresponding to an additional path length of $\sim$18 m. The first peak corresponds to direct light, the second peak apparently corresponds to a low-charge ring of a few degrees greater opening-angle than the direct spot. No such occurrence is observed of the collimators in the February 2019 dataset—therefore it can be reasonably assumed that this artefact is conditional on the properties of the injected light. Likely this structure is not a property of the pulses injected during the July 2019 dataset as no such structure is observed of those of the control fibres in the same dataset. The timing responses of the collimators were not measured before installation—though certainly they should be characterised for the HK collimators over a variety of wavelengths.

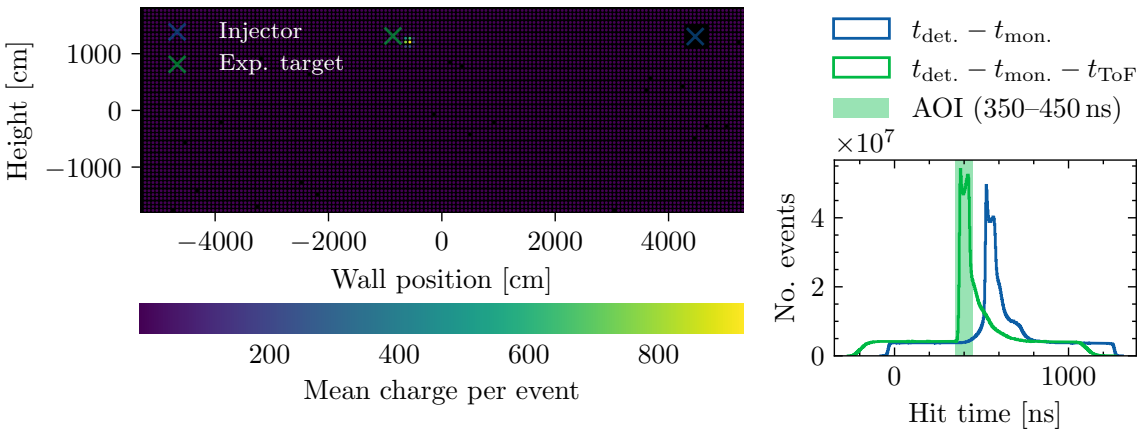

Fig. 13.9 Charge hit-map on the wall of the inner detector and corresponding timing distribution for B1 collimator run 81390 taken during the preliminary deployment in July of 2019.

The control fibres at B2 and B4 are as expected; the remainder display some unexpected properties. Namely, the hit pattern induced by the B1 control fibre is asymmetric and larger than expected, and those of B3 and B5 are more diffuse than their adjacent diffusers but symmetric and uniform. Though the properties of the fibres of the same specification as the installed control fibres are well-known, the installed ones were not measured in the laboratory before installation. Therefore, a pre-existing fault in the fibres cannot be ruled out. Judging from the hit pattern and given prior experience of broken fibres in the laboratory, it is reasonable to assume that the B1 fibre may have been damaged before being submerged. That the B5 control fibre came to resemble the B3 after six months submerged Fig. 13.14, however, raises some concerns about their longevity. A study of the optical effects of long-term water exposure to these fibres was not possible given the time scale of the development process. In retrospect, encasing the fibre ferrule within a water-tight enclosure similar to, but much smaller than, that of the diffuser would have been a sensible measure taken against such effects.

The hit patterns induced by all diffusers are as expected. A typical hit-map is given for Fig. 13.15. The transmission of the B5 diffuser is perhaps lower than the rest. The collimator assembly intrudes on the beams of the B4 and B5, ultimately causing a shadow on the bottom cap of the detector. This is the result of an oversight made during the production of the collimators. At the final stage of the assembly process, a small amount of epoxy resin is applied on top of the grub screws that fix the lens-mount within the sleeve to prevent leakages. This small amount of epoxy was however large enough for all collimators, with the exception of B2, to prevent the collimator cap from sitting relatively flush to the optical plate. Inserting the collimator into its slot from the rear was not feasible either as the bolts attaching the cap

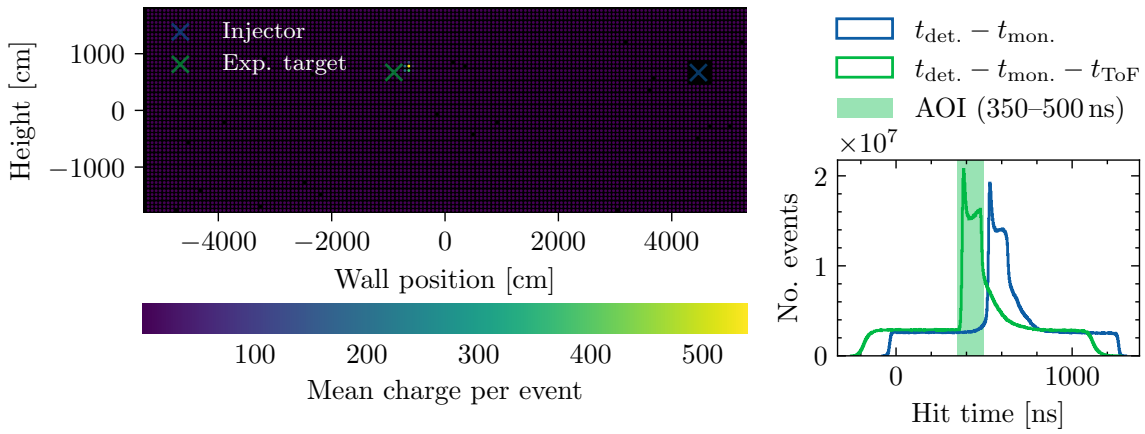

Fig. 13.10 Charge hit-map on the wall of the inner detector and corresponding timing distribution for B2 collimator run 81391 taken during the preliminary deployment in July of 2019.

to the sleeve overhang the cap. Hence, for these collimators, the sleeve protrudes ∼25 mm from the optical plate—large enough to reside within the field of view of the diffuser.

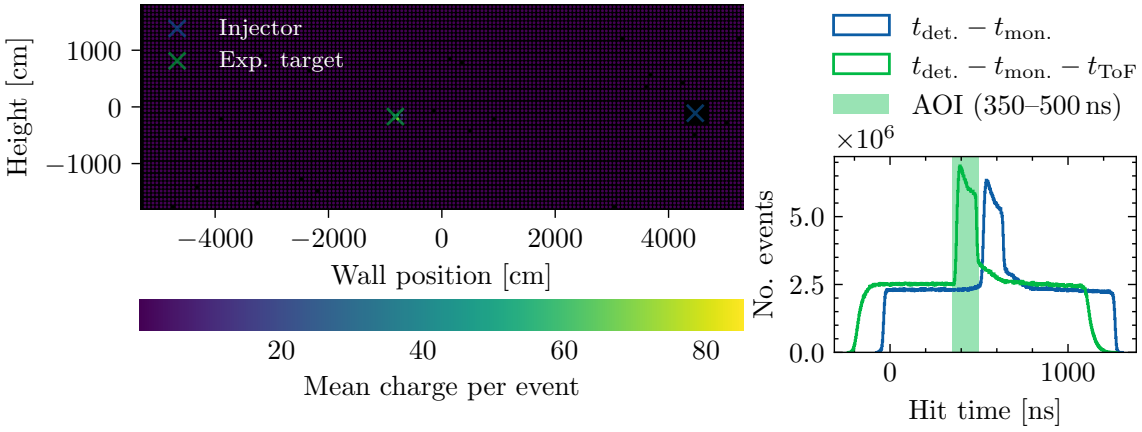

Fig. 13.11 Charge hit-map on the wall of the inner detector and corresponding timing distribution for B3 collimator run 81392 taken during the preliminary deployment in July of 2019.

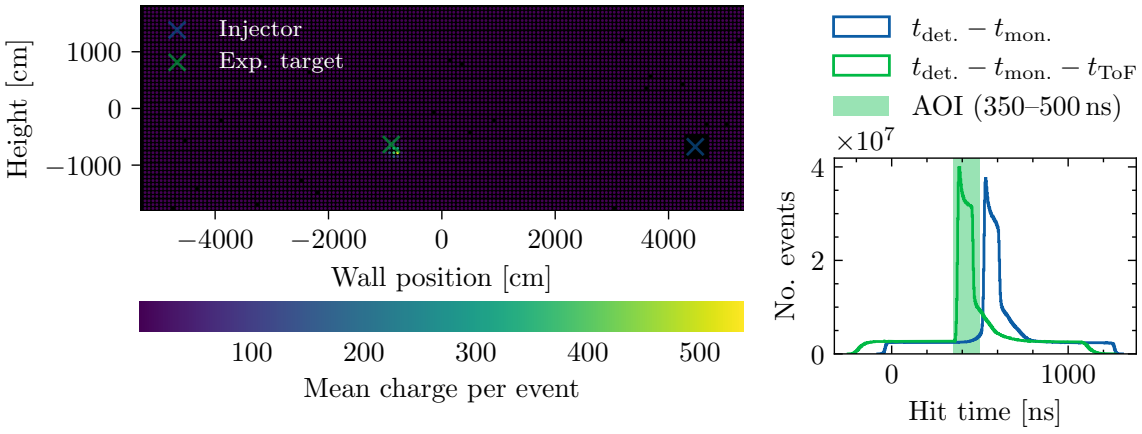

Fig. 13.12 Charge hit-map on the wall of the inner detector and corresponding timing distribution for B4 collimator run 81393 taken during the preliminary deployment in July of 2019.

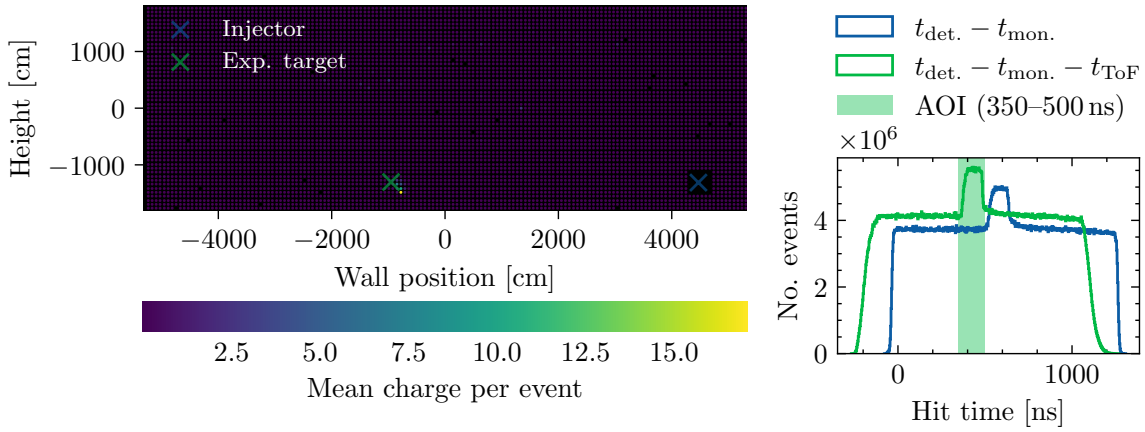

Fig. 13.13 Charge hit-map on the wall of the inner detector and corresponding timing distribution for B5 collimator run 81394 taken during the preliminary deployment in July of 2019.

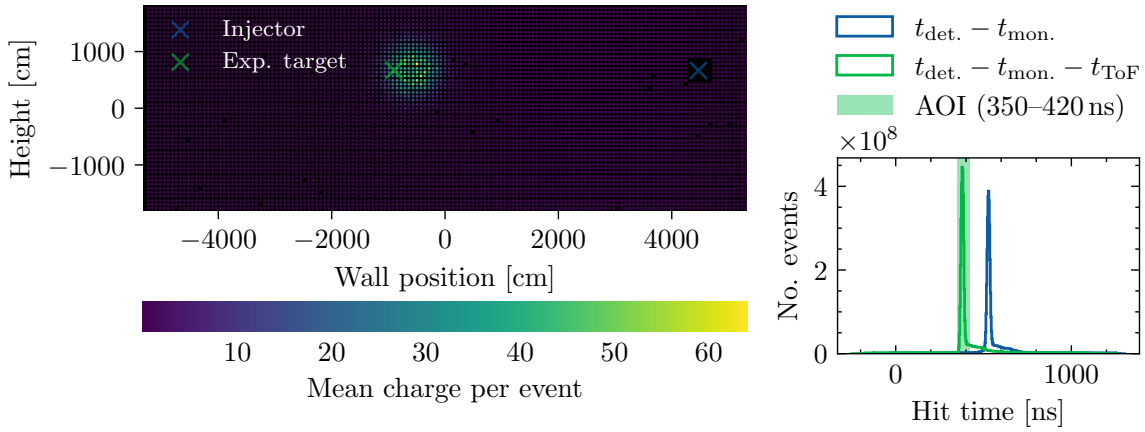

Fig. 13.14 B2 control fibre, Run 81403.

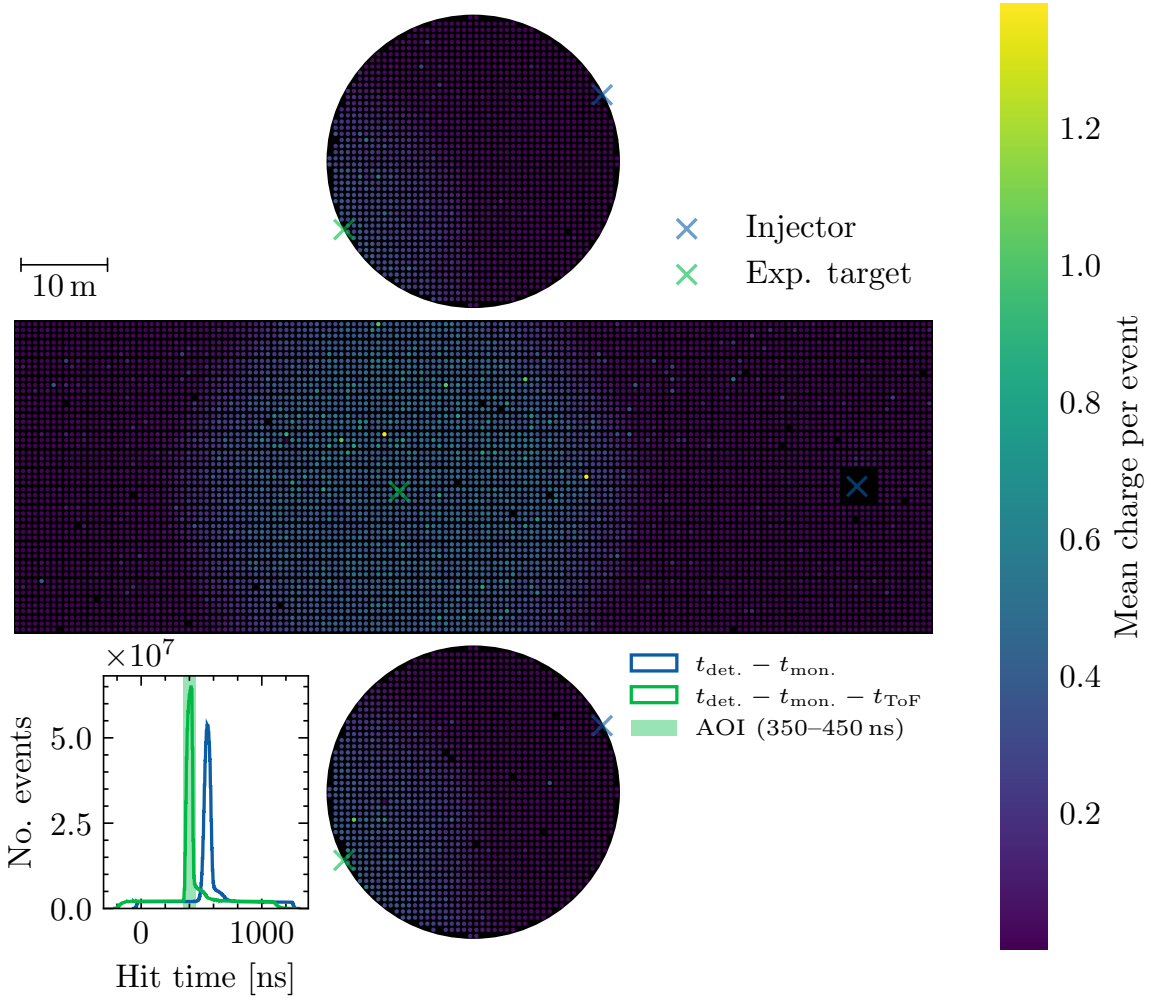

Fig. 13.15 B3 diffuser, Run 81404.

CHAPTER 14

CONCLUSION

A light injection system has been developed to calibrate the photo-sensor responses and water properties of large-scale water Cherenkov neutrino detectors. The system consists of a source that delivers light to an array of optics via optical fibres installed within the tank. When supplied with light, the optics project characteristic light sources across the volume of the tank. The detector data induced by these well-understood sources allow several calibrations to be performed. Two optical devices are in development—a diffuser, to project a uniform source with a wide field of view for the purpose of calibrating the photo-sensor responses; and a collimator, to project a high-intensity, low-angle beam for the purpose of monitoring the levels of optical scattering and absorption effected by the tank's water. The optical properties of these devices are tested extensively within their dedicated test stand systems as well as their durability when exposed to high pressures. In the summer of 2018, seven diffusers and six collimators were produced, of which five of each were permanently installed within Super-Kamiokande. The detector data later collected demonstrated that the optics were stable over the first six months of being submerged. This small light injection system has now been integrated into Super-Kamiokande's automated calibration routine, with which the optics are being continuously monitored in order to assess their longevity.

14.1 Future Prospects

Taking on board the lessons learned from this production process, the development of these devices continues in anticipation of the production of a comprehensive light injection system for Hyper-Kamiokande and possibly other future water Cherenkov detectors. Following a design review, two particular aspects of these optics have been identified to be improved for

future systems—that is, their scalability and long-term durability. Some current studies of the collimator and diffuser designs that address these issues will now be presented.

14.1.1 The diffuser

Improvements are being made to the diffuser material and its enclosure. A new candidate diffuser material is polytetrafluoroethylene (PTFE). Along with various practical benefits of using PTFE over PMMA, the transmittance is improved over the planned spectrum of the light source—below 350 nm, the transmittance of PMMA begins to fall; PTFE, however, retains that transmittance. Unlike PMMA, PTFE is also recognised for its chemical and water resistance. It is unlikely that the enclosure will be abandoned in light of these properties, however, due to its role in flattening the diffuser's angular intensity profile as already discussed. Accordingly, another area of diffuser research is the optical effects of surface treatments applied to the enclosure interior. The angular intensity profiles for different surface treatments are shown in Fig. 14.2. Practical improvements to the enclosure design are under consideration also.

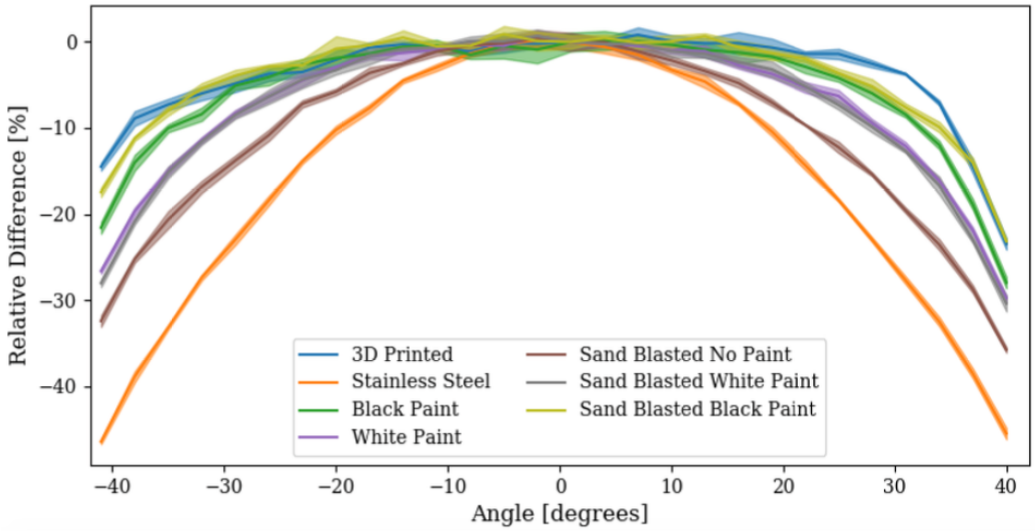

Fig. 14.1 Comparison of light intensity profile shapes for different surface treatments of the diffuser enclosure interior (reproduced from Ref. [152]).

Fig. 14.2 Observed light intensity as a function of polar angle from the diffuser for various surface treatments of the enclosure interior. The measured intensities are presented as a fraction of the peak intensity. Reproduced from Ref. [152].

14.1.2 The collimator

Scalability is an issue particularly relevant to the collimator: the commissioning of a single collimator calls for precision machining and a significant amount of tuning and testing by an expert due to the sensitivity of the optics. The production of perhaps a hundred of such units for HK would be laborious to the point of infeasibility. Another issue is that of inflexibility: unlike the diffuser, which would be effective inside both a very large detector like HK and a smaller one like the IWCD, the size of the collimator's narrow beam must be tuned in proportion to the detector size. Given that the opening angle of the optics is capped by the size of the fibre core, a larger opening angle would be impossible without changing the specifications of either the lens or the fibre—and the choice of fibre will likely be motivated by other requirements of the light injection system. The use of a lens with a lower pitch than ¼—lenses with 0.23 pitch are readily available—could be satisfactory, though not guaranteed. An ideal collimator would be fairly insensitive to the fibre specifications and its beam size would be adjustable rather by mechanical means.

Both of these are issues inherent to the lens; therefore, an alternative source of focussing is being sought. One such candidate is the replacement of the GRIN lens with a conventional lens. Preliminary simulation and laboratory studies have been conducted with a 25.4 mm diameter, 30 mm focal length achromatic doublet lens[1]. In particular, the properties of this lens model has been studied when coupled with a graded-index fibre of 62.5 μm diameter core, 125 μm diameter core and 0.275 NA[2]. The beam imaged by this fibre model and the corresponding radial intensity distribution may be seen for various beam distances in Figs. 14.3 and 14.4. The horizontal beam profile of the lens–fibre coupled system when separated by $\sim$7 mm may be found in Fig. 14.5—the shape is sufficiently Gaussian, thus following the shape of the beam projected by the coupled fibre. It is therefore necessary for an aperture to be placed upstream of this system to improve the uniformity of the radial intensity distribution. As well, simulations and laboratory tests have been conducted to probe the sensitivity of this system to the misalignment of the lens with respect to the fibre. As the lens is much larger, the optical system is far less sensitive to $\sim$1 mm misalignments transverse to the optical axis.

The focal point of the lens is $\sim$22 mm from the back face; the lens–fibre separation necessary to achieve a beam with $\sim$3.5° half opening angle is $\sim$17 mm. As the light from the fibre must propagate for some distance before reaching the lens, the beam diameter over laboratory distances ($\sim$5 cm) is larger than that of the GRIN lens collimator and thus larger than the sensor of the test stand camera. Hence, capturing the full extent of the beam cross

[1] Thorlabs® AC254-030-A-ML

[2] Thorlabs® GIF625

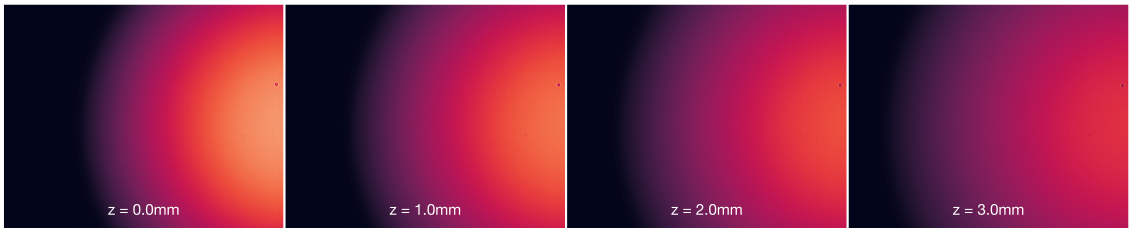

Fig. 14.3 Laboratory scans of the graded-index optical fibre.

section is not possible without an upgrade to both the test stand hardware and software. A preliminary enclosure design to house this lens is shown in Fig. 14.6. The lens and the fibre are to be mounted within a threaded tube in which their positions are mechanically adjustable. Ray tracing simulations have been conducted to inform the dimensions of this system—specifically, the separation between the lens and the aperture and the lens and fibre.

Various improvements are also being made to the test stand. The greatest flaw with the setup described is that the longitudinal scan range is limited by the beam diameter of the test collimator, as the data-acquisition software was unable to scan along multiple axes in the same run. It is optimal to sample the collimator beam over the largest range possible within the lab for extrapolations to HK distances to be as accurate as possible. Since then, the data-acquisition software was rewritten from the ground-up in Python, allowing for three-dimensional scans in the $\hat{x}$, $\hat{y}$ and $\hat{z}$ directions. For three-dimensional scans, images taken in the $\hat{x}\hat{y}$ plane are stitched together at the analysis level and processed as previously described. As well, the collimators have only been measured for one wavelength; ideally, their performances should be known for the range of wavelengths that they will be injected with when installed in HK. Accordingly, a three-colour (395 nm, 505 nm and 625 nm) LED switcher is being integrated into the test stand setup.

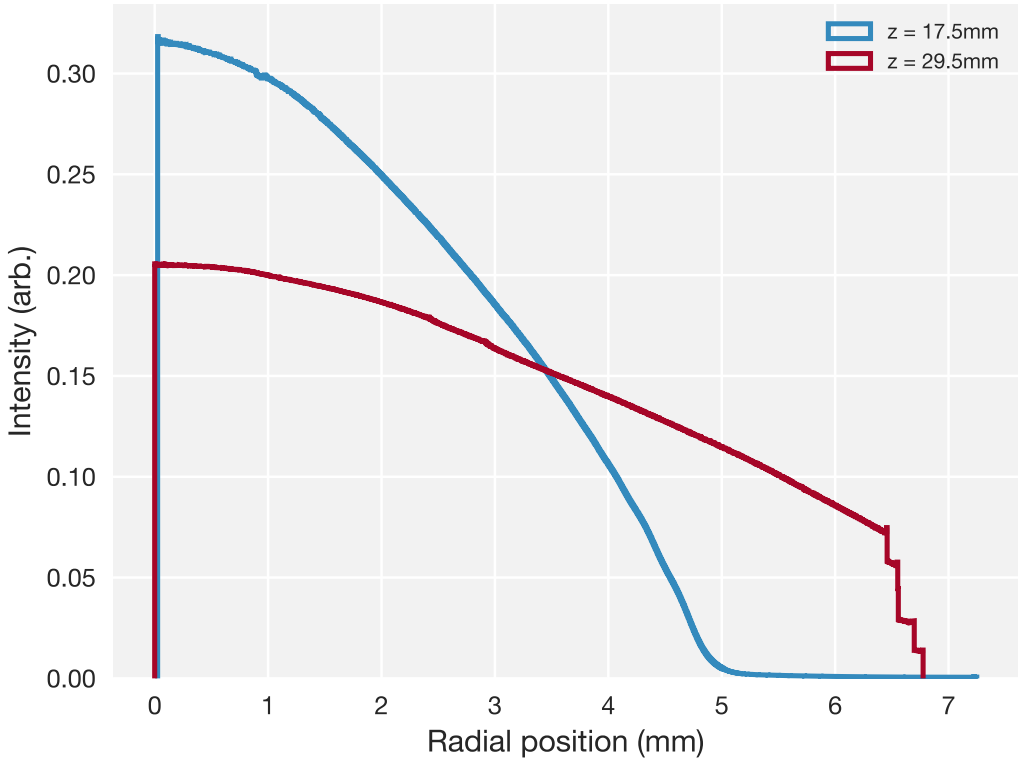

Fig. 14.4 In-air radial intensity distribution of the graded-index optical fibre as measured at two longitudinal distances in the laboratory.

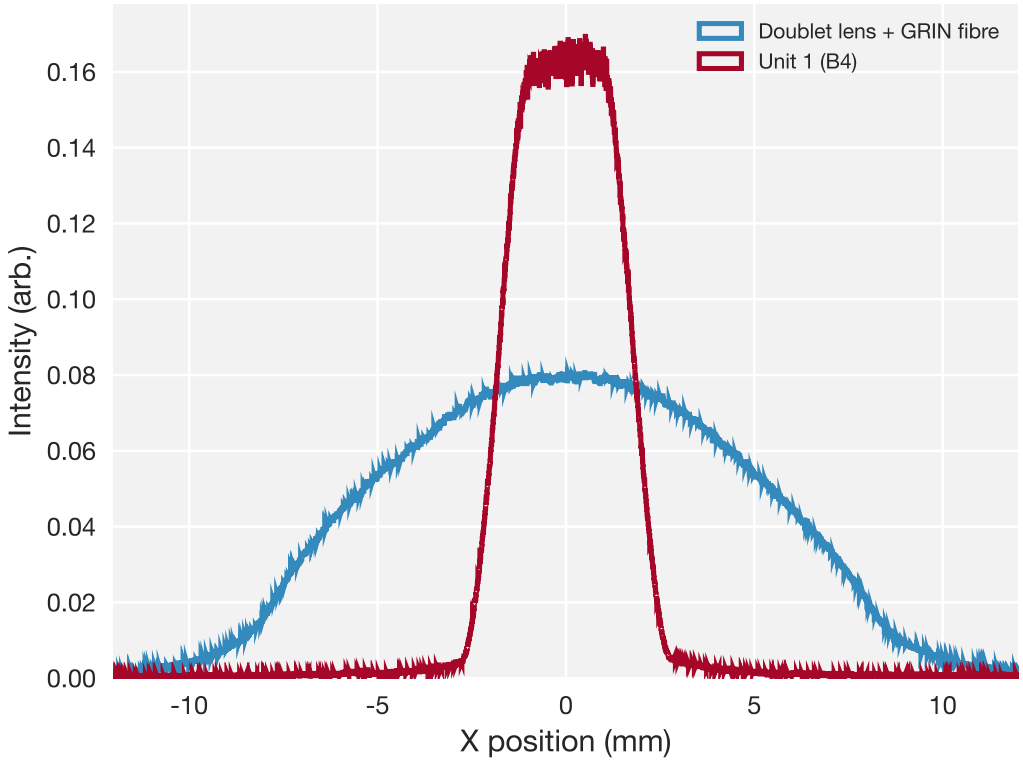

Fig. 14.5 Intensity along horizontal of centre of beam as measured of the doublet lens and collimator unit 1 at an approximately equivalent beam distance in the laboratory.

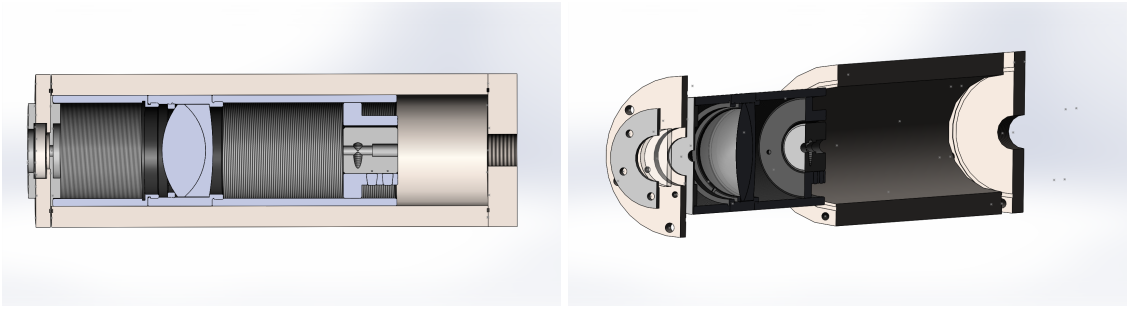

(a) Cross sectional view.

(b) Exploded view.

Fig. 14.6 Cross sectional (Fig. 14.6a) and exploded views (Fig. 14.6b) of a preliminary design of a collimator assembly to house a conventional lens.

PART IV

CONCLUSION

CHAPTER 15

CONCLUSION

This thesis presented a measurement of the neutrino–proton NC0π cross section at T2K's near detector and demonstrated its improvement with a machine-learning driven signal selection. A thorough set of conclusions have been drawn in Chapter 9. The flux-averaged value was measured as $(3.31 \pm 0.36) \times 10^{-40}\,\mathrm{cm}^2\,\mathrm{nucleon}^{-1}$, reducing the fractional error of the measurement by 28% compared to that of a nominal cut-based selection. The differential cross sections with respect to the reconstructed kinetic energy and polar angle of the primary proton candidate have also been presented.

This thesis also presented a prototype light injection system for the calibration of the future water Cherenkov detector, Hyper-Kamiokande. The key findings of these studies and proposals for future work have been offered in Chapter 14. The optical devices of said system were developed at the University of Warwick's HK laboratory and were tested extensively in their test stands and once in-situ in the Super-Kamiokande detector. During the summer of 2018, Super-Kamiokande was drained for maintenance in preparation for Gadolinium-loading. These optics were successfully installed along the vertical of the detector barrel during this period and were connected to its dedicated light source soon after. The signal from each optic was analysed from detector data collected in February and July of 2019 in which stability was displayed. Soon after that, the light injection system was integrated into the detector's automated calibration routine. The new system now allows SK to monitor the levels of scattering and absorption present in the water at five distinct depths. Since the installation, the designs of the optics have been subject to review. Following this, efforts are underway to improve the scalability of these optics as Hyper-Kamiokande's construction phase draws nearer.

BIBLIOGRAPHY

[1] A. D. Sakharov, "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe", Journal of Experimental and Theoretical Physics Letters **5**, 24–27 (1967).

[2] K. Abe et al., "Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations", Nature **580**, 339–344 (2020).

[3] W. Pauli, "Dear radioactive ladies and gentlemen", Physics Today **31N9**, 27 (1978).

[4] F. Reines and C. L. Cowan, "Detection of the free neutrino", Physical Review **92**, 830–831 (1953).

[5] C. L. Cowan et al., "Detection of the free neutrino: A confirmation", Science **124**, 103–104 (1956).

[6] R. Davis, "Attempt to Detect the Antineutrinos from a Nuclear Reactor by the $Cl^{37}(\overline{\nu}, e^{-})A^{37}$ Reaction", Physical Review **97**, 766–769 (1955).

[7] R. Davis and D. S. Harmer, "Attempt to observe the reaction $Cl^{37}(\bar{\nu}_e\, e^{-})Ar^{37}$ induced by reactor antineutrinos", in Bulletin of the American Physical Society, Vol. 4 (1959), p. 217.

[8] G. Danby et al., "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos", Physical Review Letters **9**, 36–44 (1962).

[9] K. Kodama et al., "Observation of tau neutrino interactions", Physics Letters B **504**, 218–224 (2001).

[10] M. L. Perl et al., "Evidence for anomalous lepton production in e^{+}–e^{-} annihilation", Physical Review Letters **35**, 1489–1492 (1975).

[11] M. Acciarri et al., "Determination of the number of light neutrino species from single photon production at LEP", Physics Letters B **431**, 199–208 (1998).

[12] R. Davis, "Solar Neutrinos. II. Experimental", Physical Review Letters **12**, 303–305 (1964).

[13] R. Davis, D. S. Harmer, and K. C. Hoffman, "Search for neutrinos from the sun", Physical Review Letters **20**, 1205–1209 (1968).

[14] W. Hampel et al., "GALLEX solar neutrino observations: Results for GALLEX IV", Physics Letters B **447**, 127–133 (1999).

[15] J. N. Abdurashitov et al., "Measurement of the solar neutrino capture rate with gallium metal", Physical Review C **60**, 055801 (1999).

[16] K. S. Hirata et al., "Observation of ^{8}B solar neutrinos in the Kamiokande-II detector", Physical Review Letters **63**, 1696–1699 (1989).

[17] K. Zuber, *Neutrino Physics*, 2nd ed. (Taylor and Francis, 2011), pp. 1–465.

[18] C. V. Achar et al., "Detection of muons produced by cosmic ray neutrinos deep underground", Physics Letters **18**, 196–199 (1965).

[19] F. Reines et al., "Evidence for high-energy cosmic-ray neutrino interactions", Physical Review Letters **15**, 429–433 (1965).

[20] M. F. Crouch et al., "Cosmic-ray muon fluxes deep underground: Intensity vs depth, and the neutrino-induced component", Physical Review D **18**, 2239–2252 (1978).

[21] K. Hirata et al., "Experimental study of the atmospheric neutrino flux", Physics Letters B **205**, 416–420 (1988).

[22] R. Becker-Szendy et al., "Search for muon neutrino oscillations with the Irvine-Michigan-Brookhaven detector", Physical Review Letters **69**, 1010–1013 (1992), http://link.aps.org/doi/10.1103/PhysRevLett.69.1010.

[23] W. Allison et al., "Measurement of the atmospheric neutrino flavour composition in Soudan 2", Physics Letters B **391**, 491–500 (1997).

[24] W. W. Allison et al., "The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of Soudan 2", Physics Letters B **449**, 137–144 (1999).

[25] M. Aglietta et al., "Experimental study of atmospheric neutrino flux in the NUSEX experiment", EPL **8**, 611–614 (1989).

[26] K. Daum et al., "Determination of the atmospheric neutrino spectra with the Fréjus detector", Zeitschrift für Physik C Particles and Fields **66**, 417–428 (1995).

[27] C. S. Wu et al., "Experimental test of parity conservation in beta decay", Physical Review **105**, 1413–1415 (1957).

[28] T. D. Lee and C. N. Yang, "Question of parity conservation in weak interactions", Physical Review **104**, 254–258 (1956).

[29] A. Salam and J. C. Ward, "Electromagnetic and weak interactions", Physics Letters **13**, 168–171 (1964).

[30] F. J. Hasert et al., "Search for elastic muon-neutrino electron scattering", Physics Letters B **46**, 121–124 (1973).

[31] F. J. Hasert et al., "Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment", Physics Letters B **46**, 138–140 (1973).

[32] C. Albajar et al., "Studies of intermediate vector boson production and decay in UA1 at the CERN proton-antiproton collider", Zeitschrift für Physik C Particles and Fields **44**, 15–61 (1989).

[33] P. A. Zyla et al., "Review of particle physics", Progress of Theoretical and Experimental Physics **2020**, 083C01 (2020).

[34] B. Pontecorvo, "Mesonium and anti-mesonium", Soviet Physics JETP **6**, 429 (1957).

[35] Z. Maki, M. Nakagawa, and S. Sakata, "Remarks on the unified model of elementary particles", Progress of Theoretical and Experimental Physics **28**, 870–880 (1962).

[36] B. Pontecorvo, "Neutrino experiments and the problem of conservation of leptonic charge", Soviet Physics JETP **26**, 984–988 (1967).

[37] Y. Fukuda et al., "Atmospheric ν_μ/ν_e ratio in the multi-GeV energy range", Physics Letters B **335**, 237–245 (1994).

[38] Y. Fukuda et al., "Evidence for Oscillation of Atmospheric Neutrinos", Physical Review Letters **81**, 1562–1567 (1998).

[39] 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", Physical Review Letters **87**, 071301 (2001).

[40] A. Y. Smirnov, "Where do we stand with solar neutrino oscillations?", Physical Review D **58**, 096016 (1998).

[41] A. Gando et al., "Reactor on-off antineutrino measurement with KamLAND", Physical Review D **88**, 033001 (2013).

[42] N. Agafonova et al., "Final results on neutrino oscillation parameters from the OPERA experiment in the CNGS beam", Physical Review D **100**, 051301 (2019).

[43] K. Abe et al., "Supernova Model Discrimination with Hyper-Kamiokande", The Astrophysical Journal **916**, 15 (2021).

[44] K. Abe et al., "Observation of electron neutrino appearance in a muon neutrino beam", Physical Review Letters **112**, 061802 (2014).

[45] P. Adamson et al., "Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset", Physical Review Letters **125**, 131802 (2020).

[46] M. A. Acero et al., "First measurement of neutrino oscillation parameters using neutrinos and antineutrinos by NOvA", Physical Review Letters **123**, 151803 (2019).

[47] K. Abe et al., "Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV", Physical Review D **97**, 072001 (2018).

[48] M. G. Aartsen et al., "Measurement of Atmospheric Neutrino Oscillations at 6-56 GeV with IceCube DeepCore", Physical Review Letters **120**, 071801 (2018).

[49] P. F. de Salas et al., "2020 global reassessment of the neutrino oscillation picture", Journal of High Energy Physics **2021**, 71 (2021).

[50] K. Abe et al., "Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam", Physical Review Letters **107**, 041801 (2011).

[51] M. Apollonio et al., "Search for neutrino oscillations on a long base-line at the Chooz nuclear power station", European Physical Journal C **27**, 331–374 (2003).

[52] P. Adamson et al., "Electron neutrino and antineutrino appearance in the full MINOS data sample", Physical Review Letters **110**, 171801 (2013).

[53] D. Adey et al., "Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay", Physical Review Letters **121**, 241805 (2018).

[54] C. D. Shin et al., "Observation of reactor antineutrino disappearance using delayed neutron capture on hydrogen at RENO", Journal of High Energy Physics **2020**, 29 (2020).

[55] H. de Kerret et al., "Double Chooz θ_{13} measurement via total neutron capture detection", Nature Physics **16**, 558–564 (2020).

[56] K. N. Abazajian et al., "Light Sterile Neutrinos: A White Paper", 10.48550/ARXIV. 1204.5379 (2012).

[57] C. Athanassopoulos et al., "Candidate Events in a Search for $\nu_\mu \to \nu_e$ Oscillations", Physical Review Letters **75**, 2650–2653 (1995).

[58] C. Athanassopoulos et al., "Results on $\nu_\mu \to \nu_e$ Neutrino Oscillations from the LSND Experiment", Physical Review Letters **81**, 1774–1777 (1998).

[59] A. A. Aguilar-Arevalo et al., "Updated MiniBooNE neutrino oscillation results with increased data and new background studies", Physical Review D **103**, 052002 (2021).

[60] C. Giunti and T. Lasserre, "eV-Scale Sterile Neutrinos", Annual Review of Nuclear and Particle Science **69**, 163–190 (2019).

[61] F. Kaether et al., "Reanalysis of the GALLEX solar neutrino flux and source experiments", Physics Letters B **685**, 47–54 (2010).

[62] J. N. Abdurashitov et al., "Measurement of the response of a Ga solar neutrino experiment to neutrinos from a ^{37}Ar source", Physical Review C **73**, 045805 (2006).

[63] J. Kostensalo et al., "The gallium anomaly revisited", Physics Letters B **795**, 542–547 (2019).

[64] R. Acciarri et al., "A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam", 10.48550/ARXIV.1503.01520 (2015).

[65] A. Gando et al., "Partial radiogenic heat model for Earth revealed by geoneutrino measurements", Nature Geoscience **4**, 647–651 (2011).

[66] M. Agostini et al., "Spectroscopy of geoneutrinos from 2056 days of Borexino data", Physical Review D - Particles, Fields, Gravitation and Cosmology **92**, 031101 (2015).

[67] E. Baracchini et al., "PTOLEMY: A Proposal for Thermal Relic Detection of Massive Neutrinos and Directional Detection of MeV Dark Matter", 10.48550/ARXIV.1808.01892 (2018), https://arxiv.org/abs/1808.01892.

[68] S. Abe, "First Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen", (2022).

[69] V. Albanese, "The SNO+ experiment", JINST **16**, P08059 (2021).

[70] R. Arnold et al., "Probing new physics models of neutrinoless double beta decay with SuperNEMO", The European Physical Journal C **70**, 927–943 (2010), https://doi.org/10.1140/epjc/s10052-010-1481-5.

[71] M. Aker et al., "Direct neutrino-mass measurement with sub-electronvolt sensitivity", Nature Physics **18**, 160–166 (2022), https://doi.org/10.1038/s41567-021-01463-1.

[72] B. Abi et al., "Volume I. Introduction to DUNE", Journal of Instrumentation **15**, T08008 (2020).

[73] J. A. Formaggio and G. P. Zeller, "From eV to EeV: Neutrino cross sections across energy scales", Reviews of Modern Physics **84**, 1307–1341 (2012).

[74] M. S. Athar and S. K. Singh, *The Physics of Neutrino Interactions*, 1st ed. (Cambridge University Press, 2020), p. 930.

[75] C. H. Llewellyn Smith, "Neutrino reactions at accelerator energies", Physics Reports **3**, 261–379 (1972).

[76] A. Bodek et al., "Vector and axial nucleon form factors: A duality constrained parameterization", European Physical Journal C **53**, 349–354 (2008).

[77] A. A. Aguilar-Arevalo et al., "First measurement of the muon neutrino charged current quasielastic double differential cross section", Physical Review D **81**, 092005 (2010).

[78] R. Gran et al., "Measurement of the quasielastic axial vector mass in neutrino interactions on oxygen", Physical Review D **74**, 052002 (2006).

[79] P. Adamson et al., "Study of quasielastic scattering using charged-current ν_μ-iron interactions in the MINOS near detector", Physical Review D **91**, 012005 (2015).

[80] K. Abe et al., "Measurement of the ν_μ charged-current quasielastic cross section on carbon with the ND280 detector at T2K", Physical Review D **92**, 112003 (2015).

[81] M. Martini et al., "Unified approach for nucleon knock-out and coherent and incoherent pion production in neutrino interactions with nuclei", Physical Review C **80**, 065501 (2009).

[82] J. Nieves, I. R. Simo, and M. J. Vacas, "Inclusive charged-current neutrino-nucleus reactions", Physical Review C **83**, 045501 (2011).

[83] J. G. Morfín, J. Nieves, and J. T. Sobczyk, "Recent developments in neutrino/antineutrino-nucleus interactions", Advances in High Energy Physics **2012**, 934597 (2012).

[84] D. Cline et al., "Observation of elastic neutrino-proton scattering", Physical Review Letters **37**, 252–255 (1976).

[85] W. Lee et al., "Observation of the reaction $\nu_\mu + p \to \nu_\mu + p$", Physical Review Letters **37**, 186–189 (1976).

[86] M. Pohl et al., "Study of neutrino proton elastic scattering in the "Gargamelle" freon experiment", Physics Letters B **72**, 489–492 (1978).

[87] J. Horstkotte et al., "Measurement of neutrino-proton and antineutrino-proton elastic scattering", Physical Review D **25**, 2743 (1982).

[88] L. A. Ahrens et al., "Measurement of neutrino-proton and antineutrino-proton elastic scattering", Physical Review D **35**, 785–809 (1987).

[89] A. A. Aguilar-Arevalo et al., "Measurement of the neutrino neutral-current elastic differential cross section on mineral oil at $E_\nu \sim 1$ GeV", Physical Review D **82**, 092005 (2010).

[90] A. A. Aguilar-Arevalo et al., "Measurement of the antineutrino neutral-current elastic differential cross section", Physical Review D **91**, 1–12 (2015).

[91] K. Abe et al., "Measurement of the neutrino-oxygen neutral-current interaction cross section by observing nuclear deexcitation γ rays", Physical Review D **90**, 072012 (2014).

[92] K. Abe et al., "Measurement of neutrino and antineutrino neutral-current quasielasticlike interactions on oxygen by detecting nuclear deexcitation γ rays", Physical Review D **100**, 112009 (2019).

[93] G. Garvey et al., "Role of strange quarks in quasielastic neutrino scattering", Physical Review C **48**, 1919–1925 (1993).

[94] J. Ashman et al., "A measurement of the spin asymmetry and determination of the structure function g_1 in deep inelastic muon-proton scattering", Physics Letters B **206**, 364–370 (1988).

[95] S. Pate and D. Trujillo, "Strangeness vector and axial-vector form factors of the nucleon", in EPJ web of conferences, Vol. 66 (2014), p. 06018.

[96] K. Nishikawa, "Letter of intent for a neutrino oscillation experiment at JHF", KEK Report **35** (2003).

[97] I. Atsuko, *T2K Intranet*, 2020, t2k.org (visited on 08/18/2020).

[98] K. Abe et al., "The T2K experiment", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **659**, 106–135 (2011).

[99] S. Nagamiya, "Introduction to J-PARC", Progress of Theoretical and Experimental Physics **1**, 1–13 (2012).

[100] J-PARC, *Accelerator Layout*, 2011, http://j-parc.jp/Acc/en/layoutE.html (visited on 04/19/2021).

[101] K. Abe et al., "T2K neutrino flux prediction", Physical Review D **87**, 012001 (2013).

[102] K. Abe et al., "First combined measurement of the muon neutrino and antineutrino charged-current cross section without pions in the final state at T2K", Physical Review D **101**, 112001 (2020).

[103] P. A. Amaudruz et al., "The T2K fine-grained detectors", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **696**, 1–31 (2012).

[104] N. Abgrall et al., "Time projection chambers for the T2K near detectors", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **637**, 25–46 (2011).

[105] D. Allan et al., "The electromagnetic calorimeter for the T2K near detector ND280", Journal of Instrumentation **8**, P10019 (2013).

[106] S. Aoki et al., "The T2K Side Muon Range Detector (SMRD)", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **698**, 135–146 (2013).

[107] K. Abe et al., "Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **694**, 211–223 (2012).

[108] K. Abe et al., "Measurement of the muon neutrino charged-current cross sections on water, hydrocarbon and iron, and their ratios, with the T2K on-axis detectors", Progress of Theoretical and Experimental Physics **2019**, 093C02 (2019).

[109] T. Ovsiannikova et al., "The new experiment WAGASCI for water to hydrocarbon neutrino cross section measurement using the J-PARC beam", in Journal of physics: conference series, Vol. 675, 1 (2016), p. 012030.

[110] K. Abe et al., "Measurements of ν_μ and $\nu_\mu + \bar{\nu}_\mu$ charged-current cross-sections without detected pions nor protons on water and hydrocarbon at mean antineutrino energy of 0.86 GeV", Progress of Theoretical and Experimental Physics **2021**, 043C01 (2021), https://doi.org/10.1093/ptep/ptab014.

[111] K. Abe et al., "Measurements of neutrino oscillation in appearance and disappearance channels by the T2K experiment with 6.6×10^{20} protons on target", Physical Review D **91**, 072010 (2015).

[112] K. Abe et al., "J-PARC Neutrino Beamline Upgrade Technical Design Report", 1–127 (2019).

[113] D. Attié et al., "Characterization of resistive Micromegas detectors for the upgrade of the T2K Near Detector Time Projection Chambers", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1025**, 166109 (2022).

[114] K. Abe et al., "T2K ND280 Upgrade - Technical Design Report", (2019), http://arxiv.org/abs/1901.03750.

[115] K. Abe et al., "Measurement of double-differential muon neutrino charged-current interactions on C_8H_8 without pions in the final state using the T2K off-axis beam", Physical Review D **93**, 112012 (2016).

[116] K. Abe et al., "Characterization of nuclear effects in muon-neutrino scattering on hydrocarbon with a measurement of final-state kinematics and correlations in charged-current pionless interactions at T2K", Physical Review D **98**, 032003 (2018).

[117] T2K Internal Technical Note, *T2K-TN-221*, tech. rep. (2015).

[118] T2K Internal Technical Note, *T2K-TN-222*, tech. rep. (2014).

[119] T2K Internal Technical Note, *T2K-TN-001*, tech. rep. (2009).

[120] V. M. Abazov et al., "Evidence for production of single top quarks", Physical Review D **78**, 012005 (2008).

[121] H. J. Yang, B. P. Roe, and J. Zhu, "Studies of boosted decision trees for MiniBooNE particle identification", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **555**, 370–385 (2005).

[122] B. Abi et al., "Neutrino interaction classification with a convolutional neural network in the DUNE far detector", Physical Review D **102**, 092003 (2020).

[123] N. Choma et al., "Graph Neural Networks for IceCube Signal Classification", in Proceedings - 17th IEEE international conference on machine learning and applications, ICMLA 2018 (2019), pp. 386–391.

[124] P. Adamson et al., "Constraints on Oscillation Parameters from ν_e Appearance and ν_μ Disappearance in NOvA", Physical Review Letters **118**, 231801 (2017).

[125] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python", Journal of Machine Learning Research **12**, 2825–2830 (2011).

[126] J. R. Quinlan, "Induction of Decision Trees", Machine Learning **1**, 81–106 (1986).

[127] J. Quinlan, *C4.5: Programs for Machine Learning*, 1st ed. (Morgan Kaufmann, 1992).

[128] L. Breiman et al., *Classification and Regression Trees*, 1st ed. (Chapman and Hall, 2017).

[129] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system", in Proceedings of the acm sigkdd international conference on knowledge discovery and data mining (2016), pp. 785–794.

[130] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: A statistical view of boosting", Annals of Statistics **28**, 337–407 (2000).

[131] T. Hastie et al., "Multi-class AdaBoost", Statistics and Its Interface **2**, 349–360 (2009).

[132] L. Breiman, "Random forests", Machine Learning **45**, 5–32 (2001).

[133] J. H. Ward, "Hierarchical Grouping to Optimize an Objective Function", Journal of the American Statistical Association **58**, 236–244 (1963).

[134] C. Andreopoulos et al., "The GENIE Neutrino Monte Carlo Generator: Physics and User Manual", 10.48550/ARXIV.1510.05494 (2015).

[135] B. G. Tice et al., "Measurement of ratios of ν_μ charged-current cross sections on C, Fe, and Pb to CH at neutrino energies 2-20 GeV", Physical Review Letters **112**, 231801 (2014).

[136] M. Betancourt et al., "Direct Measurement of Nuclear Dependence of Charged Current Quasielasticlike Neutrino Interactions Using MINERvA", Physical Review Letters **119**, 082001 (2017).

[137] T2K Internal Technical Note, *T2K-TN-131*, tech. rep. (2012).

[138] T2K Internal Technical Note, *T2K-TN-080*, tech. rep. (2011).

[139] T2K Internal Technical Note, *T2K-TN-310*, tech. rep. (2016).

[140] F. James and M. Roos, "Minuit - a system for function minimization and analysis of the parameter errors and correlations", Computer Physics Communications **10**, 343–367 (1975).

[141] T. Kutter, *Private Communication*, 2019.

[142] P. C. Hansen and D. P. O'Leary, "The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems", SIAM Journal on Scientific Computing **14**, 1487–1503 (1993).

[143] T2K Internal Technical Note, *T2K-TN-091*, tech. rep. (2011).

[144] K. Abe et al., "Hyper-Kamiokande Design Report", (2018).

[145] K. Abe et al., "Physics potentials with the second Hyper-Kamiokande detector in Korea", Progress of Theoretical and Experimental Physics **2018**, 102 (2018).

[146] The DUNE Collaboration, "Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF", (2015), http://arxiv.org/abs/1512.06148.

[147] The Hyper-Kamiokande Working Group, *The access tunnel excavation for the Hyper-Kamiokande started*, 2021, http://www.hyper-k.org/en/news/news-20210510.html (visited on 09/13/2021).

[148] The Hyper-Kamiokande Working Group, *The production of Hyper-Kamiokande photosensors started*, 2021, http://www.hyper-k.org/en/news/news-20210112.html (visited on 09/13/2021).

[149] K. Abe et al., "Calibration of the Super-Kamiokande detector", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **737**, 253–272 (2014).

[150] L. H. Anthony, "A model independent determination of the π_0 background in Super-Kamiokande in the T2K ν_e appearance measurement", Doctoral Thesis (University of Liverpool, 2020).

[151] Thorlabs, *0.50 NA Step-Index Multimode Fibers*, 2021, https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=362&pn=FP200URT (visited on 04/16/2021).

[152] S. L. Valder, "Measurement of the ν_e CC π+ cross-section using the ND280 tracker and development of optical diffuser calibration systems for Hyper-Kamiokande", Doctoral Thesis (University of Warwick, 2020), p. 186, http://wrap.warwick.ac.uk/160562/.

[153] F. Bociort, "Imaging properties of gradient-index lenses", Doctoral Thesis (1994), pp. 1–178, http://homepage.tudelft.nl/q1d90/fbweb/diss.pdf.

[154] J. Canny, "A Computational Approach to Edge Detection", IEEE Transactions on Pattern Analysis and Machine Intelligence **PAMI-8**, 679–698 (1986).

[155] T2K Internal Technical Note, *T2K-TN-146*, tech. rep. (2013).

APPENDICES

APPENDIX A

DISTRIBUTIONS OF TRAINING VARIABLES OF THE NC0π MULTIVARIATE SIGNAL SELECTION

The variables with the following distributions were used to train the multivariate selection as described in Chapter 7. The sample is comprised of events which pass the preselection given in Table 7.2.

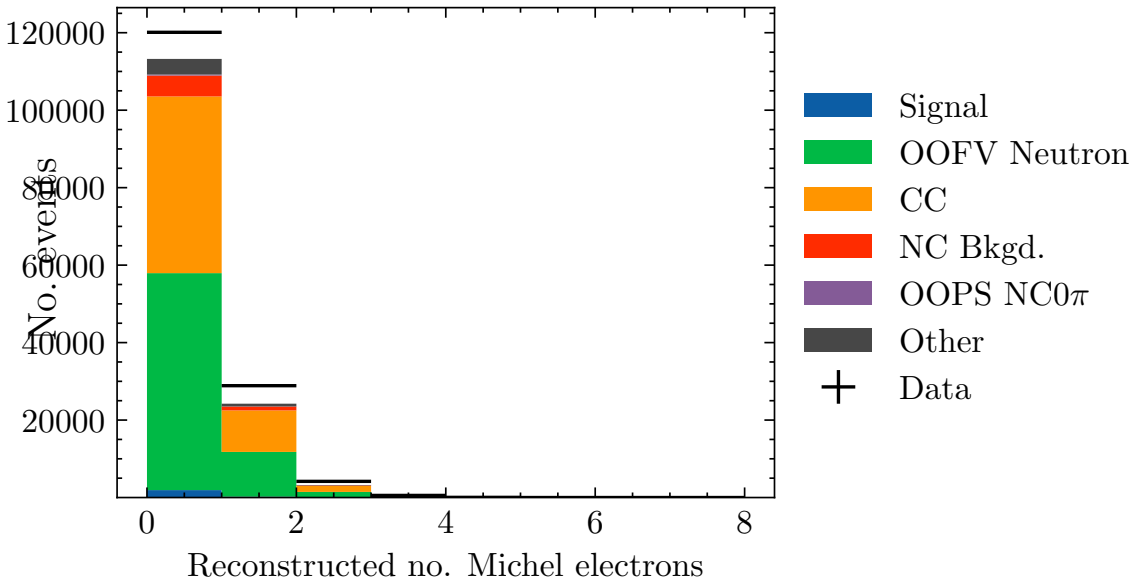

Fig. A.1 Number of reconstructed Michel electrons in FGD1.

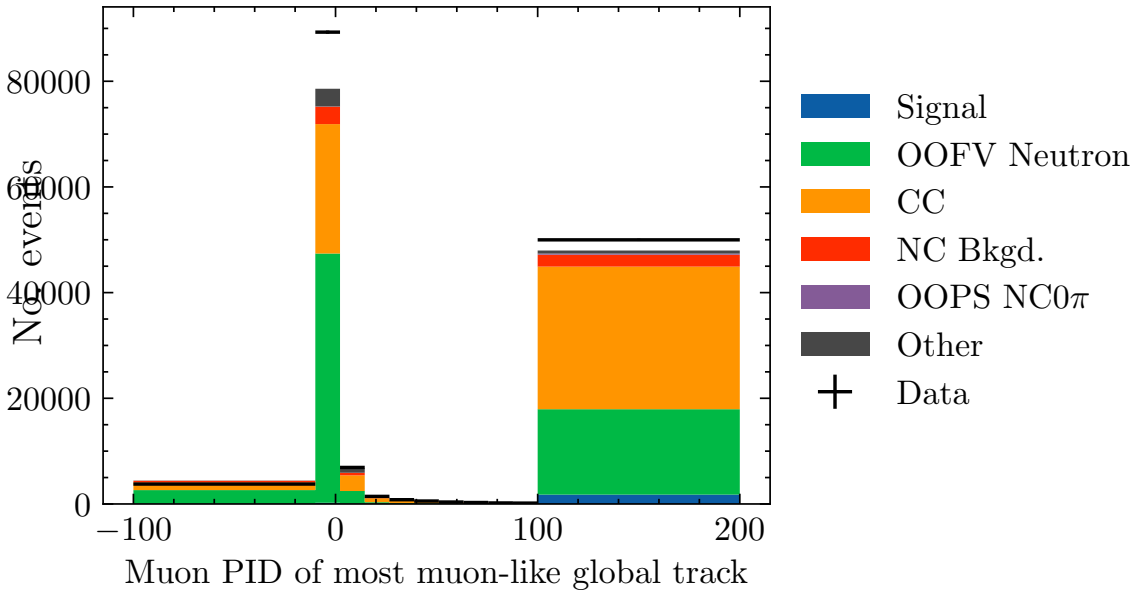

Fig. A.2 Muon PID of most muon-like global track.

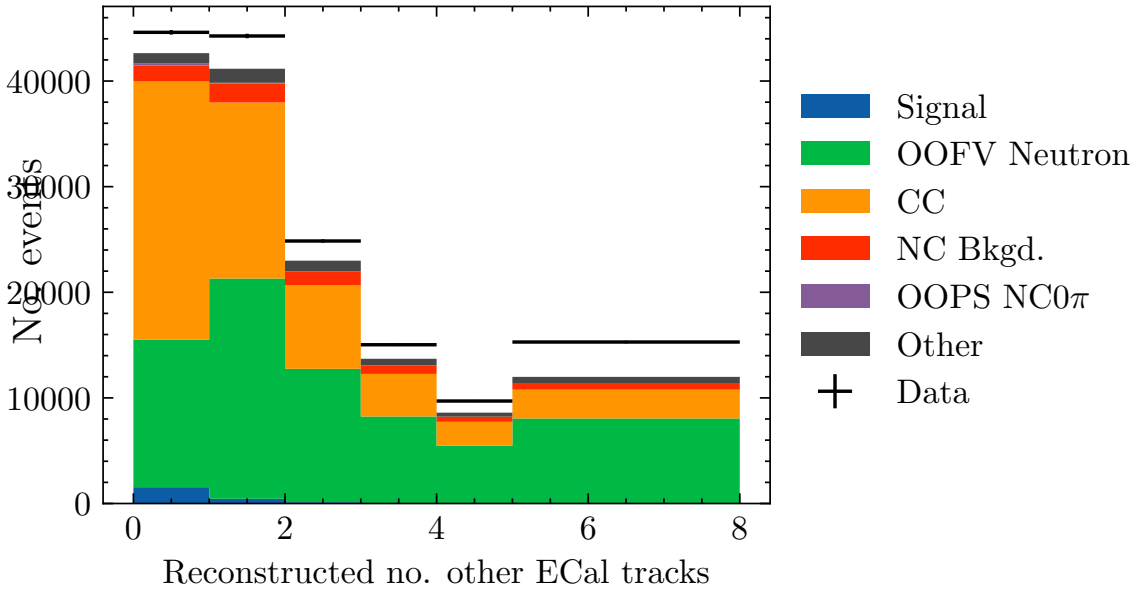

Fig. A.3 Reconstructed no. other ECal tracks

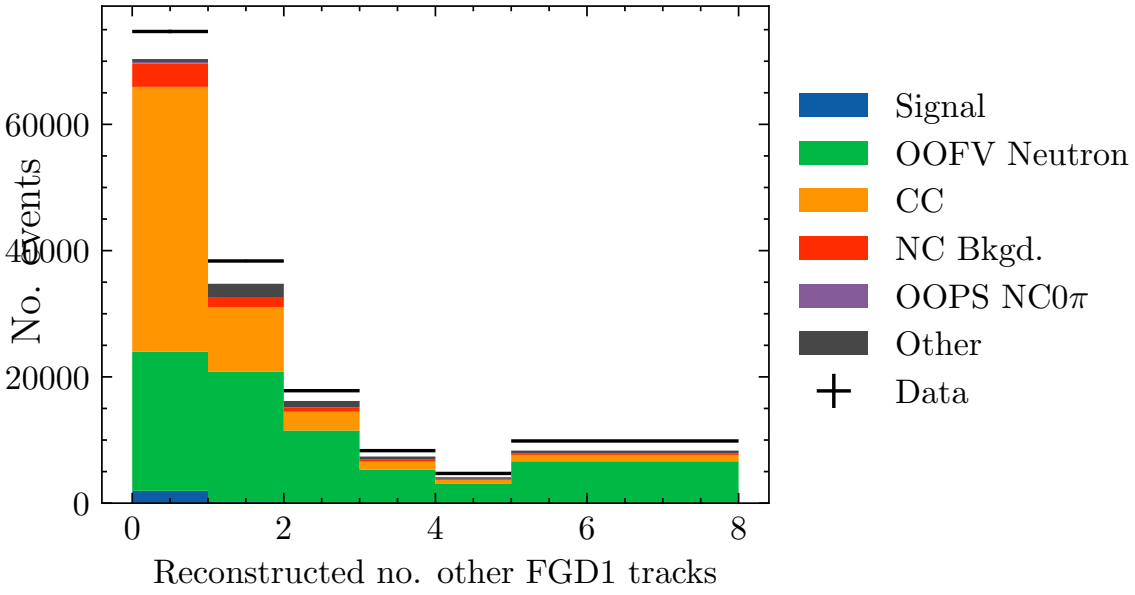

Fig. A.4 Reconstructed no. other FGD1 tracks

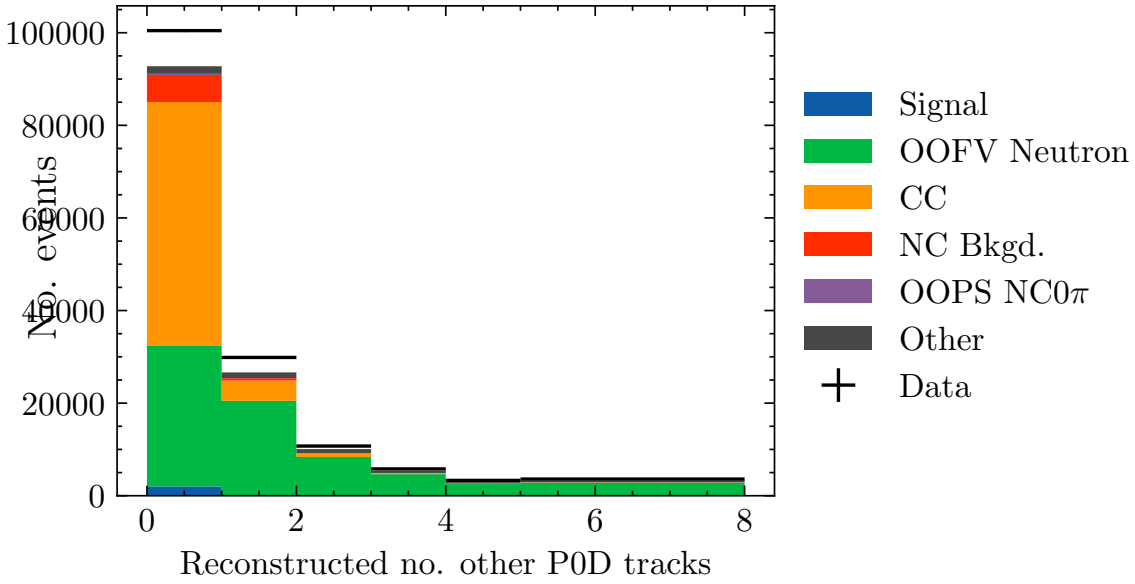

Fig. A.5 Reconstructed no. other P0D tracks

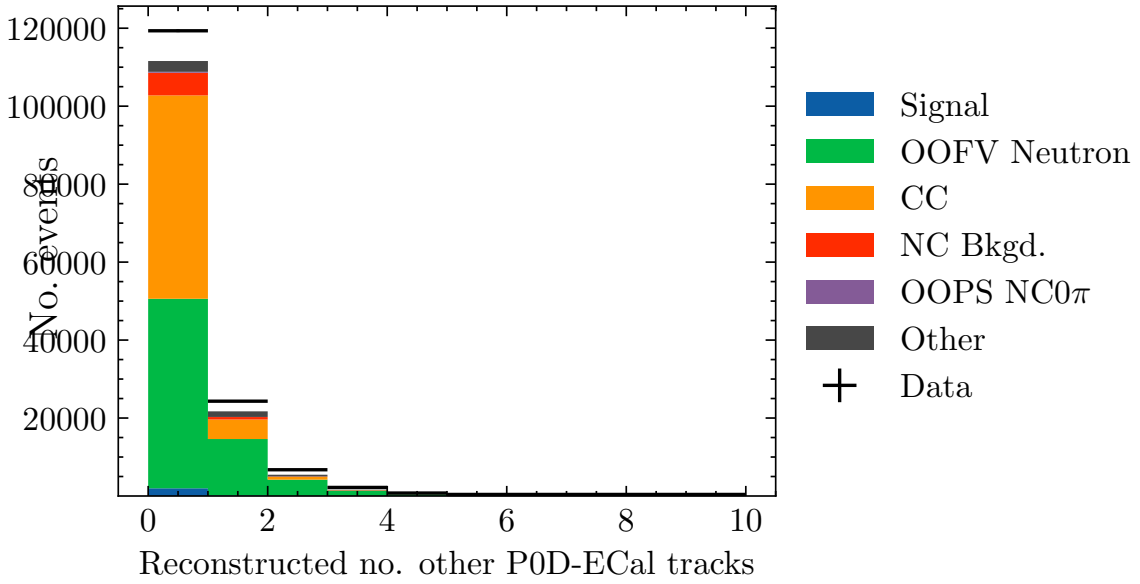

Fig. A.6 Reconstructed no. other P0D-ECal tracks

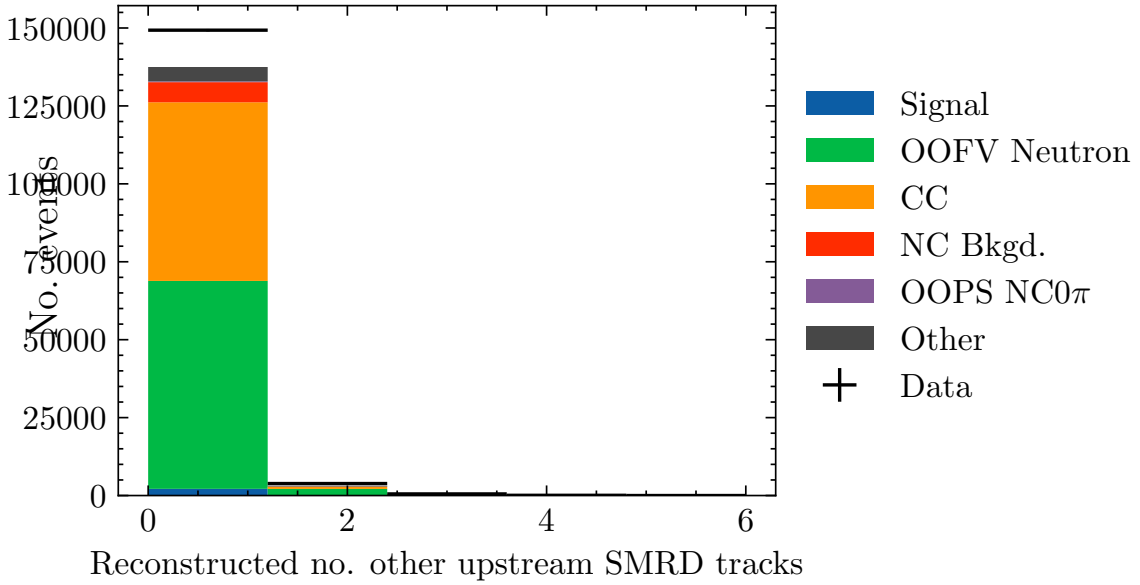

Fig. A.7 Reconstructed no. other upstream SMRD tracks

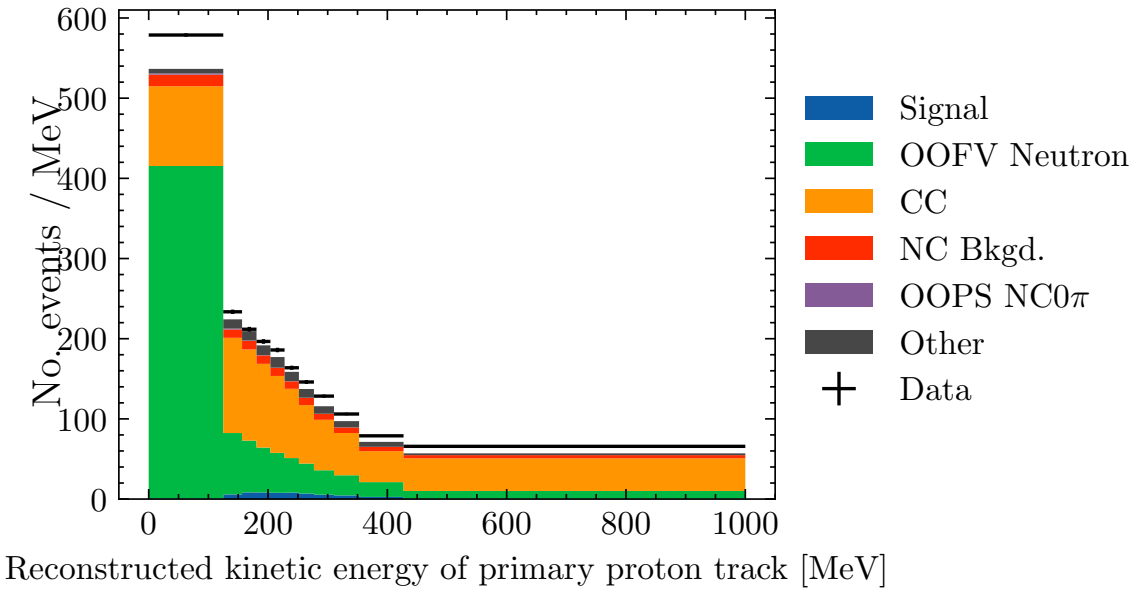

Fig. A.8 Reconstructed kinetic energy of primary proton track

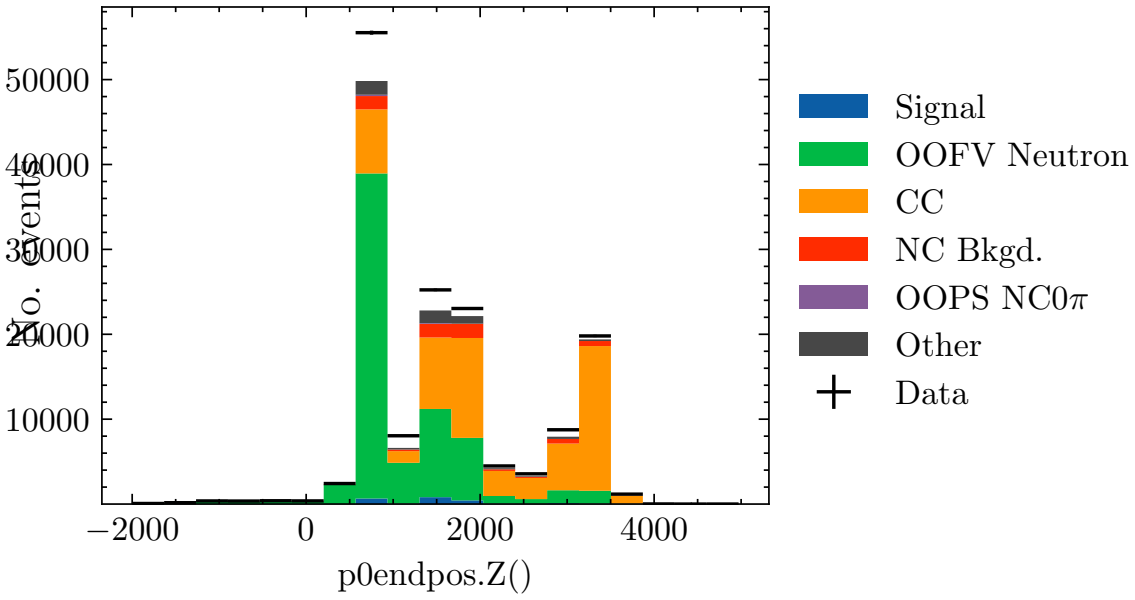

Fig. A.9 Z-axis component of reconstructed primary proton track end position

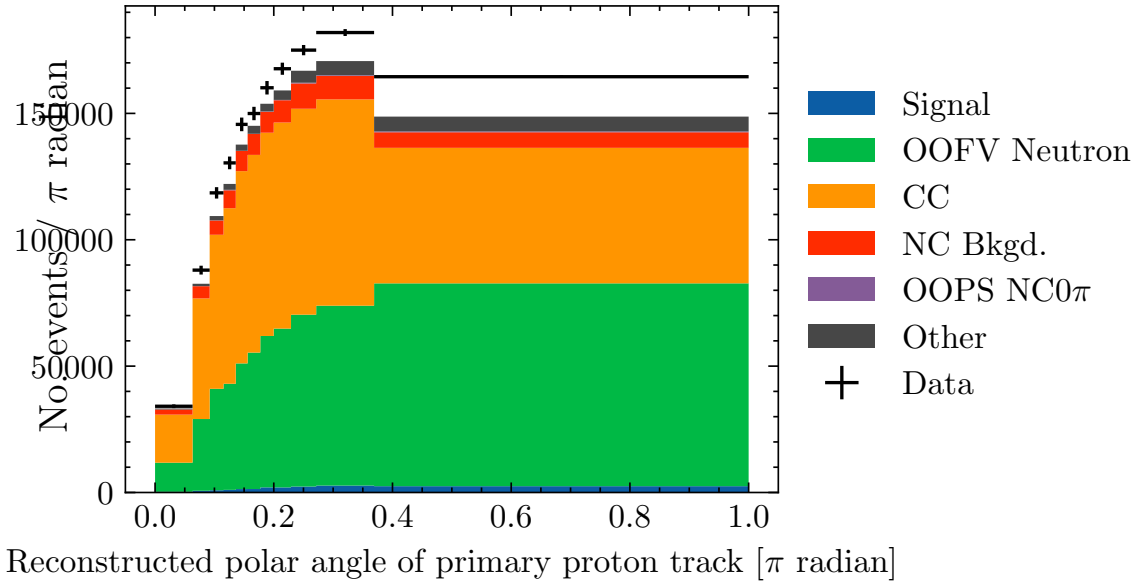

Fig. A.10 Reconstructed polar angle of primary proton track

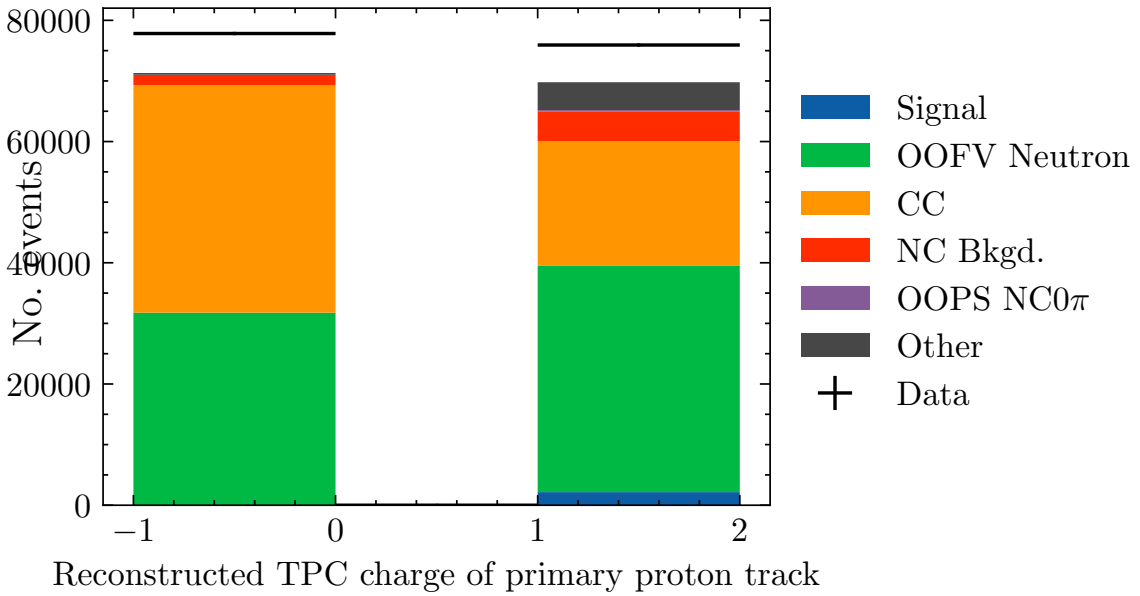

Fig. A.11 Reconstructed TPC charge of primary proton track

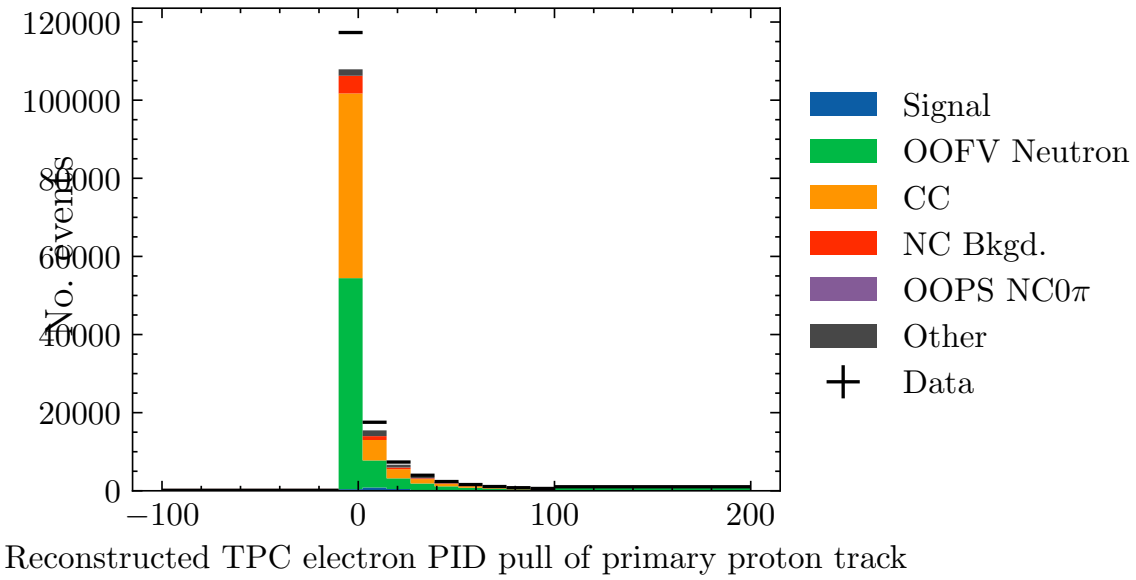

Fig. A.12 Reconstructed TPC electron particle identification pull of primary proton track

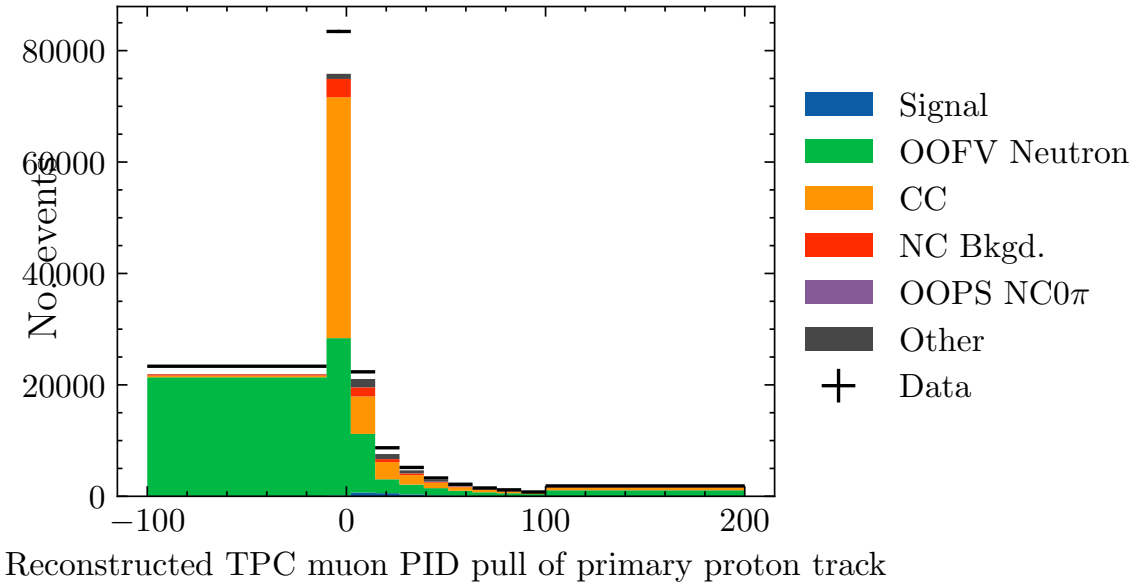

Fig. A.13 Reconstructed TPC muon particle identification pull of primary proton track

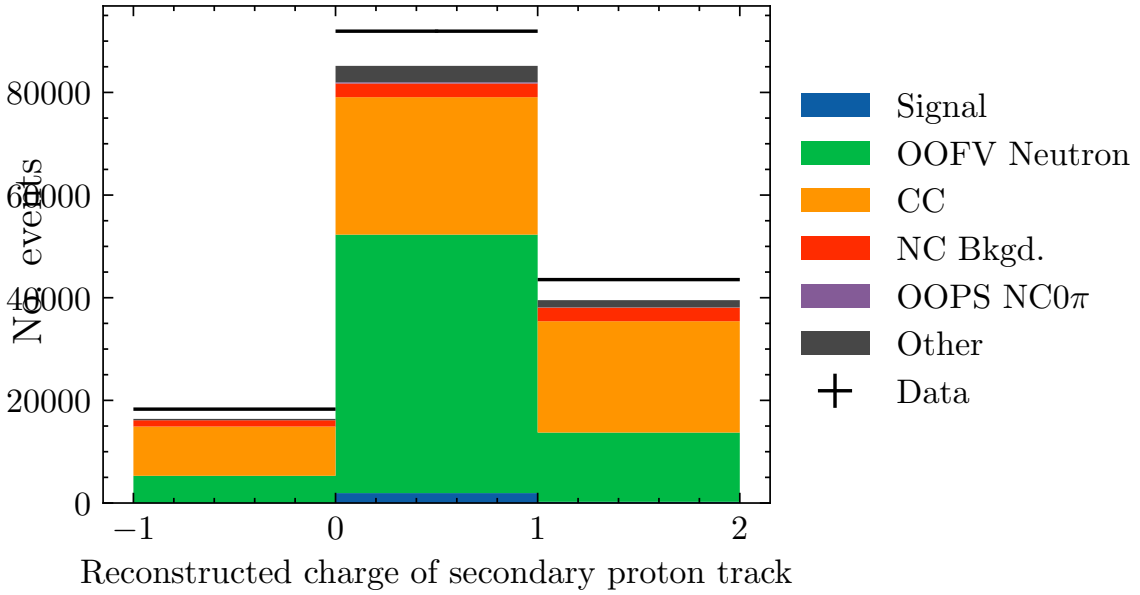

Fig. A.14 Reconstructed charge of secondary proton track

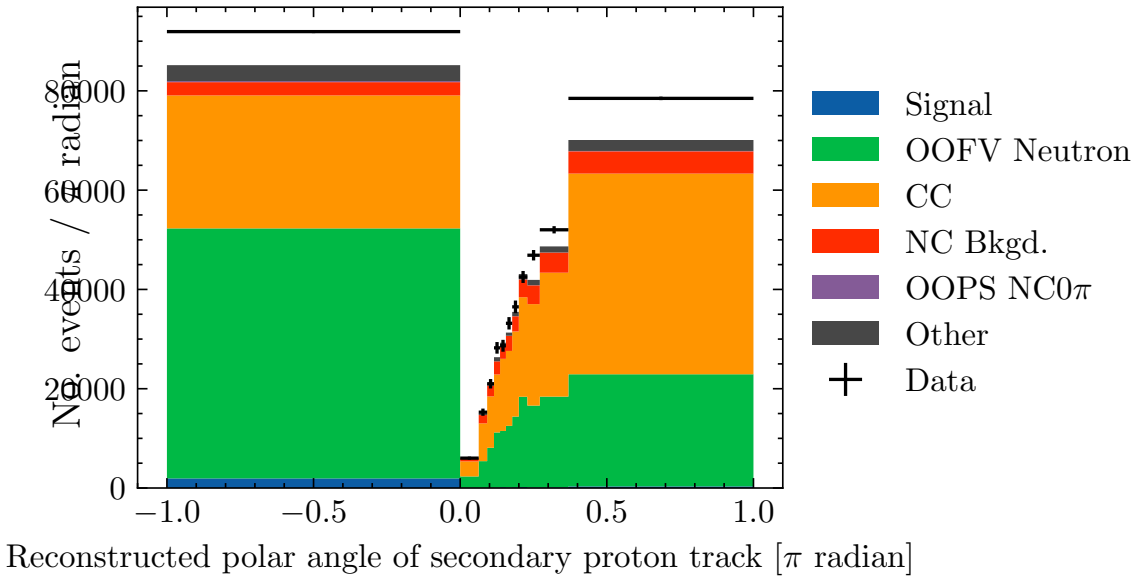

Fig. A.15 Reconstructed polar angle of secondary proton track