

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Edward John Peter Larkin

1258942

2. THESIS DEPOSIT

2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.

2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service. [At present, theses submitted for a Masters degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]

2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the Guide to Examinations for Higher Degrees by Research.)

2.4 (a) Hard Copy I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers immediately.

I agree that my thesis may be photocopied.

(b) Digital Copy I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.

My thesis can be made publicly available online.

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following: Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hardbound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:


- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature:Date:



**A Cross Section Measurement From Neutrino
Interactions on Argon Gas**

by

Edward John Peter Larkin

Thesis

Submitted to the University of Warwick

for the degree of Physics

Doctor of Philosophy

Department of Physics

March 2016

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments	vii
Declarations	viii
Abstract	ix
Abbreviations	x
List of Figures	xvi
List of Tables	xx
Chapter 1 Historical Background	1
1.1 The search for neutrinos	2
1.1.1 Cowan–Reines neutrino experiment	2
1.1.2 Muon neutrinos at the AGS	2
1.1.3 DONUT and tau neutrinos	3
1.2 Solar and atmospheric anomalies	3
1.2.1 Chlorine as a detector	3
1.2.2 The Homestake experiment	4
1.2.3 Further solar neutrino experiments	5
1.2.4 The atmospheric anomaly	6
Chapter 2 Theoretical Background	10
2.1 Neutrino masses	10
2.1.1 Mass terms	11
2.1.2 Mass terms and hierarchy	12
2.2 Neutrino oscillation physics	12
2.2.1 Two neutrino case	13
2.2.2 Solar neutrinos and the MSW effect	15

2.2.3	The PMNS mixing matrix	18
2.2.4	Three flavour oscillation	18
2.2.5	Specific applications	21
2.2.6	Mass hierarchy, CP violation and octant	23
2.3	Neutrino interactions	24
2.3.1	Elastic scattering	24
2.3.2	Charged current quasi-elastic interactions	26
2.3.3	2p2h	28
2.3.4	Resonant pion production	29
2.3.5	Coherent processes	30
2.3.6	Deep inelastic scattering	30
2.3.7	Final state interactions	31
Chapter 3 Contemporary Neutrino Physics		33
3.1	Neutrino oscillation experiments	34
3.1.1	Resolution of solar and atmospheric anomalies	34
3.1.2	Reactor neutrinos	35
3.1.3	Accelerator neutrinos	37
3.1.4	Solar neutrinos	41
3.1.5	Atmospheric neutrinos	42
3.1.6	Global fit to mixing parameters	42
3.2	Neutrino interactions with matter	42
3.2.1	MiniBooNE results	44
3.2.2	Argon TPCs	44
Chapter 4 The T2K Detectors		47
4.1	Beam	48
4.1.1	Off-axis configuration	48
4.1.2	Accelerator	50
4.1.3	Secondary beamline	51
4.1.4	Muon flux	51
4.2	Super-Kamiokande	52
4.2.1	The detector	53
4.2.2	Particle identification	53
4.3	INGRID	54
4.3.1	The detector array	55
4.3.2	The proton module	55
4.4	ND280	55

4.4.1	The UA1 magnet	57
4.4.2	SMRD	57
4.4.3	ECals	58
4.4.4	P0D	59
4.4.5	FGDs	59
4.4.6	TPCs	60
Chapter 5 T2K Software		63
5.1	Simulating neutrino interactions	63
5.1.1	NEUT	64
5.1.2	GENIE	65
5.1.3	Disagreements between generators	66
5.2	Interactions outside the ND280	66
5.3	ND280 simulation	67
5.4	ND280 electronics simulation	67
5.5	ND280 calibration	68
5.6	ND280 reconstruction	68
5.6.1	TPC reconstruction	69
5.6.2	FGDs	69
5.6.3	Tracker	69
5.6.4	P0D	70
5.6.5	ECals	70
5.6.6	SMRD	72
5.6.7	Global reconstruction	73
5.7	Analysis	73
Chapter 6 TREx Reconstruction		74
6.1	Framework	74
6.1.1	TREx base classes	75
6.2	Hit preparation	76
6.3	Pattern recognition	78
6.3.1	The A* algorithm	79
6.3.2	Cell definition	83
6.3.3	Charge and occupancy flags	84
6.3.4	Pattern object preparation	85
6.3.5	Edge detection	86
6.3.6	Junction object detection	88
6.3.7	Path object formation	90

6.3.8	Cell association	91
6.3.9	Kink finding	91
6.3.10	Determining cluster object orientation	94
6.3.11	Interpolation and extrapolation	96
6.3.12	Horizontal and vertical clustering	97
6.3.13	Cell association corrections	97
6.3.14	Junction object expansion	99
6.3.15	Translation	100
6.4	First pass seeding	100
6.5	Matching and merging	101
6.5.1	Horizontal gap merge	101
6.5.2	Vertical gap merge	101
6.6	t_0 determination	102
6.6.1	Default t_0	102
6.6.2	t_0 from constituents	103
6.6.3	Fitting from x - z hits	103
6.6.4	Fitting from y - z hits	103
6.6.5	Cathode crossers	104
6.7	Tracking	106
6.7.1	Cluster correction	106
6.7.2	Likelihood fit	106
6.8	PID	106
6.9	Likelihood matching	107
6.9.1	Intra-pattern matching	108
6.9.2	Inter-pattern matching	109
6.10	Likelihood merging	109
6.11	Cathode merging	109
6.12	Tracker and global reconstruction	110
6.13	Validation	110
Chapter 7 Selection		114
7.1	Proto-analysis	114
7.1.1	Preselection	114
7.1.2	Main selection	115
7.1.3	Issues raised	115
7.1.4	Improvements since the proto-analysis	116
7.2	This analysis	116

7.2.1	Cut flow	116
7.3	Selection performance	125
Chapter 8 Systematic uncertainties		127
8.1	Corrections	127
8.1.1	Gas Monte Carlo correction	127
8.1.2	Proton range correction	128
8.1.3	PID corrections	128
8.1.4	Momentum resolution corrections	129
8.2	Propagation of systematic uncertainties	129
8.2.1	Variations	130
8.2.2	Weights	130
8.2.3	Efficiencies	130
8.3	Event level systematic uncertainties	131
8.3.1	Sand and cosmic muon background	131
8.3.2	Flux related systematic uncertainties	132
8.4	Standard TPC variation uncertainties	133
8.4.1	Momentum scale	133
8.4.2	Momentum resolution	133
8.4.3	B field distortion	133
8.4.4	PID	134
8.5	Standard TPC efficiency uncertainties	134
8.5.1	Track efficiency	134
8.5.2	Cluster efficiency	136
8.5.3	Charge ID efficiency	137
8.6	Gas interaction specific uncertainties	137
8.6.1	Vertex uncertainties	137
8.6.2	t_0 determination	141
8.6.3	Hairy track systematic uncertainties	143
8.7	Out of fiducial volume uncertainties	146
8.7.1	Rate uncertainties	146
8.7.2	Bad t_0 uncertainties	146
8.7.3	Broken track uncertainties	146
Chapter 9 Argon Cross Section		148
9.1	Single bin analysis	148
9.1.1	Methodology	148
9.1.2	Predictions from Monte Carlo	149

9.1.3	Cross section calculation	150
9.2	Distributions in other variables	151
9.2.1	Muon kinematics	151
9.2.2	Secondary particle kinematics	151
9.3	Hand scan	155
9.3.1	Methodology	155
9.3.2	Results	156
9.3.3	Background analysis	157
9.3.4	Overview of signal	162
Chapter 10 Conclusions		166

Acknowledgments

I would like to start by thanking my supervisors Dr Gary Barker and Dr Steve Boyd without whose help and support none of this would have been possible.

I would also like to thank members of office P450, past and present. Particular thanks go to Dr David Hadley who took on the unenviable task of guiding me through the T2K software during the start of my PhD.

During my work on T2K I have enjoyed the assistance and fellowship of a large portion of our collaborators. Specific thanks go to those I spent extended time with during my long term placement.

Final thanks go to my friends and family whose support certainly helped smooth over any rough patches during the course of my PhD.

For these acknowledgements I have attempted to avoid naming people specifically, to mitigate any offence caused when I inevitably forget somebody. If your name is not here but you feel it should be, please send me an e-mail and I shall felt tip it into my copy at the earliest convenient opportunity.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

Chapter 1, Chapter 2 and Chapter 3 provide historical, theoretical and experimental background to the results presented in this thesis and are the result of a review of the literature.

Chapter 4 and Chapter 5 cover the T2K experiment's hardware and software specifically. These details are derived from published documents, internal documentation and personal communications with others in the collaboration.

The `TREx` software I was involved in developing is discussed in Chapter 6. This software was primarily developed by myself and Dr Anthony Hillairet in collaboration with Lukas Koch, Dr Tom Feusels, Dr Stefania Bordoni, Dr Pip Hamilton and others. The pattern recognition algorithms detailed therein are my own original work where not specifically cited otherwise. The remainder of `TREx` was primarily developed by Dr Anthony Hillairet building on existing T2K software.

The selection and systematic uncertainties described in Chapter 7 and Chapter 8 were developed by myself and Lukas Koch. Some uncertainties were based on computation by other collaborators as stated in the text. This work would not have been possible without the gas interaction proto-analysis by Dr Pip Hamilton. The analysis covered in this thesis used the HighLAND2 framework developed by a number of T2K collaborators.

Finally the results presented in Chapter 9 are entirely my own, derived for the purposes of this thesis, and should not be taken as official T2K results.

Abstract

T2K is a neutrino oscillation experiment using an off-axis beam of muon neutrinos to study rates of $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. We use its ND280 near detector to make measurements of neutrino interactions in argon gas. Long baseline neutrino experiments are approaching the point of being limited by systematic uncertainties on neutrino-nucleus cross sections rather than only statistics. This makes it imperative for a proper understanding of such processes to reduce their uncertainties. A gas interaction analysis provides a way to gain such an understanding.

This thesis describes the first successful neutrino–nucleus cross section measurement performed on gas. The **TREx** reconstruction algorithms developed for this goal feature a novel path finding based pattern recognition algorithm and have now become the official TPC reconstruction software for the T2K experiment. The analysis itself uses real data and Monte Carlo based on the NEUT and GENIE neutrino interaction generators. It extracts a cross section for charged current quasi-elastic ν_μ interactions on argon, integrated over the T2K flux, of $\sigma_{\nu_\mu\text{CC}}^{40\text{Ar}} = (4.19 \pm 0.70(\text{stat}) \pm 1.04(\text{sys})) \times 10^{-39} \text{ cm}^2$ or $\sigma_{\nu_\mu\text{CC}}^{40\text{Ar}} = (4.61 \pm 0.78(\text{stat}) \pm 0.94(\text{sys})) \times 10^{-39} \text{ cm}^2$, depending on the simulation used for efficiency and purity calculations.

The distributions of these results with respect to muon and proton kinematics and proton multiplicity are discussed in the context of differences between the two generators. The ND280 TPCs have the advantage of detailed reconstruction which affords unique opportunities to test the agreement with data of nuclear models used in simulation. Future generations of this analysis are expected to produce more varied and precise measurements in this vein.

Abbreviations

The following abbreviations are used through the thesis and presented here for clarity:

- **AGKY** model: Model for hadronisation.
- **AGS**: Alternating Gradient Synchrotron.
- **ArgoNeuT**: ARGOn NEUtrino Test.
- **BEBC**: Big European Bubble Chamber.
- **BNL**: Brookhaven National Laboratory.
- **Borexino**: Diminutive of BORon solar neutrino EXperiment (BOREX).
- **CC**: Charged Current.
- **CCFR**: Chicago-Columbia-Fermilab-Rochester.
- **CCQE**: Charged Current Quasi-Elastic.
- **CDHSW**: Cern-Dortmund-Heidelberg-Saclay-Warsaw.
- **CERN**: Conseil Européen pour la Recherche Nucléaire, now European Organization for Nuclear Research.
- **CHARM**: CERN-Hamburg-Rome-Moscow.
- **CHORUS**: CERN Hybrid Oscillation Research apparatus.

- **CHOOZ**: Not actually an acronym despite often being capitalised; comes from the Chooz region and nuclear power plant.
- **CORSIKA**: COsmic Ray SIMulations for KAscade.
- **DIS**: Deep Inelastic Scattering.
- **DONUT**: Direct Observation of the NU Tau.
- **Double-CHOOZ**: See **CHOOZ**.
- **DUNE**: Deep Underground Neutrino Experiment.
- **ECal**: Electromagnetic CALorimeter.
- **EM**: ElectroMagnetic.
- **E531**: Generic Fermilab experiment number.
- **FGD**: Fine Grained Detector.
- **FLUKA**: FLUktuierende KAskade.
- **FNAL**: Fermi National Accelerator Laboratory.
- **GALLEX**: GALLium EXperiment.
- **GEANT / Geant**: GEometry ANd Tracking.
- **GENIE**: Generates Events for Neutrino Interaction Experiments.
- **GNO**: Gallium Neutrino Observatory.
- **GRV98**: Glück-Reya-Vogt 1998.
- **HIP**: Highly Ionising Particle.
- **HighLAND**: HIGH Level Analysis at the Near Detector.
- **Hyper-Kamiokande**: See **Kamiokande**.
- **HV cluster**: Horizontal or Vertical cluster.

- **ICAL**: Iron CALorimeter.
- **ICARUS T600**: Imaging Cosmic And Rare Underground Signals (ICARUS) detector filled with ~ 600 tonnes of liquid argon.
- **ILL**: Institut Laue-Langevin.
- **IMB**: Irvine-Michigan-Brookhaven.
- **INGRID**: Interactive Neutrino GRID.
- **INO**: India-based Neutrino Observatory.
- **J-PARC**: Japan Proton Accelerator Research Complex.
- **JNUBEAM**: J-PARC ν BEAM.
- **JUNO**: Jiangmen Underground Neutrino Observatory.
- **Kamiokande**: KAMIOKA Nucleon Decay Experiment.
- **KamLAND**: KAMioka Liquid scintillator Anti Neutrino Detector.
- **KARMEN**: KARlsruhe Rutherford Medium Energy Neutrino.
- **KEK**: High Energy Accelerator Research Organization (the abbreviation works better in Japanese).
- **KM3NeT**: km³ NEutrino Telescope.
- **K2K**: KEK-to-Kamioka (see **KEK**).
- **LArTPC**: Liquid ARgon TPC.
- **LAMPF**: Los Alamos Meson Physics Facility.
- **LINAC**: LINear ACcelerator.
- **LSND**: Liquid Scintillator Neutrino Detector.
- **MACRO**: Monopole, Astrophysics and Cosmic Ray Observatory.

- **MINOS**: Main Injector Neutrino Oscillation Search.
- **MINOS+**: See **MINOS**.
- **MEC**: Meson Exchange Current.
- **MiniBooNE**: (Mini-) BOOster Neutrino Experiment.
- **MicroBooNE**: (Micro-) BOOster Neutrino Experiment.
- **MINERVA**: Main INjector ExpeRiment for ν -A.
- **MINUIT**: A tool for finding the minimum value of a multi-parameter function.
- **MIP**: Minimum Ionising Particle.
- **MPPC**: Multi-Pixel Photon Counter.
- **MR**: Main Ring.
- **MSW** effect: Mikheyev-Smirnov-Wolfenstein effect.
- **MUMON**: MUon MONitor.
- **mwe**: Meter Water Equivalent.
- **NA61/SHINE**: CERN experiment NA61; SPS Heavy Ion and Neutrino Experiment.
- **NC**: Neutral current.
- **ND280**: Near Detector at 280 m.
- **NEOS**: NEutrino Oscillation at Short baseline.
- **NEUT**: Neutrino interaction generator used by T2K.
- **NOMAD**: Neutrino Oscillation MAgnetic Detector.
- **NuMi**: ν at the Main Injector.

- **NO ν A**: NuMi Off-axis ν_e Appearance.
- **n ν nh**: n Proton n Hole.
- **NUSEX**: NUcleon Stability EXperiment.
- **NuTeV**: ν at the TEVatron.
- **OPERA**: Oscillation Project with Emulsion tRacking Apparatus.
- **ORCA**: Oscillation Research with Cosmics in the Abyss.
- **POD**: Pi-Zero Detector.
- **PAI**: PhotoAbsorption Ionisation model.
- **PDF**: Probability Density Function.
- **PID**: Particle IDentification.
- **PINGU**: Precision IceCube Next Generation Upgrade.
- **PMNS** matrix: Pontecorvo-Maki-Nakagawa-Sakata matrix.
- **PROSPECT**: Precision Reactor Oscillation and SPECTrum.
- **PYTHIA**: Code for simulating hadronic interactions.
- **PYTHIA/JETSET**: Combination of two sets of code for simulating hadronic interactions.
- **QE**: Quasi-Elastic.
- **RCS**: Rapid Cycling Synchrotron.
- **RENO**: Reactor Experiment for Neutrino Oscillation.
- **RENO-50**: RENO (see **RENO**), located ~ 50 km from a reactor.
- **ROOT**: Not an acronym despite capitalisation.
- **SAGE**: Soviet-America Gallium Experiment.

- **sbc**at: SciBar Cellular AuTomaton.
- **SBND**: Short Baseline Near Detector.
- **SciBooNE**: SciBar BOOster Neutrino Experiment.
- **SciBar**: SCIntillator BAR.
- **SMRD**: Side Muon Range Detector.
- **SNU**: Solar Neutrino Unit.
- **SNO**: Sudbury Neutrino Observatory.
- **SNO+**: See **SNO**.
- **SoLid**: Short baseline Oscillation search with Lithium-6 Detector.
- **SOX**: Short distance neutrino Oscillations with BOREXINO, see **BOREXINO**.
- **SPS**: Super Proton Synchrotron.
- **SRIM**: Stopping Range of Ions in Matter.
- **Stereo**: Search for sterile neutrinos at ILL reactor.
- **Super-Kamiokande**: See **Kamiokande**.
- **T2K**: Tōkai-to-Kamioka.
- **TPC**: Time Projection Chamber.
- t_0 : Absolute time used to calculate x position in an ND280 TPC.
- **TREx**: TPC Reconstruction EXtension.
- **UA1**: Underground Area experiment 1, the CERN experiment from which the ND280 magnet was borrowed.
- **WLS** fibre: WaveLength Shifting fibre.
- **2p2h**: Two Proton Two Hole.

List of Figures

1.1	Solar neutrino energy spectrum using solar model BS05(OP)	5
1.2	Kamiokande solar neutrino results	7
1.3	SNO ^8B neutrino results	7
1.4	Super-Kamiokande atmospheric neutrino results	8
2.1	Normal and inverse neutrino mass hierarchy	12
2.2	Mixing angles as a function of matter electron number density	17
2.3	Total charged current cross sections per nucleon	25
2.4	Elastic scattering of (anti) neutrinos on electrons and nucleons	25
2.5	Quasi-elastic scattering of (anti) neutrinos on nucleons	26
2.6	Feynman diagram for 2p2h processes	28
2.7	Resonant pion production with (anti) neutrinos on nucleons	29
2.8	Deep inelastic scattering of (anti) neutrinos on nucleons	31
3.1	KamLAND results	35
3.2	Daya Bay results	37
3.3	LSND anomaly	38
3.4	MiniBooNE low energy excess	39
3.5	MINOS results for atmospheric mixing	40
3.6	Cross sections data and M_A for MinoBooNE and NOMAD	44
3.7	ArgoNeuT back to back proton event	45
4.1	The T2K experiment layout	48
4.2	The T2K beamline	49
4.3	Neutrino oscillation probability as a function of energy	49
4.4	Effect of off-axis angle on J-PARC fluxes	49
4.5	The J-PARC neutrino facility physical location	50
4.6	The Super-Kamiokande detector	52
4.7	The INGRID on-axis detector	54

4.8	The INGRID modules	55
4.9	The ND280 off-axis near detector	56
4.10	The ND280 P0D	59
4.11	A ND280 TPC	60
4.12	Hits left by particles passing through the TPCs	62
5.1	Differences in proton properties between GENIE and NEUT	66
6.1	Data structure used for TREx reconstruction	76
6.2	Illustration of TREx hit preparation	77
6.3	Overview of TREx pattern recognition	78
6.4	Illustration of the A* algorithm	81
6.5	Illustration of the issues with different connect distances	82
6.6	Illustration of TREx flags for non-physical effects	84
6.7	Illustration of formation of pattern objects	86
6.8	Illustration of the finding of edges in a pattern object	87
6.9	Illustration of the second pass finding for edges in a pattern object	89
6.10	Illustration of junction object detection	90
6.11	Illustration of path object construction in events with multiple junctions	91
6.12	Illustration of cell association	92
6.13	Illustration of TREx kink finding	93
6.14	Illustration of cluster object angle determination via the dichotomy technique	95
6.15	Illustration of path object extrapolation and interpolation	96
6.16	Illustration of TREx clustering	97
6.17	Illustration of cell association corrections	98
6.18	Illustration of junction object expansion	99
6.19	Illustration of TREx matching and merging	102
6.20	Illustration of t_0 determination with $x-z$ hits in an upstream detector	105
6.21	Illustration of t_0 determination with the side ECal	105
6.22	Illustration of t_0 determination with $y-z$ hits in an upstream detector	105
6.23	Illustration of $\frac{dE}{dx}$ for particles produced by the ND280 beam	107
6.24	Illustration of TREx likelihood matching	108
6.25	Illustration of TREx performance compared to old TPC reconstruction	110
6.26	Illustration of TREx track efficiency and purity compared to old TPC reconstruction	112
6.27	Illustration of TREx vertex position accuracy	112
6.28	Illustration of TREx vertex multiplicity	112

6.29	Momentum residual distributions for old reconstruction and TREx	113
7.1	Muon PID in selection	120
7.2	MIP PID in selection	120
7.3	Momentum cut in selection	121
7.4	x -angle cut in selection	121
7.5	Projected vertex position calculation	122
7.6	Likelihood matching cut in selection	123
7.7	Passing by tracks	124
7.8	Efficiency and purity of selection cuts	125
8.1	Proton ranges in simulation and data and resulting corrections	129
8.2	T2K flux uncertainties	132
8.3	TPC efficiency as a function of momentum	135
8.4	Normal and sub-optimal	138
8.5	Vertex efficiency systematic uncertainties for two- and multi-track events	138
8.6	Vertex resolution systematic uncertainties for two- and multi-track events	139
8.7	Secondary particle reconstruction efficiency uncertainties at different true trajectory lengths	140
8.8	t_0 candidates in x , y and z	141
8.9	t_0 efficiencies for tracker detectors using beam data and Monte Carlo	142
8.10	Two hairy events as handled by reconstruction	143
8.11	Normal and artificially hairy version of the same gas interaction event	144
8.12	Hairy track muon efficiency systematic	145
8.13	Vertex resolution systematic uncertainties for hairy track events	145
8.14	Broken track uncertainties	147
9.1	Selected events in real data, GENIE and NEUT	150
9.2	Selected muon momentum in real data and Monte Carlo	152
9.3	Selected muon angle in real data and Monte Carlo	152
9.4	Selected vertex multiplicity in real data and Monte Carlo	154
9.5	Selected vertex proton multiplicity in real data and Monte Carlo	154
9.6	Selected highest proton momentum in real data and Monte Carlo	154
9.7	Event displays for undetermined events	155
9.8	Multiplicity of charged secondary particles from hand scan	156
9.9	Hand scan events of difficult to determine multiplicity	159

9.10	Hand scan events with backwards going products	159
9.11	Hand scan background events with delta-rays	159
9.12	Hand scan background events with colinear tracks	160
9.13	Hand scan background events with overlapping tracks	160
9.14	Hand scan background events with messy out of fiducial volume in- teractions	160
9.15	Hand scan background events with unassociated hits	161
9.16	Hand scan background events (misc)	161
9.17	Hand scan signal events with two tracks	163
9.18	Hand scan signal events with back to back tracks	163
9.19	Hand scan signal events with hammer-like topology	163
9.20	Hand scan signal event with secondary interaction close to vertex . .	164
9.21	Hand scan signal events with high multiplicity	164
9.22	Hand scan signal events with hairy appearance	164
9.23	Hand scan signal events with high vertex activity	165

List of Tables

3.1	Global fit values for neutrino oscillation parameters	43
8.1	TPC gas mixture corrections	128
8.2	TPC track finding efficiencies	135
8.3	TPC cluster finding efficiencies as a function of angle	136
8.4	Summary of vertex uncertainties	140
8.5	t_0 efficiency difference between data and Monte Carlo for different sources	142
8.6	Summary of hairy track uncertainties	145
9.1	Signal and background events found by hand scan	156




Figure 1: The first signal event seen during this analysis.

Chapter 1

Historical Background

The history of the theory and discovery of neutrinos has been discussed in a plethora of works for both lay and scientific audiences, (see for example [1][2][3]).

Wolfgang Pauli first postulated the existence of what is now known as the neutrino in 1930. At the time his ‘desperate remedy’ was needed to patch up some problems with energy conservation and spin statistics in β decay. The energy conservation issue was most troubling to physicists at the time; the products of two body decays should be monochromatic but β decay electrons were observed with a continuous spectrum of energies.

The neutrino was by no means the only proposed explanation. Niels Bohr was prepared to accept violation of conservation of energy rather than the addition of a third wheel to the existing model of protons and electrons[4] and more exotic models were also proposed. Pauli himself was somewhat reluctant to publish the proposal, mainly discussing it at conferences.

A new electrically neutral particle became less unwelcome following the discovery of the neutron, as a result of which Pauli’s particle was called the ‘neutrino’. Enrico Fermi’s β decay theory[5][6] incorporated this neutrino. Despite growing theoretical acceptance it looked like the neutrino was to all intents and purposes invisible from an experimental standpoint. Pauli himself wagered a box of champagne against its detection. It would be a quarter of a century before he had to pay up.

Since they were first proposed, neutrinos have necessitated increasingly clever and intricate experiments as researchers attempt to uncover more and more of their properties. This chapter looks at noteworthy experiments leading up to the discovery of neutrino oscillations, as well as providing background and motivation to the theories discussed in Chapter 2 and modern experiments discussed in Chapter 3.

1.1 The search for neutrinos

We know from the invisible decay width of the Z boson that there are at most three flavours of light (mass $< \frac{1}{2}m_Z$), active (weakly interacting) neutrinos[7]. Confirming their existence through experiment was not easy. The electric neutrality of neutrinos makes them inherently difficult to observe, hence the early pessimism over the chances of detecting them. Experimentalists nonetheless found ways around this with initial detections using high flux neutrino sources.

1.1.1 Cowan–Reines neutrino experiment

The Nobel prize winning result confirming the existence of (anti) neutrinos was published in 1956[8]. Frederick Reines, Clyde L. Cowan, Jr. et al conducted the experiment at the Savannah River nuclear power plant following a preliminary experiment with a tentative positive result at Hanford[9].

Initial plans to use a nuclear bomb as a neutrino source never came to fruition. Their spirit lived on in the use of a nuclear reactor, providing $\bar{\nu}_e$ from nuclear fission in a much more controlled manner. These were detected with the inverse β decay reaction

$$\bar{\nu}_e + p \rightarrow e^+ + n, \quad (1.1)$$

on a ^{108}Cd infused scintillator. It is characterised by a prompt γ pair from e^+ annihilation on an e^- followed by a delayed γ once the neutron thermalises and is captured.

A reactor power dependent signal was witnessed at $20\times$ expected background, confirming the existence of (electron) neutrinos.

1.1.2 Muon neutrinos at the AGS

Until 1956 only one type of neutrino had been detected, producing e^\pm in weak interactions. An experiment at the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory (BNL)[10] led to the identification of another distinct flavour which instead produced μ^\pm .

The $(\bar{\nu})_\mu$ produced in the decays of secondary particles from a proton beam impinging on a target were picked up by a well shielded detector downstream. Here spark chamber modules were able to observe μ^\pm produced by neutrino interactions.

A total of 51 events were found where the identity of the lepton (μ^\pm track rather than e^\pm shower) was clear. This indicated the existence of at least two

neutrino types; $\bar{\nu}_e$ coupling with e^\pm and $\bar{\nu}_\mu$ coupling with μ^\pm .

1.1.3 DONUT and tau neutrinos

Although the τ lepton was discovered in 1975[11] observation of a $\bar{\nu}_\tau$ was not published until 2000, after the DONUT experiment[12] at Fermilab observed their interactions with emulsion targets.

Fermilab's Tevatron created a beam of high energy (800 GeV) protons which impinged on a tungsten beam dump, producing $\bar{\nu}_\tau$ neutrinos primarily through the strange D meson decay

$$D_S^- \rightarrow \tau^- + \bar{\nu}_\tau. \quad (1.2)$$

Just as earlier experiments required distinctions to be made between the e^\pm and μ^\pm produced by their respective neutrinos, DONUT required τ^\pm to be reliably identified. At their beam energies τ^\pm decayed within around 2 mm of production, mostly into a state with a single visible particle. $\bar{\nu}_\tau$ interactions were thus identified from a vertex at the τ^\pm production point followed by a kink from its decay.

Post selection, DONUT found evidence for 4 ν_τ interaction candidate events. Although small, this number was far in excess of the null hypothesis background of 0.34 candidate events and consistent with Standard Model predictions.

1.2 Solar and atmospheric anomalies

Initial discoveries in neutrino physics required man made, high flux sources of neutrinos close to a detector. The sun is certainly a huge source of neutrinos but the observed flux is limited by its great distance from our planet and a lack of funds to place a large neutrino detector any closer to it.

Another significant neutrino source is the upper atmosphere where interactions of high energy particles produce secondaries decaying into neutrinos.

Both solar and atmospheric neutrinos were the source of controversy in the latter half of the 20th century which was eventually resolved by the theory of neutrino oscillations.

1.2.1 Chlorine as a detector

The concept of using radiochemical means to detect neutrinos was initially suggested by Bruno Pontecorvo[13] who proposed detecting neutrinos from inverse β decay in ^{37}Cl ,



requiring neutrinos of 814 keV or more.

Raymond Davis, Jr., with his background in radiochemistry, took on the task of detecting neutrinos with ${}^{37}\text{Cl}$. This started in 1955 with vats placed near nuclear reactors at Brookhaven and Savannah river (the site of the Cowan–Reines result from subsection 1.1.1). Since these reactors produced $\bar{\nu}_e$ while the detector was sensitive only to ν_e both efforts produced negative results.

Davis was also interested in using his detector to measure solar neutrinos. A first attempt in 1955 (still before any neutrino had been detected) using a 4000 L detector buried underground to reduce cosmic μ background found nothing; the number of neutrinos above the 0.814 MeV threshold was not great enough. Later efforts at the Baberton limestone mine also failed to detect anything.

1.2.2 The Homestake experiment

Davis contacted John Bahcall in 1962, prompted by the latter’s work calculating beta decay rates in stellar interiors[14]. Bahcall had an interest in solar neutrinos stemming from their utility as a test of theory on nuclear reactions in stars[15]. The two began a lengthy collaborative effort, with Davis’ group running experiments and Bahcall and collaborators providing theoretical calculations.

Bahcall and his collaborators provided a first calculation of solar neutrino fluxes in 1963[14], predicting capture in a chlorine detector at a rate of ~ 5 SNU[†].

A much larger volume of ${}^{37}\text{Cl}$ than had previously been used was clearly needed. After several years of setbacks and attempts to find funding, work began on the huge 400000 L detector in the Homestake mine.

Despite the tiny interaction rate and huge number of ${}^{37}\text{Cl}$ atoms to sift through, Davis was able to identify and count individual ${}^{37}\text{Ar}$ atoms at a rate of a few tens of atoms per week. The tank’s location deep underground in the Homestake gold mine shielded it from cosmic ray muon background which could have also produced ${}^{37}\text{Ar}$.

Initial observations detected substantially fewer neutrino interactions than expected, and the gradual accumulation of data and refinement of theory made the disagreement more stark. By the late 1970s sufficient data had been collected to report an experimental neutrino interaction rate of (1.8 ± 0.4) SNU, while the

[†]The small number of expected solar neutrino interactions led to the definition of the Solar Neutrino Unit (SNU) for convenience, equivalent to 10^{-36} neutrino interactions per target atom per second.

expectation from solar models was around (7.5 ± 1.5) SNU[17][18]. The experiment kept collecting data in a stable manner to the 1990s[19] with the large statistical uncertainties eventually becoming comparable to systematics.

1.2.3 Further solar neutrino experiments

Radiochemical experiments with ^{71}Ga rather than ^{37}Cl have the advantage of a lower energy threshold[20] giving an large boost to the number of observable neutrinos (see Figure 1.1 for the solar ν_e energy spectrum). Comparing ^{71}Ga and ^{37}Cl data can reveal at which energies the ν_e deficit appears.

Experiments such as SAGE[21][22] and GALLEX[23] (later GNO[24]) examined solar neutrinos from the late 1980s onwards utilising ^{71}Ga targets. These experiments found about half the neutrinos predicted by standard solar models. Though higher than the number observed by Homestake it was still a substantial deficit.

A fortuitous contribution to neutrino physics came from the large Čerenkov detectors initially built to search for proton decay. These detectors look for a cone of




Figure 1.1: Solar neutrino energy spectrum using solar model BS05(OP). pp chain processes are illustrated by solid black lines whilst CNO processes are illustrated by dashed blue lines. From Bahcall and Serenelli[16].

Čerenkov light produced by charged particles moving faster than the local speed of light, such as those produced by neutrino interactions. Kamiokande and its successor Super-Kamiokande used neutrino–electron elastic scattering to provide extremely useful data on solar neutrinos[25][26]. With higher energy thresholds they were only able to observe the most energetic ^8Be neutrinos but their access to kinematic and temporal information was unprecedented in solar neutrino experiments and allowed confirmation that ‘solar neutrinos’ were indeed coming from the direction of the sun (Figure 1.2).

The Sudbury Neutrino Observatory (SNO) provided the ultimate solution to the solar neutrino problem. The experiment[27] used a D_2O Čerenkov detector to distinguish between charged current, neutral current and elastic interactions, providing sensitivity to rates of ν_e and ν_l (any flavour) interactions separately. As shown in Figure 1.3 they found rates of charged current (sensitive only to ν_e) and elastic (sensitive primarily to ν_e) interactions consistent with previous experiments, while neutral current (sensitive to any ν_l) interactions occurred at the rate predicted by solar models[28].

All of these results taken together implied neutrino oscillation as a solution to the anomaly.

1.2.4 The atmospheric anomaly

Atmospheric neutrinos are produced as a result of interactions between high energy cosmic rays and nuclei in the upper atmosphere. Decays of the resulting secondary particles produce around twice as many $\bar{\nu}_\mu$ as $\bar{\nu}_e$ at low energies or large distances.

The first observations of atmospheric neutrinos came from scintillator arrays in gold mines in Karnataka state, India[29] and South Africa[30]. Both measured muons from $\bar{\nu}_\mu$ interactions in surrounding rock and removed cosmic ray muon background by locating the detectors deep in the mines (under 7500 mwe[†] and 8800 mwe rock respectively) and limiting the angular acceptance to muons travelling in the horizontal direction. Further measurements made at the two locations[31][32], at the University of Utah[33] and at the Baksan detector in Russia[34] also observed atmospheric neutrino induced muons.

Although some of these early experiments measured slightly lower than expected $\bar{\nu}_\mu$ fluxes the atmospheric anomaly became a major issue when it became possible to measure the ratio of $\bar{\nu}_\mu$ to $\bar{\nu}_e$ flux. Due to many uncertainties cancelling

[†]Meter Water Equivalent (mwe) is a unit expressing in a consistent way the fraction of cosmic rays reaching an underground facility. A laboratory under 1000 mwe of rock will see the same flux of cosmic rays as a laboratory under 1000 m of water.




Figure 1.2: Kamiokande solar neutrino fluxes as a function of angle to the sun showing a deficit in data relative to prediction. Solid line shows prediction assuming a standard solar model and no oscillations, dashed lines shows best fit to data. From [26].




Figure 1.3: SNO ^8B neutrino fluxes as measured through charged current, neutral current and elastic scattering processes. Diagonal dashed lines show standard solar model prediction for total flux (as measured by neutral current channel). Results are consistent with standard solar model predictions if ν_e are allowed to oscillate into ν_μ or ν_τ . From [28].

out a deviation in this ratio is more suggestive than a simple deficit in observed $\overline{\nu}_\mu$.

Alongside their solar measurements Kamiokande examined this ratio[35] as did the other Čerenkov nucleon decay experiment IMB[36]. Both found that the ratio of $(\nu_\mu + \bar{\nu}_\mu)$ to $(\nu_e + \bar{\nu}_e)$ was about half the expectation.

Later experiments Soudan-2[37] (a tracking calorimeter, also co-opted from the search for nucleon decay) and MACRO[38] (a search for magnetic monopoles) confirmed this deviation. On the other hand, Fréjus [39] and NUSEX[40] both found the ratio to be in line with expectation, further confounding the issue.

The breakthrough came when Super-Kamiokande[42], Kamiokande's neutrino oscillation dedicated successor, reported a strongly zenith angle dependent deficit in atmospheric neutrinos[41]. As shown in Figure 1.4 the huge water Čerenkov detector observed that $\overline{\nu}_e$ and (high energy) downwards going $\overline{\nu}_\mu$ appeared at roughly the rate expected but upwards going $\overline{\nu}_\mu$ appeared at half this rate. In addition, at lower energies the number of $\overline{\nu}_\mu$ was lower than expected over all angles.

The K2K long baseline accelerator experiment backed up Super-Kamiokande's results with indications[43], and eventual 4.3σ confirmation[44], of the disappearance of ν_μ from an accelerator source at 250 km with neutrino energies ~ 1.5 GeV. Accelerators provide ν_μ in analogous interactions to those in the upper atmosphere but in a more controlled environment.

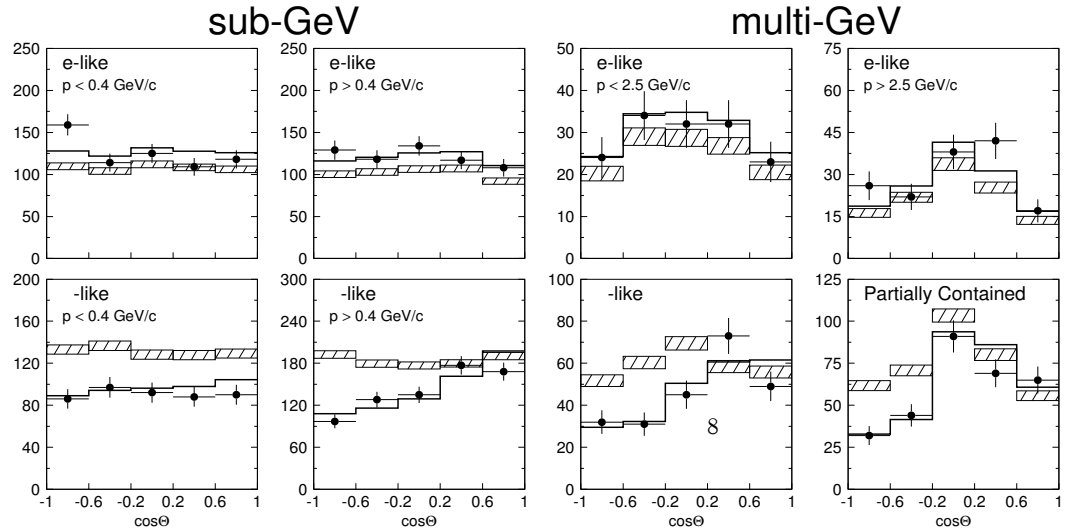


Figure 1.4: Super-Kamiokande results for atmospheric neutrino observations as a function of zenith angle. Hatched regions show predictions with no oscillations. A deficit in low energy ν_μ and a zenith angle dependent deficit in higher energy ν_μ is visible. From [41].

These results strongly implied oscillations as a solution. The distance travelled by atmospheric neutrinos coming from above was much less than an oscillation length whilst those coming from below saw many oscillation lengths — since the exact number was beyond the resolution of detectors they were essentially washed out into an average of half $\bar{\nu}_\mu$ and half some other flavour.

Because of these anomalies and the experiments resolving them in them we live in a world where neutrino oscillations are accepted as fact. The physics of neutrino oscillations are discussed in Chapter 2. In Chapter 3 we come back to the realm of experimental neutrino physics and discuss contemporary neutrino experiments.

Chapter 2

Theoretical Background

Pontecorvo's plans for a ^{37}Cl based neutrino experiment had been halted by his defection to the USSR, which lacked suitable sites, but he remained interested[45]. Reports reached him of a positive result from Ray Davis' reactor experiment with a ^{37}Cl detector (subsection 1.2.1) implying that the reactor's $\bar{\nu}_e$ had been detected as ν_e . These reports turned out to be false but they provoked an interest in neutrino oscillations resulting in hypotheses[46] which by the time of the late 1960s had been refined. Meanwhile Maki, Nakagawa and Sakata extended the idea of neutrino mixing in 1962[47].

Neutrino mixing and oscillations are delved into in this chapter. We also discuss the specifics of neutrino interactions on nuclei, which are particularly relevant for the cross section analysis covered by this thesis.

The theory behind neutrino physics has been covered in for example [48][49][50]. Neutrinos are unique among standard model fermions in that they can only interact via the weak interaction. As its name suggests this is relatively weak, hence the extreme unwillingness of neutrinos to interact with other matter.

2.1 Neutrino masses

In the standard model neutrinos are massless. However, for reasons which will soon become clear, they are required to be massive in order for oscillations to occur. The additional mass terms can represent a slight extension to the standard model or be part of a wider theory.

2.1.1 Mass terms

Mass terms in the Standard Model Lagrangian have the general form

$$\mathcal{L}_{m_\nu} = m\bar{\nu}_i\nu_i. \quad (2.1)$$

A massive neutrino has two components, $\nu = \nu_R + \nu_L$, where ν_R and ν_L are right and left chiral projections respectively ($P_R\nu = \nu_R$, $P_L\nu = \nu_L$). If the two are independent the neutrino is a Dirac particle with mass terms

$$\mathcal{L}_{m_\nu} = m_D\bar{\nu}_{iR}\nu_{iL} + \text{H. c.} \quad (2.2)$$

On the other hand since the neutrino has no charge the two components can be conjugates, $\nu_R = (\nu_L)^c$, where the superscript denotes particle–anti particle conjugation. The general Majorana mass terms are

$$\mathcal{L}_{m_\nu} = m\bar{\nu}_L^c M_M \nu_L + \text{H. c.}, \quad (2.3)$$

where M_M denotes the symmetric matrix of Majorana masses linking neutrino fields $\nu = (\nu_1 \dots \nu_n)$. These are not invariant under U(1) and so are only allowed for particles with no (conserved) charge.

The most general neutrino mass matrix includes both Dirac and Majorana mass terms as

$$\mathcal{L}_{m_\nu} = \frac{1}{2} \begin{pmatrix} \nu_L^T & \nu^c_L{}^T \end{pmatrix} C \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix} \begin{pmatrix} \nu_L \\ \nu^c_L \end{pmatrix} + \text{H. c.}, \quad (2.4)$$

where M_D , M_L and M_R denote Dirac, left handed Majorana and right handed Majorana mass matrices respectively C the particle–anti particle conjugation matrix $i\gamma_2\gamma_0$ and ν the n dimensional vector of neutrino fields.

The massive neutrino fields can be found by diagonalising this mass matrix. An interesting consequence of this is the possibility of a ‘see-saw mechanism’, where the smallness of observed neutrino masses is explained by their suppression by a large right handed mass terms.

Majorana mass terms also break lepton number, allowing for processes specific to Majorana neutrinos such as neutrinoless double beta decay.

2.1.2 Mass terms and hierarchy

The three known generations of neutrino provide three mass states usually denoted ν_1 , ν_2 and ν_3 . Current experiments are sensitive to the differences between their masses but not the absolute ordering.

We know that there exists one smaller mass squared difference, usually identified with solar oscillations, and one larger difference, usually identified with atmospheric oscillations. By convention the ν_1 - ν_2 mass difference is taken to be the smaller of the two while the ν_2 - ν_3 mass difference (approximately the same as the ν_1 - ν_3 mass difference) is taken to be the larger. As illustrated in Figure 2.1 this is compatible with two hierarchies, with ν_3 as either the most or least massive state.

2.2 Neutrino oscillation physics

Neutrino oscillations are a consequence of having two sets of neutrino eigenstates, mass and flavour, which do not coincide with each other. Neutrinos are created

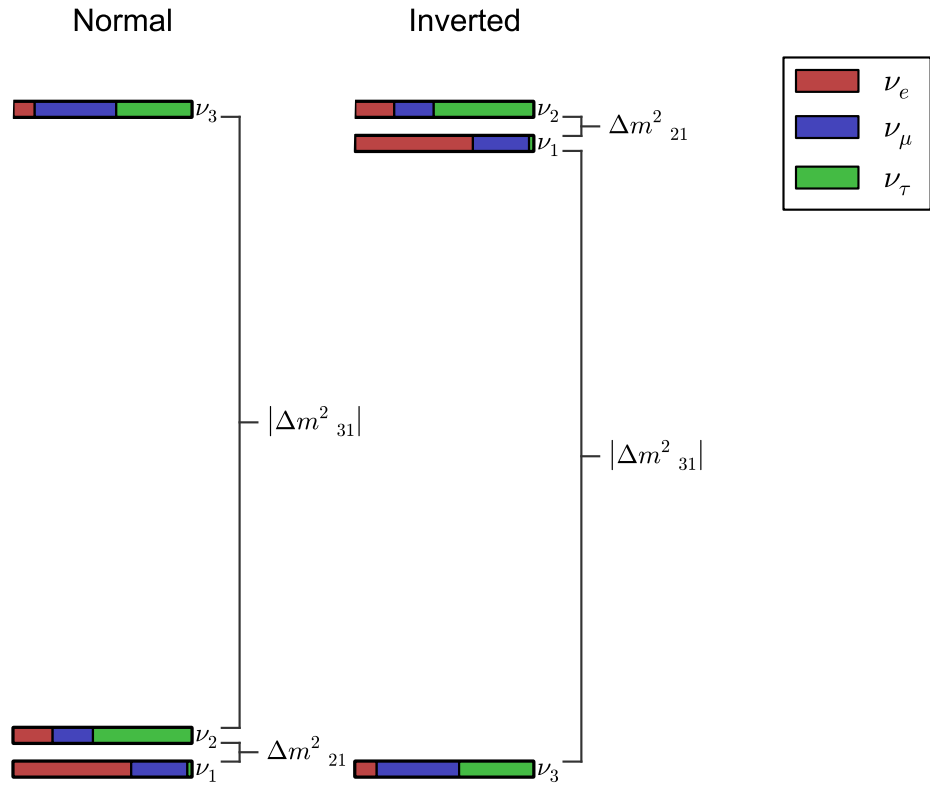


Figure 2.1: Illustration of normal and inverted mass hierarchies, either of which is compatible with current data. Assuming oscillation parameters from [51].

and detected as flavour eigenstates but the mass states in which they propagate are some mixture of these. Different masses ensure the states propagate at different rates allowing them to interfere. The ultimate consequence of this is that a neutrino created as one flavour state has a chance of being detected as another.

2.2.1 Two neutrino case

Oscillations are easiest to understand in the case of two flavours oscillating in a vacuum. The principles outlined here can be generalised to oscillation between three or more flavours.

We consider flavour eigenstates, ν_e and ν_μ , and mass eigenstates, ν_1 and ν_2 . These states are not identical but are linked by a unitary mixing matrix as

$$\begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} \\ U_{\mu1} & U_{\mu2} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \quad (2.5)$$

From a practical standpoint it is sufficient to model the neutrino states as plane waves (though a wave packet treatment is more complete[50]). As neutrinos propagate through vacuum in mass eigenstates the evolution of a neutrino created in some flavour eigenstate ν_a is

$$|\nu(\tilde{x})\rangle = \sum_i e^{-i\tilde{p}_i \cdot \tilde{x}} U_{ai} |\nu_i\rangle, \quad (2.6)$$

where $\tilde{x} \equiv (t, \vec{x})$ is the four-position of the neutrino relative to its starting point and $\tilde{p}_i \equiv (E_i, \vec{p}_i)$ the four-momentum of mass state i . The probability of its being detected in flavour eigenstate ν_b at some later time t is

$$\begin{aligned} P(\nu_a \rightarrow \nu_b) &= |\langle \nu_b | \nu(\tilde{x}) \rangle|^2 \\ &= \sum_i \sum_j U_{bi}^* U_{ai} U_{bj} U_{aj}^* e^{-i(\tilde{p}_i - \tilde{p}_j) \cdot \tilde{x}}, \end{aligned} \quad (2.7)$$

where ν_a, ν_b denote flavour eigenstates and $\nu_i, \nu_j, \tilde{p}_i, \tilde{p}_j$ denote mass eigenstates and their four-momenta.

Neutrinos have sufficiently low mass to be relativistic. If we assume the two mass states to have the same momentum, $p = |\vec{p}_1| = |\vec{p}_2|$, we have

$$\begin{aligned}
(\tilde{p}_2 - \tilde{p}_1) \cdot \tilde{x} &= (E_2 - E_1)t \\
&= \left(p\sqrt{1 + \frac{m_2^2}{p^2}} - p\sqrt{1 + \frac{m_1^2}{p^2}} \right) t \\
&\approx \frac{(m_2^2 - m_1^2)}{2p}t,
\end{aligned} \tag{2.8}$$

which, taking $L \approx t$ as the neutrino's distance from its source, $E \approx p$ as the average energy of the mass states and $\Delta m^2 \equiv m_2^2 - m_1^2$, becomes

$$\Delta m^2 \frac{L}{2E}. \tag{2.9}$$

We have

$$\begin{aligned}
P(\nu_a \rightarrow \nu_b) &= \sum_i \sum_j U_{bi}^* U_{ai} U_{bj} U_{aj}^* e^{-i\Delta m^2 \frac{L}{2E}} \\
&= |U_{a1} U_{b1}|^2 \\
&\quad + |U_{a2} U_{b2}|^2 \\
&\quad + 2U_{a1} U_{b1}^* U_{a2}^* U_{b2} e^{-i\Delta m^2 \frac{L}{2E}} \\
&\quad + 2U_{a2} U_{b2}^* U_{a1}^* U_{b1} e^{-(-i\Delta m^2 \frac{L}{2E})} \\
&= |U_{a1} U_{b1}|^2 + |U_{a2} U_{b2}|^2 + 2U_{a1} U_{b1}^* U_{a2}^* U_{b2} \cos\left(\Delta m^2 \frac{L}{2E}\right),
\end{aligned} \tag{2.10}$$

and the probability of observing ν_b clearly oscillates with distance.

Using the parametrisation of Equation 2.5 we obtain

$$\begin{aligned}
P(\nu_a \rightarrow \nu_b) &= 2 \cos^2 \theta \sin^2 \theta - 2 \cos^2 \theta \sin^2 \theta \cos\left(\Delta m^2 \frac{L}{2E}\right) \\
&= \sin 2\theta \sin^2\left(\Delta m^2 \frac{L}{4E}\right) \\
&= \sin 2\theta \sin^2\left(1.27 \Delta m^2 / \text{eV}^2 \frac{L/\text{km}}{E/\text{GeV}}\right),
\end{aligned} \tag{2.11}$$

providing the oscillation probability in an intuitive form.

2.2.2 Solar neutrinos and the MSW effect

Weak interactions of neutrinos with matter cause oscillations to differ from their expectation in vacuum. The most significant manifestations of this, for its impact on solar neutrino observations, is the resonant enhancement of oscillation probability known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect[52].

Various interactions affect neutrino propagation in matter. The least negligible (proportional to G_F rather than one of its higher powers) is forward scattering. This is a coherent process somewhat analogous to refraction of light[53] and while its direct observation is implausible the mean potentials it effects can significantly alter neutrino oscillations, allowing their probability to exceed the upper limit of $\sin^2(2\theta)$ in vacuum oscillations.

All active neutrino flavours experience neutral current interactions with leptons and hadrons in matter. Charged current interactions are also possible for neutrinos of the appropriate flavour. Standard electrically neutral, electron dense matter allows for charged current interactions with ν_e . The potentials experienced by different flavours are

$$\begin{aligned} V_e &= \sqrt{2}G_F \left(N_e - \frac{1}{2}N_n \right), & V_{\bar{e}} &= -\sqrt{2}G_F \left(N_e - \frac{1}{2}N_n \right), \\ V_\mu &= \sqrt{2}G_F \left(\frac{1}{2}N_n \right), & V_{\bar{\mu}} &= -\sqrt{2}G_F \left(-\frac{1}{2}N_n \right), \\ V_\tau &= \sqrt{2}G_F \left(\frac{1}{2}N_n \right), & V_{\bar{\tau}} &= -\sqrt{2}G_F \left(-\frac{1}{2}N_n \right). \end{aligned} \tag{2.12}$$

Once again we consider the two neutrino case (this can be generalised to further flavours). Propagation of mass eigenstates is given by

$$i \frac{d}{dt} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \begin{pmatrix} p + \frac{m_1^2}{2E} & 0 \\ 0 & p + \frac{m_2^2}{2E} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}, \tag{2.13}$$

and using the mixing matrix described in subsection 2.2.1 we find an evolution of

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \begin{pmatrix} p + \frac{m_1+m_2}{4E} - \frac{\Delta m^2}{4E} \cos 2\theta & \frac{\Delta m^2}{4E} \sin 2\theta \\ \frac{\Delta m^2}{4E} \sin 2\theta & p + \frac{m_1+m_2}{4E} + \frac{\Delta m^2}{4E} \cos 2\theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}, \tag{2.14}$$

in the flavour basis.

Matter effects are incorporated by adding $\text{diag}(V_e, V_\mu)$ to the flavour Hamiltonian. The part of the Hamiltonian proportional to the identity does not affect oscillations and so can be subtracted out, giving

$$i \frac{d}{dt} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \begin{pmatrix} -\frac{\Delta m^2}{4E} \cos 2\theta + \sqrt{2}G_F & \frac{\Delta m^2}{4E} \sin 2\theta \\ \frac{\Delta m^2}{4E} \sin 2\theta & \frac{\Delta m^2}{4E} \cos 2\theta \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}, \quad (2.15)$$

Propagation eigenstates are produced by diagonalising this effective Hamiltonian with

$$\begin{pmatrix} \nu_A \\ \nu_B \end{pmatrix} = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix}. \quad (2.16)$$

The mixing angle ϕ for these propagation states with flavour states is defined as

$$\tan 2\phi = \frac{\frac{\Delta m^2}{2E} \sin 2\theta}{\frac{\Delta m^2}{2E} \cos 2\theta - \sqrt{2}G_F N_e}. \quad (2.17)$$

This shows that the actual mixing depends on the electron number density of the medium being propagated through. The following cases are useful to consider, visible in Figure 2.2:

- $\sqrt{2}G_F N_e \gg \frac{\Delta m^2}{2E} \cos 2\theta$: In this case, where density is extremely high, ϕ becomes close to 90° and the flavour eigenstates roughly align with the propagation eigenstates.
- $\sqrt{2}G_F N_e = \frac{\Delta m^2}{2E} \cos 2\theta$: In this case mixing is maximum. This is known as the MSW resonance condition.
- $\sqrt{2}G_F N_e = 0$: In this case the vacuum mixing angles are clearly recovered.

During the course of their journey from the sun's core to its surface neutrinos in the right energy range will pass through all of these regimes. At production the neutrinos are in a high density environment where propagation eigenstates ν_A, ν_B can almost completely overlap with the flavour eigenstates ν_e, ν_μ . These neutrinos propagate through the resonance until they reach a vacuum, where ν_A, ν_B now almost completely overlap with the mass eigenstates ν_1, ν_2 .

As a result ν_e produced in the core start out almost entirely aligned with the ν_B state. Provided that conditions are right as the neutrino propagates by the time it reaches the vacuum of space the neutrino will still be in the ν_B state which




Figure 2.2: Mixing angles between flavour eigenstates and matter propagation states as a function of electron number density, showing initially minimal mixing followed by a resonance and final mixing aligning roughly with vacuum mixing angles. Using oscillation parameters from [51].

now corresponds to ν_2 , with a ν_e component of $-\sin \theta$ and a ν_μ component of $\cos \theta$. Therefore, the transition probability is

$$P(\nu_e \rightarrow \nu_\mu) = |\langle \nu_\mu | \nu(t) \rangle|^2 = \cos^2 \theta. \quad (2.18)$$

The necessary condition for this to hold is *adiabaticity*. Stellar density must change slowly enough for the neutrino state to adjust. A small value of θ can break this condition. Our current knowledge of the solar mixing angle and mass splitting suggests that higher energy solar neutrinos experience resonance, with around two thirds of ν_e converted to ν_μ , explaining the deficit detailed in Section 1.2.

2.2.3 The PMNS mixing matrix

So far we have considered two neutrino oscillations. The three neutrino flavours we can observe are linked through the unitary mixing matrix

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (2.19)$$

known as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. This matrix can also be parametrised as a rotation through three mixing angles; θ_{12} , θ_{23} and θ_{13} . There are also up to three CP violating phases of which δ_{CP} is relevant to oscillation experiments. Using the shorthands s_{ij} and c_{ij} for the respective sines and cosines of the θ_{ij} mixing angles,

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\text{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{\text{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (2.20)$$

where CP violating phase δ_{CP} is placed with the odd θ_{13} terms by convention.

In the case of three Majorana neutrinos the remaining two CP violating phases appear as

$$U_{\text{General, Majorana}} = U \times \text{diag}(1, e^{i\alpha}, e^{i\beta}), \quad (2.21)$$

where U is the PMNS matrix as defined in Equation 2.20 and α and β are two new phases, neither of which is observable through oscillation experiments.

Further generalisations to higher dimensions, including heavy or sterile neutrinos, are viable but not discussed here.

2.2.4 Three flavour oscillation

The case of three oscillating species of neutrino is analogous to the two neutrino case discussed in subsection 2.2.1. A greater number of terms are interfering and the resulting equations are more complex.

The transition probability in this case is

$$\begin{aligned}
P\left(\nu_a \rightarrow \nu_b; \frac{L}{E}\right) &= \sum_i \sum_j U_{ai}^* U_{bi} U_{aj} U_{bj}^* e^{-i \frac{\Delta m^2_{ji} L}{2E}} \\
&= \sum_i |U_{ai}|^2 |U_{bi}|^2 \\
&\quad + 2 \sum_{i>j} \operatorname{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \operatorname{Re}\left(e^{-i \frac{\Delta m^2_{ji} L}{2E}}\right) \\
&\quad - 2 \sum_{i>j} \operatorname{Im}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \operatorname{Im}\left(e^{-i \frac{\Delta m^2_{ji} L}{2E}}\right),
\end{aligned} \tag{2.22}$$

assuming relativistic neutrinos with small mass. Notably the Majorana phases in Equation 2.21 disappear in the quartic product $U_{ai}^* U_{bi} U_{aj} U_{bj}^*$. This directly shows that Majorana phases are invisible to neutrino oscillation experiments.

Since U is unitary

$$\begin{aligned}
\sum_i U_{ai}^* U_{bi} &= \delta_{ab} \\
\left| \sum_i U_{ai}^* U_{bi} \right|^2 &= \sum_i \sum_j U_{ai}^* U_{bi} U_{aj} U_{bj}^* \\
&= \sum_i |U_{ai}|^2 |U_{bi}|^2 + 2 \sum_{i>j} \operatorname{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \\
&= \delta_{ab},
\end{aligned}$$

which combines with Equation 2.22 to give

$$\begin{aligned}
P\left(\nu_a \rightarrow \nu_b; \frac{L}{E}\right) &= \delta_{ab} \\
&\quad + 2 \sum_{i>j} \operatorname{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \left(\cos\left(\frac{\Delta m^2_{ji} L}{2E}\right) - 1 \right) \\
&\quad + 2 \sum_{i>j} \operatorname{Im}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \left(\sin\left(\frac{\Delta m^2_{ji} L}{2E}\right) \right),
\end{aligned} \tag{2.23}$$

which can also be written as

$$\begin{aligned}
\text{P}\left(\nu_a \rightarrow \nu_b; \frac{L}{E}\right) &= \delta_{ab} \\
&\quad -4 \sum_{i>j} \text{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \sin^2\left(\frac{\Delta m_{ji}^2 L}{4E}\right) \\
&\quad +2 \sum_{i>j} \text{Im}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \sin\left(\frac{\Delta m_{ji}^2 L}{2E}\right).
\end{aligned} \tag{2.24}$$

The case with antineutrinos is analogous but with an extra minus sign acquired by the terms derived from imaginary parts as

$$\begin{aligned}
\text{P}\left(\bar{\nu}_a \rightarrow \bar{\nu}_b; \frac{L}{E}\right) &= \delta_{ab} \\
&\quad +2 \sum_{i>j} \text{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \left(\cos\left(\frac{\Delta m_{ji}^2 L}{2E}\right) - 1\right) \\
&\quad -2 \sum_{i>j} \text{Im}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \left(\sin\left(\frac{\Delta m_{ji}^2 L}{2E}\right)\right).
\end{aligned} \tag{2.25}$$

In searches for neutrino survival (i.e. $\nu_a \rightarrow \nu_a$) the $U_{ai}^* U_{bi} U_{aj} U_{bj}^*$ are clearly real, so

$$\text{P}\left(\bar{\nu}_a \rightarrow \bar{\nu}_a; \frac{L}{E}\right) = 1 + 2 \sum_{i>j} \text{Re}(U_{ai}^* U_{bi} U_{aj} U_{bj}^*) \left(\cos\left(\frac{\Delta m_{ji}^2 L}{2E}\right) - 1\right), \tag{2.26}$$

with a symmetry between particle and antiparticle survival rates as required by CPT invariance.

In the limit $\Delta m^2_{ij} \frac{L}{2E} \approx 0^\dagger$ the corresponding oscillation terms disappear

$$\begin{aligned}
& 2\text{Re}(U_{ai}^* U_{bi} U_{aj}^* U_{bj}) \left(\cos\left(\frac{\Delta m^2_{ij} L}{2E}\right) - 1 \right) \\
& \pm 2\text{Im}(U_{ai}^* U_{bi} U_{aj}^* U_{bj}) \left(\sin\left(\frac{\Delta m^2_{ij} L}{2E}\right) \right) \\
& \approx 0,
\end{aligned} \tag{2.27}$$

and only constant terms are relevant for oscillations. On the other hand if the detector covers a large number of oscillation lengths[‡] the terms average out to

$$\begin{aligned}
& 2\text{Re}(U_{ai}^* U_{bi} U_{aj}^* U_{bj}) \left\langle \cos\left(\frac{\Delta m^2_{ij} L}{2E}\right) - 1 \right\rangle \\
& \pm 2\text{Im}(U_{ai}^* U_{bi} U_{aj}^* U_{bj}) \left\langle \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right) \right\rangle \\
& \approx -2\text{Re}(U_{ai}^* U_{bi} U_{aj}^* U_{bj}),
\end{aligned} \tag{2.28}$$

contributing a constant enhancement or deficit in the number of observed neutrinos.

2.2.5 Specific applications

Neutrino oscillation experiments in general are not sensitive to the entire PMNS matrix. The specific terms which are relevant depend on the experiment's $\frac{L}{E}$, the neutrino flavours being created and detected, and the presence of matter effects.

The survival probability of $(\bar{\nu})_e$ is perhaps the least complicated to derive, particularly at scales where Δm^2_{21} can be neglected. This is the case for medium baseline reactor experiments such as RENO and Daya Bay. From Equation 2.24, using the PMNS matrix parametrisation of Equation 2.20,

[†]This is trivially the case at zero distance from the neutrino source. It also approximately occurs by design in near detectors for long baseline oscillation experiments such as our own ND280.

[‡]In general this applies where the range of L or E required for an oscillation length is on a scale smaller than a detector's length or energy resolution. In atmospheric experiments neutrinos approaching from very large zenith angles are effectively averaged out since in this region the detector's angular resolution gives spatial a resolution larger than the scale of an oscillation length.

$$\begin{aligned}
\text{P}\left(\bar{\nu}_e \rightarrow \bar{\nu}_e; \frac{L}{E}\right) &\approx 1 - 4 \sum_{i>j; i,j \neq 2,1} \text{Re}(U_{ei}^* U_{ei} U_{ej} U_{ej}^*) \sin^2\left(\frac{\Delta m_{ji}^2 L}{4E}\right) \\
&\approx 1 - 4 |U_{e1}|^2 |U_{e3}|^2 \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \\
&\quad - 4 |U_{e2}|^2 |U_{e3}|^2 \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right) \\
&\approx 1 - \sin^2 2\theta_{13} \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right),
\end{aligned} \tag{2.29}$$

where $\Delta m_{31}^2 \approx \Delta m_{32}^2$ given the small value of Δm_{21}^2 . This shows that $(\bar{\nu}_e \rightarrow \bar{\nu}_e)$ measurements with the relevant baselines (medium baseline reactor experiments) can unambiguously provide the value of θ_{13} .

From long baseline or accelerator experiments $(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu)$ measurements can improve our knowledge on θ_{23} and Δm_{23}^2 , while precise $(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)$ results from accelerators can also shed light on the value of θ_{13} . Again assuming negligible Δm_{21}^2 effects,

$$\begin{aligned}
\text{P}\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu; \frac{L}{E}\right) &\approx -4\text{Re}(U_{\mu 1}^* U_{\mu 1} U_{\mu 3} U_{\mu 3}^*) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \\
&\quad - 4\text{Re}(U_{\mu 2}^* U_{\mu 2} U_{\mu 3} U_{\mu 3}^*) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right),
\end{aligned} \tag{2.30}$$

and

$$\begin{aligned}
\text{P}\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_e; \frac{L}{E}\right) &\approx -4\text{Re}(U_{e1}^* U_{\mu 1} U_{e3} U_{\mu 3}^*) \sin^2\left(\frac{\Delta m_{31}^2 L}{4E}\right) \\
&\quad - 4\text{Re}(U_{e2}^* U_{\mu 2} U_{e3} U_{\mu 3}^*) \sin^2\left(\frac{\Delta m_{32}^2 L}{4E}\right) \\
&\quad \pm 2\text{Im}(U_{e1}^* U_{\mu 1} U_{e3} U_{\mu 3}^*) \sin\left(\frac{\Delta m_{31}^2 L}{2E}\right) \\
&\quad \pm 2\text{Im}(U_{e2}^* U_{\mu 2} U_{e3} U_{\mu 3}^*) \sin\left(\frac{\Delta m_{32}^2 L}{2E}\right).
\end{aligned} \tag{2.31}$$

These terms are more fiddly than the $\bar{\nu}_e$ appearance terms of Equation 2.29, but eventually yield[54][55][56]

$$P\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu; \frac{L}{E}\right) \approx 1 - 4 \sin^2 \theta_{23} \cos^2 \theta_{13} \left(1 - \sin^2 \theta_{23} \cos^2 \theta_{13}\right) \sin^2 \left(\frac{\Delta m_{\text{eff}}^2 L}{4E}\right), \quad (2.32)$$

for effective mass squared difference

$$\Delta m_{\text{eff}}^2 = \Delta m_{32}^2 + \Delta m_{21}^2 \sin^2 \theta_{12} + \Delta m_{21}^2 \cos \delta_{\text{CP}} \sin \theta_{13} \tan \theta_{23} \sin 2\theta_{12}, \quad (2.33)$$

which simplifies to

$$P\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu; \frac{L}{E}\right) \approx 1 - \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{\text{eff}}^2 L}{4E}\right),$$

if the deficit is assumed to be due entirely to $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations.

For $\bar{\nu}_e$ appearance[57] the probability approximates to

$$P\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_e; \frac{L}{E}\right) \approx \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{32}^2 L}{4E}\right),$$

ignoring subleading terms (which include matter effects and δ_{CP} dependence).

Solar neutrino oscillations are dependent on the MSW effect discussed in in subsection 2.2.2.

2.2.6 Mass hierarchy, CP violation and octant

The three remaining unknowns in neutrino physics are the mass hierarchy, the CP violating phase δ_{CP} and the octant of θ_{23} .

The mass hierarchy has been mentioned in subsection 2.1.2. Its identity as normal or inverse is given by the sign of Δm_{31}^2 . We currently only know this parameter's magnitude.

CP violation in the neutrino sector is of huge theoretical interest. CP violation is one of the Sakharov conditions needed for baryogenesis in the early universe. Currently known mechanisms for CP violation are insufficient for baryogenesis in

the early universe on the scale needed to explain our present abundance of matter. Non-zero δ_{CP} could be determined by confirming a difference between, for example, oscillations of $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$.

The remaining unknown, the octant of θ_{23} , emerges since current experiments mainly measure $\sin^2 2\theta_{23}$. Through measurements of $\sin^2 2\theta_{23}$ alone it is impossible to determine whether θ_{23} is above or below 45° .

These unknowns are not independent of each other. Mass hierarchy and θ_{23} octant depend on the values of δ_{CP} and θ_{13} [58].

2.3 Neutrino interactions

An understanding of neutrino interactions is obviously important for neutrino physics. Since direct measurements of neutrinos' kinematic properties are impossible experimentalists have to rely on the observable particles produced by their interactions.

At present the most precise data available is at higher energies, above a few GeV, where both quasi-elastic cross sections and inclusive cross sections (including the deep inelastic scattering effects which dominate) are measured. Below this range, in the region where T2K operates, data is relatively sparse.

At very low energies (up to about 10 MeV) interactions such as inverse beta decay are important. Such interactions have historically been significant (see for example subsection 1.2.2) and are still useful in searches for low energy neutrinos.

In the intermediate energy range (around 100 MeV to around 20 GeV[†]) several new processes become relevant. The lower portion of this energy range (to about 1 GeV) is dominated by elastic and quasi-elastic scattering. At intermediate energies (a few GeV) the most significant processes involve resonant production of mesons. For higher energies deep inelastic scattering dominates. The dominant processes for charged current neutrino and anti neutrino interactions in this range are shown in Figure 2.3 and discussed further in this section.

2.3.1 Elastic scattering

Elastic interactions, as shown in Figure 2.4, leave the identities of the incoming and outgoing particles unchanged and redistribute their energy and momentum. The processes which are feasible to observe involve scattering on electrons,

[†]This range is relevant to contemporary accelerator experiments including T2K and thus our own analysis.




Figure 2.3: Total ν_μ (Figure 2.3a) and $\bar{\nu}_\mu$ (Figure 2.3b) charged current cross sections per nucleon divided by energy. Quasi-elastic events can be seen to dominate at low energies and deep inelastic scattering at high energies, with resonance effects in between. Experimental data is also shown. From [59]




Figure 2.4: Elastic scattering of (anti) neutrinos on electrons and nucleons, with $\bar{\nu}_l$ denoting any neutrino flavour and N denoting nucleons.

$$\begin{aligned}
\nu + e^- &\rightarrow \nu + e^- \\
\bar{\nu} + e^- &\rightarrow \bar{\nu} + e^-,
\end{aligned}
\tag{2.34}$$

and nucleons,

$$\begin{aligned}
\nu + N &\rightarrow \nu + N \\
\bar{\nu} + N &\rightarrow \bar{\nu} + N.
\end{aligned}
\tag{2.35}$$

Most of these interactions are neutral current (NC), although ν_e and $\bar{\nu}_e$ scattering on electrons as in Equation 2.34 do have charged current channels.

2.3.2 Charged current quasi-elastic interactions

Quasi-elastic (QE) interactions leave the incoming and outgoing particles mostly unchanged apart from exchanging particles for their weak isospin partners, with $\nu_l \leftrightarrow l$ and $p \leftrightarrow n$, as shown in Figure 2.5.




Figure 2.5: Quasi-elastic scattering of (anti) neutrinos on nucleons, with $\bar{\nu}_l$ denoting any neutrino flavour and l^\pm its corresponding lepton.

As with elastic scattering, charged current quasi-elastic interactions (CCQE) can take place on electrons

$$\nu_l + e^- \rightarrow \nu_e + l^-, \quad (2.36)$$

(inverse muon decay in the case of ν_μ) and nucleons

$$\begin{aligned} \nu_l + p &\rightarrow l^+ + n \\ \bar{\nu}_l + n &\rightarrow l^- + p. \end{aligned} \quad (2.37)$$

The CCQE interactions on nucleons of Equation 2.42 are of particular utility in neutrino physics. Assuming no interference from nuclear effects, they allow neutrino kinematics to be accurately reconstructed from the outgoing lepton. They occur mainly where neutrino energy is insufficient to probe nuclear structure.

The formalism of Llewellyn Smith[60] presents CCQE cross sections as

$$\begin{aligned} \frac{d\sigma}{d(Q^2)} &= \frac{M^2 G_F \cos^2 \theta_C}{8\pi E_\nu^2} \left(A(Q^2; M, m, \xi) \right. \\ &\quad + B(Q^2; M, m, \xi) \frac{s-u}{M^2} \\ &\quad \left. + C(Q^2; M, \xi) \left(\frac{s-u}{M^2} \right)^2 \right), \end{aligned} \quad (2.38)$$

for nucleon mass M , lepton mass m and anomalous magnetic moment ξ . The $A(Q^2; M, m, \xi)$, $B(Q^2; M, m, \xi)^\dagger$ and $C(Q^2; M, \xi)$ can be further parametrised in terms of form factors F_1 , F_2 , F_{3A} , F_{3V} , F_A and F_P .

The F_{3A} and F_{3V} factors represent G-parity violating terms (second class currents) which are zero in the standard model. F_P can be directly related to F_A given the nucleon and pion mass. The remaining three form factors can be further reduced with electron scattering data, which provides reliable values for the vector terms F_1 and F_2 .

[†]An extra minus sign is acquired by the $B(Q^2; M, m, \xi)$ term in the case of anti neutrinos on nuclei

This leaves one unknown form factor, F_A . Its value at zero is known from β decay experiments. For the remainder of its domain the function is assumed to have a dipole form[61] of

$$F_A(Q^2) = \frac{F_A(0)}{1 + \frac{Q^2}{M_A^2}}. \quad (2.39)$$

The final unknown, M_A , is known as the axial mass and can be determined through neutrino cross section experiments.

The formalism of Smith and Monitz[62] for CCQE interactions in the context of a relativistic Fermi gas model is often used by generators. A similar model for neutral current elastic scattering is described by Ahrens et al[63] where the axial form factor is modified to

$$F_A(Q^2) = \frac{F_A(0)}{1 + \frac{Q^2}{M_A^2}} (1 + \eta), \quad (2.40)$$

where η accounts for (small but non-zero) additional isoscalar contributions to the cross section.

2.3.3 2p2h

2p2h describes processes where the neutrino interaction involves multiple correlated nucleons[64], as illustrated in Figure 2.6. This can lead an extra nucleon in the final state along with the extra contribution to total cross sections. The processes are also referred to as meson exchange current (MEC), since they are mediated by the




Figure 2.6: Feynman diagram for 2p2h processes in neutrino interactions on nuclei. The N are correlated nucleons while the $(\bar{\nu}_l)$ is any (anti) neutrino and l^\pm its corresponding lepton.

exchange of mesons between nucleons, and can also be generalised to n nucleons (nph).

Such processes are of interest because of their impact on physics analyses. Sometimes it is impossible to detect all final state nucleons due to final state interactions (see subsection 2.3.7) whilst others it is impractical to resolve them. Unless all final nucleons can be detected multi-nucleon processes are difficult to distinguish from vanilla quasi-elastic interactions. If CCQE kinematics are assumed in a dataset which includes these processes then a significant bias in reconstructed neutrino energy is consequent.

2.3.4 Resonant pion production

At energies of the order a few GeV neutrinos still lack the energy to break up individual nucleons but are able to raise them to excited states from which they relax by emitting particles (usually single pions). A number of such interactions exist, for example

$$\nu_l + p \rightarrow l^- + \Delta^{++} \rightarrow l^- + \pi^+ + p,$$

provides a channel for charged current resonant pion production on protons illustrated in Figure 2.7a. Both charged and neutral current processes are allowed producing π^+ , π^- and π^0 . At higher energies resonances producing multiple particles become accessible.

Rein and Sehgal provided calculations for single pion production from reso-




Figure 2.7: Examples of resonant pion production with (anti) neutrinos on nucleons, with $\bar{\nu}_l$ denoting any neutrino flavour and l^\pm its corresponding lepton. Figure 2.7a shows charged current single π^+ production whilst Figure 2.7b shows neutral current π^0 production.

nances up to 2 GeV including interference terms[65]. This work is implemented in many currently used neutrino interaction generators.

As with CCQE resonant pion production cross sections can be calculated by describing the nucleus using form factors. In this case the remaining axial form factors provide two parameters, C_5^A and m_A^{RES} , to be fitted.

2.3.5 Coherent processes

Coherent interactions, where the neutrino interacts coherently with the entire nucleus, are allowed. In these cases the nucleus remains intact and recoils unchanged. Momentum transfer is low and pions are produced at highly forward going angles which makes it possible to distinguish such events from resonant production. Rein and Sehgal have produced a model for coherent interactions[66][67] currently in use by generators.

2.3.6 Deep inelastic scattering

In deep inelastic scattering (DIS) interactions the neutrino has sufficient energy to interact directly with individual partons and break up the nucleon, producing a final state with hadronic showers,

$$\begin{aligned}\nu_l + N &\rightarrow \nu_l + X, \\ \bar{\nu}_l + N &\rightarrow \bar{\nu}_l + X,\end{aligned}\tag{2.41}$$

(neutral current), or

$$\begin{aligned}\nu_l + N &\rightarrow l^- + X, \\ \bar{\nu}_l + N &\rightarrow l^+ + X,\end{aligned}\tag{2.42}$$

(charged current). In both of these cases the X denotes some set of final state hadrons.

This process implies a large energy transfer to the hadrons from the neutrino. At higher energies, above ~ 10 GeV, DIS processes dominate neutrino cross sections. Between this region and the sub-1 GeV regime where elastic and quasi-elastic scat-




Figure 2.8: Deep inelastic scattering of (anti) neutrinos on nucleons, with $\bar{\nu}_l$ denoting any neutrino flavour, l^\pm its corresponding lepton, N any nucleon (of which only the relevant quark is labelled) and X any hadronic final state.

tering dominates is an intermediate range of energies where other processes such as resonance become significant.

Deep inelastic scattering can be described with leading order parton distribution functions. The 1998 distributions of Glück, Reya and Vogt (GRV98)[68] are an example in current use by neutrino interaction generators.

Bodek and Yang used modifications[69][70] to allow for valid descriptions of inelastic lepton scattering at high and low energies. New scaling variables and modified PDFs at low Q^2 were needed for this purpose.

2.3.7 Final state interactions

Interactions with nuclei are difficult to model and rely on experimental data. One substantial challenge is the possibility of final state interactions. This describes a broad range of processes involving outgoing hadrons re-interacting with the nucleus and can cause the absorption of outgoing particles[71] or the ejection of additional ones[72].

Such processes can effect an observable final state with substantially different topology to the initial neutrino interaction. For example, if the pion produced in a resonant charged current interaction is reabsorbed the final state will appear to be that of a CCQE interaction. This is also a concern for 2p2h interactions where the outgoing protons may be absorbed. In either case, taking an interaction which produced two or more secondary particles as CCQE creates a bias in reconstructed neutrino energy.

Chapter 3

Contemporary Neutrino Physics

Today's neutrino physics experiments have gone beyond the questions discussed in Chapter 1 of whether or not neutrinos exist, how many they are and whether they oscillate. They now quantify the specifics of neutrino oscillations and interaction.

There are many sources of neutrino available to experimenters providing a range of energies and baselines:

- **Nuclear fission** produces unstable products decaying to $\bar{\nu}_e$ of a few MeV.
- **Radioactive isotopes** also decay, producing ν_e and some $\bar{\nu}_e$ below 3.5 MeV.
- **Astrophysical sources** such as supernovae, active galactic nuclei and gamma ray bursts can produce huge neutrino fluxes at great distances from us.
- **Solar neutrinos** are ν_e produced by fusion reactions in the sun, with some neutrinos up to 10 MeV and very high fluxes below a 1.9 MeV.
- **Atmospheric neutrinos** are produced by the decays resulting from cosmic rays impinging on the upper atmosphere. π^\pm decays dominate until higher energies and produce mostly $\bar{\nu}_\mu$. Muon decays are also significant, producing $\bar{\nu}_\mu$ and $\bar{\nu}_e$. Flux is most significant in the GeV range.
- **Particle accelerators** produce neutrinos in different ways:
 - **Beam dumps** produce neutrinos when high energy protons are abruptly stopped and are mostly $\bar{\nu}_e$ and $\bar{\nu}_\mu$ with energies around 100 GeV.
 - **Decay at rest** produces neutrinos when pions are stopped in a target and then decay. They are mostly $\bar{\nu}_\mu$ with some ν_e and ν_μ but very little $\bar{\nu}_e$ background and have energies of 10s of MeV.

- **Decay in flight** produces neutrinos from pion decay in flight. These are mostly ν_μ or $\bar{\nu}_\mu$ (selection is possible by choosing π^\pm sign) with some wrong-sign and $(\bar{\nu})_e$ background. There are different configurations:
 - * **Wide band** beams do not selected pions for momentum and provide a large neutrino flux spanning a wide range of energies.
 - * **Narrow band** beams selected pions of specific momentum and can be used for precision measurements over a small range of energies.
 - * **Off-axis** configurations direct a wide band beam a small angle away from detectors. The T2K experiment adopts this configuration as described in detail in subsection 4.1.1.

In this chapter we cover contemporary neutrino oscillation experiments. We also examine experiments on neutrino interactions with matter, specifically those like our gas interaction analysis which are dedicated to extracting cross sections. A range of interesting topics such as neutrino mass, neutrinoless double beta decay and astrophysical neutrinos are sadly beyond the scope of this thesis.

3.1 Neutrino oscillation experiments

In subsection 1.2.3 and subsection 1.2.4 we discussed the anomalous results leading to the theory and discovery of neutrino oscillations. We now cover in more detail the subsequent experiments quantifying neutrino oscillations.

3.1.1 Resolution of solar and atmospheric anomalies

The solar neutrino problem is complicated by the fact that matter, especially the dense matter of the sun, affects neutrino oscillations (subsection 2.2.2). Later data, particularly from the KamLAND[73] reactor experiment (which showed a deficit in reactor $\bar{\nu}_e$ as shown in Figure 3.1), ruled out all but a large mixing angle solution with matter resonance effects (MSW) and a mass squared difference between neutrino mass states of $\Delta m^2 \sim 7.6 \times 10^{-5} \text{ eV}^2$. By the time they reach the Earth a large fraction of solar ν_e have oscillated into a mixture of ν_μ and ν_τ .

The atmospheric neutrino problem is more straightforward since matter in the Earth has a much smaller effect on oscillation than that of the sun. As discussed in subsection 1.2.4 oscillation effects are not significant for neutrinos reaching the detector from above, over distances ~ 10 s of km. Neutrinos reaching the detector from below, however, have travelled over 10000 km by which point an average of




Figure 3.1: KamLAND results compared with results from other reactor experiments, showing expectation assuming oscillation parameters provided by global fit of solar neutrino data. As shown, other reactor experiments are over baselines which are too short to observe oscillation. From [73].

half the ν_μ have oscillated into other flavours. Atmospheric and accelerator neutrino experiments back up the hypothesis of ν_μ oscillating into mostly ν_τ with near maximal mixing and $\Delta m^2 \sim 2.5 \times 10^{-3} \text{ eV}^2$.

Atmospheric and solar results demonstrate two distinct regimes for neutrino oscillations. Two parameters, θ_{13} and $|\Delta m^2_{31}|$, have been associated with solar neutrino oscillation whilst two, θ_{23} and Δm^2_{32} have been associated with atmospheric neutrino oscillation. The final mixing angle, θ_{12} , has been measured by reactor experiments and T2K, whilst the sign of Δm^2_{31} , the octant of θ_{23} and a possible CP violating phase, δ_{CP} , are as yet unknown.

Quantifying these parameters is a goal of past, current and future experiments using neutrinos from a range of different sources.

3.1.2 Reactor neutrinos

Reactors provide an excellent source of man-made $\bar{\nu}_e$ for use in oscillation experiments. The long baseline experiment, KamLAND, has been mentioned in subsection 1.2.3 and provided excellent data relevant to the solar neutrino problem. Other experiments use shorter baselines of $\sim 1 \text{ km}$ [74] or lower. These include searches for oscillations at much higher Δm^2 and observations of the probability of $\bar{\nu}_e$ survival ($P(\bar{\nu}_e \rightarrow \bar{\nu}_e)$), usually with the inverse beta decay reaction described in Equation 1.1.

Several reactor experiments at short ($< 100 \text{ m}$) baselines began publish-

ing results from the early 1980s onwards. These include experiments at the ILL reactor[75], the Gösgen nuclear power reactor[76], the Bugey reactor[77][78], the Savannah River plant[79] and the Rovno nuclear power plant[80].

Of these short baseline results only the first measurements at Bugey[77] reported results inconsistent with no oscillations, and repeated measurements with improved shielding failed to back even these up[78]. This suggests an oscillation length above that accessible to reactor energies at very short baselines.

Further experiments on ~ 1 km baselines were performed in the late 1990s at the CHOOZ nuclear power station[81] and the Palo Verde Nuclear Generating Station[82]. These did provide useful information on oscillation parameters, ruling out $\nu_\mu \rightarrow \nu_e$ as a solution to the atmospheric anomaly (subsection 1.2.4) and constraining the parameter space for Δm_{32}^2 and θ_{13} .

Until recently short baseline reactor data appeared to agree with the hypothesis of no oscillations. A recent recalculation of reactor flux, though, resulted in a scaling up of predicted $\bar{\nu}_e$ which in turn implied a deficit in most previous reactor anti neutrino experiments. This has been dubbed the ‘reactor anti neutrino anomaly’[83] and oscillation with an additional sterile species of neutrino has been proposed as a possible solution. Future and proposed short baseline reactor experiments such as PROSPECT[84], SoLid[85], NEOS[86], SOX[87] and Stereo[88] aim to test the oscillation hypothesis as an explanation for this anomaly.

Reactor experiments over medium baselines have recently provided excellent information on the previously unknown θ_{13} mixing angle. In April 2012 RENO reported a 4.9σ $\bar{\nu}_e$ disappearance result[90], followed later that year by a $> 5\sigma$ result from Daya Bay[89] (shown in Figure 3.2. These were the first experimental confirmation of non-zero θ_{13} [†]. RENO and Daya Bay along with Double-CHOOZ[91] are fairly recent reactor experiments. They all involve shielded detectors containing Gd doped scintillator for detecting inverse β decay surrounded by undoped scintillator for catching γ s and a buffer of mineral oil. The experiments utilise near and far detectors, for measuring $\bar{\nu}_e$ before and after significant oscillation. RENO utilizes single near far detectors and Daya Bay uses three near and three far detectors positioned relative to six reactors. The use of identical near and far detectors allows uncertainties arising from reactor flux to be minimised.

Future and proposed experiments such as JUNO[92][93] and RENO-50[94] aim to use sharp energy resolution and high statistics to resolve the mass hierarchy at high significance. They also plan to determine known mixing parameters to higher

[†]T2K had data consistent with $\theta_{13} > 0$ at this point in time but had not reached the 5σ level of significance.

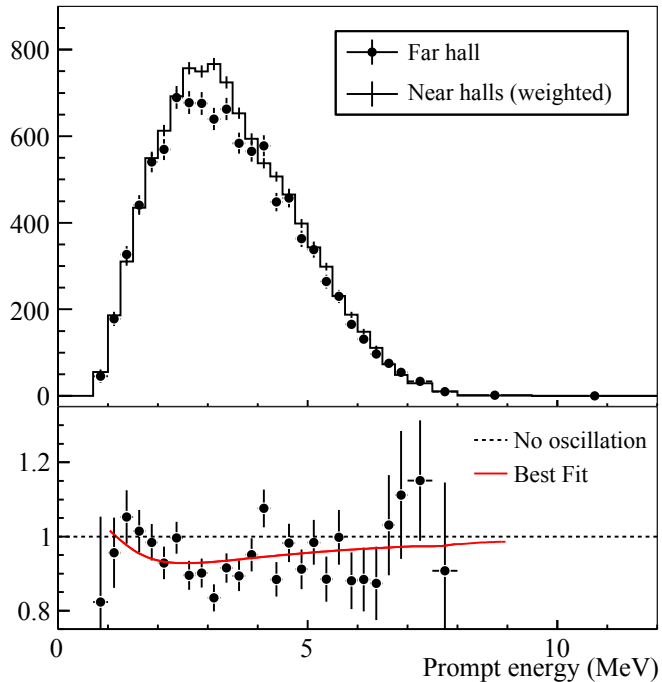


Figure 3.2: Daya Bay results confirming non-zero δ_{CP} at $> 5\sigma$. From [89].

than ever accuracy.

Measurements on $\bar{\nu}_e$ disappearance from reactors have been hugely important in solving the solar neutrino problem. More recently they have allowed us to determine that θ_{13} is indeed non-zero, a result confirmed by accelerator experiments such as T2K.

3.1.3 Accelerator neutrinos

Accelerators are the other man-made neutrino source. Their advantages are a larger degree of experimental control over neutrino energies and the ability to direct the beam allowing for measurable neutrino flux even at distances of many hundreds of kilometres. There remains difficulty in knowing precisely the neutrino content and energy spectrum of a beam and this is one of the larger sources of experimental uncertainty for accelerator experiments.

Prompt neutrino experiments such as those by CDHSW[95] and CHARM[96], were performed in the early 1980s. These found ratios of $(\bar{\nu}_e)$ and $(\bar{\nu}_\mu)$ events and of ν_μ and $\bar{\nu}_\mu$ events consistent with lepton universality. Because neutrino energies were high (above ~ 20 GeV) and baselines short, such experiments were only sensitive to oscillations with high Δm^2 , which were ruled out.

Perhaps the best known muon decay at rest experiment was LSND[97]. Using data taken during the 1990s they reported an apparent $\bar{\nu}_e$ appearance from a $\bar{\nu}_\mu$ source. Their short baseline and energies up to 200 MeV implied $\Delta m^2 \sim 1$ eV. By the time they published their final results solar and atmospheric mixing parameters of $\Delta m^2 \sim 1 \times 10^{-4}$ eV² and $\Delta m^2 \sim 1 \times 10^{-3}$ eV² respectively were already known. If the LSND appearance result was correct it would imply a third mass squared difference and therefore a fourth (sterile) neutrino species. The phase space suggested by LSND was restricted but not eliminated by reactor experiments, decay at rest experiments such as LAMPF[98] and KARMEN[99] and wide angle beams at the BNL[100][101], NOMAD[102] and CCFR[103], none of which found any evidence for $\nu_\mu \rightarrow \nu_e$ oscillations. Favoured LSND parameters and restrictions on their phase space are illustrated in Figure 3.3.

Finally MiniBooNE[105] was set up to search for $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam with similar $\frac{L}{E}$ to LSND. Their $\nu_\mu \rightarrow \nu_e$ results found no evidence for oscillation but their $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ results[106] appeared to suggest an excess consistent with oscillations in the $0.1 \text{ eV}^2 - 1 \text{ eV}^2$ range. Furthermore their final results[104] suggested a low energy excess in both modes as shown in Figure 3.4, but only marginally




Figure 3.3: Favoured oscillation parameters implied by the LSND anomaly, with exclusion regions from other experiments indicated. From [97].

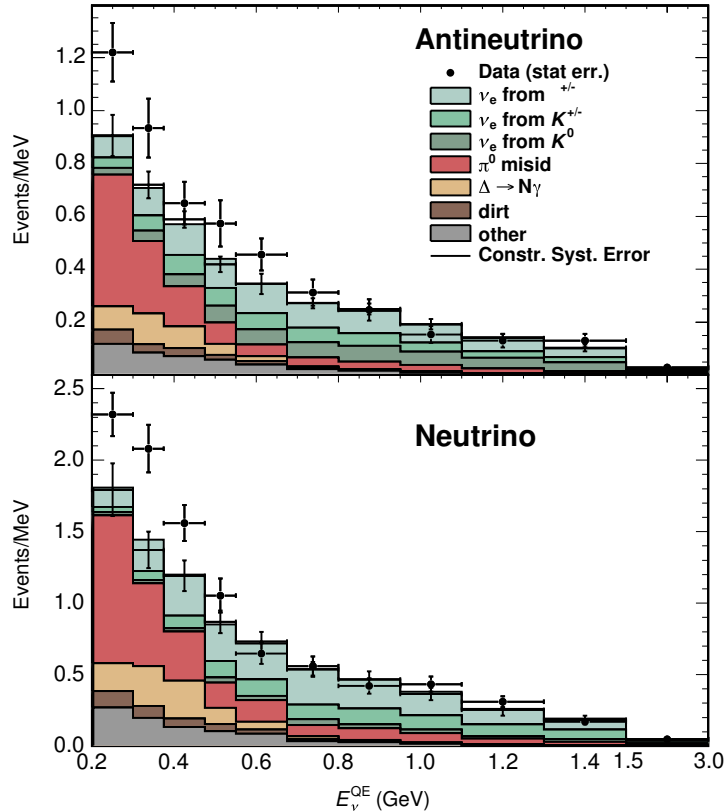


Figure 3.4: Excess of $\bar{\nu}_e$ at low energies as seen by the MiniBooNE experiment. From [104].

compatible with two-neutrino oscillations.

MicroBooNE[107], a new liquid argon TPC neutrino R&D experiment on the Booster beamline at FNAL, is due to start taking neutrino data soon. The experiment should be able to tell if the MiniBooNE excess is electron like or photon like. In doing so it will tell us whether the excess could be from some new oscillations with sterile neutrinos or simply a previously unmodelled background. It will be joined by the ICARUS T600[108] and SBND[109] detectors to form a full short baseline program addressing the LSND and MiniBooNE anomalies[110].

Many decay in flight accelerator experiments throughout the 1980s and 1990s used short baselines which were insensitive to small values of Δm^2 and thus unable to solve the solar and atmospheric neutrino problems. Such experiments included those the BNL during the 1980s[111][100], decay in flight results from CHARM[112], BEBC[113] and CDHSW[114] at CERN and CCFR[115] and the high energy NuTeV[116] at Fermilab. CHORUS[117] at CERN and E531 at[118] FNAL aimed to find ν_τ from $\nu_\mu \rightarrow \nu_\tau$ oscillations. Unsurprisingly none of these short baseline accelerator experiments were able to find evidence for oscillations,

placing instead strong limits on the available parameter space.

When it came to oscillation parameters nature was kinder to long baseline accelerator experiments. These are more recent efforts which utilise near detectors to measure unoscillated neutrinos as well as a far detector, vastly reducing systematic uncertainties relating to neutrino flux in the unoscillated beam. The K2K experiment, as mentioned in subsection 1.2.4 was one example.

The NuMi beam at Fermilab[119] provides neutrinos for both long baseline oscillation and short baseline cross section and R&D experiments. MINOS[56] was one example which supported the Super-Kamiokande atmospheric oscillation results as seen in Figure 3.5. Its successor MINOS+ [120] is the only currently operating wide-band on axis oscillation experiment, providing unique physics opportunities. Both use a near detector 1 km and a far detector 735 km from the source. NO ν A [121] uses the same source but a different off-axis angle (0.8° from the beam centre) giving a narrow band of energies around 2 GeV seen by a far detector 810 km away. NO ν A examines the same parameters as T2K, aiming to measure $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam and also precisely measure $\bar{\nu}_\mu$ survival.

Other experiments have searched for ν_τ appearance in a beam of ν_μ . Both CHORUS[117] and NOMAD[102] searched unsuccessfully for such oscillations using




Figure 3.5: MINOS results from atmospheric mixing, also showing the Super-Kamiokande results. From [56].

protons produced by CERN’s SPS proton source over a short baseline. The later OPERA experiment[122] used a much longer baseline and were able to observe ν_τ appearance at 5σ [123].

The T2K accelerator experiment uses Super-Kamiokande as a far detector, like K2K, and a source and near detector at J-PARC in Tōkai-mura . The detectors and beam providing their neutrinos are discussed in detail in Chapter 4.

Hyper-Kamiokande[124] is a proposed next generation Čerenkov detector using 1 Mton of ultra pure water. The experiment aims to use an upgraded J-PARC beam to make high statistics measurements of δ_{CP} and $\sin^2 \theta_{23}$, as well as continuing the observations of atmospheric, solar and astrophysical neutrinos made by Super-Kamiokande.

DUNE[125] is another future experiment, using a 40 kton liquid argon time projection chamber to measure neutrino interactions. The experiment plans to use a 1300 km baseline with neutrinos from a high intensity beam at Fermilab in Illinois to the Sanford lab in Lead, South Dakota, to obtain high precision measurements on δ_{CP} .

3.1.4 Solar neutrinos

Modern solar neutrino experiments often aim to observe more than just neutrino oscillation, using neutrinos to probe nuclear interactions inside the sun and check stellar models. Even where this is the case, disappearance of $\bar{\nu}_e$ makes entanglement with neutrino oscillation theory unavoidable.

Super-Kamiokande has continued observing solar neutrinos[126]. As well as improving the fit on solar oscillation parameters they have found evidence for non-zero asymmetry between neutrino flux at day and night at 2.8σ . This results from matter effects encountered by neutrinos propagating through the Earth, leading to some regeneration of oscillated ν_e . The Super-Kamiokande solar neutrino results also currently give the world’s most precise ν_e -only measurement of Δm^2_{21} .

Borexino[127] is a liquid scintillator based experiment designed to observe sub-MeV neutrinos, specifically mono-energetic ${}^7\text{Be}$ solar neutrinos. They have been able to detect, at over 5σ , geoneutrinos from the Earth[128] produced by the decay chains of U, Th and K contributing to low energy $\bar{\nu}_e$ flux[129]. Borexino’s geoneutrino measurements are supported by results from KamLAND[130].

SNO+ [131] is a follow-up to SNO, replacing the deuterium target with liquid scintillator. This allows neutrinos at much lower energies to be examined. SNO+ aims to observe neutrinoless double beta decay, low energy solar neutrinos, geoneutrinos and reactor neutrinos.

Hyper-Kamiokande[132] is predicted to observe around 115000 ^8B neutrinos a year from elastic scattering and also to have high sensitivity to day–night asymmetry.

3.1.5 Atmospheric neutrinos

Future experiments on atmospheric neutrinos propose to probe the remaining mysteries of neutrino oscillation parameters.

One proposed neutrino experiment is India’s INO[133]. A main cavern is to house the main iron calorimeter (ICAL) detector while surrounding caverns can provide space for future experiments, such as neutrinoless double beta decay experiments. They aim to use the ICAL to make precision measurements on atmospheric neutrinos including measurements of matter effects.

ORCA[134] is a study to examine the feasibility of determining mass hierarchy through two dimensional arrival pattern of ν_μ in energy and zenith angle. The detector plans to use a deep sea neutrino telescope developed for KM3NeT and distinguish between the hierarchies by examining matter induced oscillations.

PINGU[135] is a proposed detector to be located with the IceCube Deep-Core detectors. It aims to be sensitive to mass hierarchy by observing atmospheric neutrinos. The detector will closely follow the IceCube design, with optical modules deposited in the ice.

3.1.6 Global fit to mixing parameters

With the large number of experiments on neutrino oscillation parameters available some authors have brought data together in a global fit. A recent paper[51] determined the values shown in table 3.1.

Notably the atmospheric mixing angle, θ_{23} , is consistent at 1σ with the hypothesis of maximal mixing ($\theta_{23} = 45^\circ$) for a normal hierarchy, and at 1.3σ in the case of an inverted hierarchy.

3.2 Neutrino interactions with matter

Because of the challenges inherent in observing neutrinos, experimental information on neutrino interactions and cross sections is difficult to acquire. The processes contributing to neutrino cross sections are detailed in Section 2.3. Much of the information used for predicting and simulating neutrino interactions comes instead from the much larger supply of electron scattering data.

Parameter	Best fit value $^{+1\sigma}_{-1\sigma}$		
Normal hierarchy	Δm^2_{21}	$7.60^{+0.19}_{-0.18} \times 10^{-5} \text{ eV}^2$	
	$ \Delta m^2_{31} $	$2.48^{+0.05}_{-0.07} \times 10^{-3} \text{ eV}^2$	
	$\sin^2 \theta_{12}$	$3.23^{+0.16}_{-0.16} \times 10^{-1}$	
	$\sin^2 \theta_{23}$	$5.67^{+0.32}_{-1.24} \times 10^{-1}$	
	$\sin^2 \theta_{13}$	$2.26^{+0.12}_{-0.12} \times 10^{-2}$	
	θ_{12}	$34.6^{+1.0}_{-1.0} \text{ }^\circ$	
	θ_{23}	$48.9^{+1.8}_{-7.2} \text{ }^\circ$	
	θ_{13}	$8.6^{+0.3}_{-0.2} \text{ }^\circ$	
	δ_{CP}	$1.41^{+0.55}_{-0.40} \pi$	
	δ_{CP}	$254^{+99}_{-72} \text{ }^\circ$	
	Inverted hierarchy	Δm^2_{21}	$7.60^{+0.19}_{-0.18} \times 10^{-5} \text{ eV}^2$
		$ \Delta m^2_{31} $	$2.43^{+0.05}_{-0.06} \times 10^{-3} \text{ eV}^2$
$\sin^2 \theta_{12}$		$3.23^{+0.16}_{-0.16} \times 10^{-1}$	
$\sin^2 \theta_{23}$		$5.73^{+0.25}_{-0.39} \times 10^{-1}$	
$\sin^2 \theta_{13}$		$2.29^{+0.12}_{-0.12} \times 10^{-2}$	
θ_{12}		$34.6^{+1.0}_{-1.0} \text{ }^\circ$	
θ_{23}		$49.2^{+1.5}_{-2.3} \text{ }^\circ$	
θ_{13}		$8.7^{+0.2}_{-0.2} \text{ }^\circ$	
δ_{CP}		$1.48^{+0.31}_{-0.31} \pi$	
δ_{CP}		$266^{+56}_{-56} \text{ }^\circ$	

Table 3.1: Table showing most recent global fit neutrino oscillation parameters from[51].

The growing use of powerful accelerators to produce high neutrino fluxes provides many opportunities for experiments examining in detail the interactions of neutrinos with nuclei. This is a two way exchange, with information on neutrino cross sections feeding into oscillation experiments. Detectors such as MINERvA [136] and SciBooNE[137] can examine cross sections relevant to oscillation experiments and T2K's ND280 near detector itself acts as a useful source of cross section information.

3.2.1 MiniBooNE results

Cross section measurements were produced alongside MiniBooNE’s short baseline oscillation analyses. These covered many types of low energy neutrino interactions on mineral oil (CH_2)[139].

The CCQE results are of particular interest since the extracted value of the parameter M_A (see subsection 2.3.2) were about 30% higher than the global average from experiments such as NOMAD, albeit when operating at lower energies than other experiments. This discrepancy is illustrated in Figure 3.6. It prompted further development of nuclear interaction models to explain the discrepancy and highlights the need for detailed data on neutrino interactions to feed in to theoretical predictions.

3.2.2 Argon TPCs

Also of interest are experiments using liquid argon TPC (LArTPC) technology, the development of which is particularly important in anticipation of the DUNE long baseline experiment. The short baseline program at Fermilab[110] has also been




Figure 3.6: ν_μ and $\bar{\nu}_\mu$ charged current quasi-elastic cross sections for MiniBooNE and NOMAD. Note the two experiments use different topologies in defining such events so not all discrepancies are necessarily physical. From [138].

mentioned with regards to the LSND and MiniBooNE excesses and should feature the SBND[109], MicroBooNE[107] and ICARUS T600[108] detectors acting as near, intermediate and far detectors respectively.

Of particular note is the ArgoNeuT [140] LArTPC. This in conjunction with the MINOS near detector determined an integrated $\bar{\nu}_\mu$ cross section on argon of

$$\sigma/E_{\nu_\mu} = (0.66 \pm 0.03 \pm 0.08) \times 10^{-38} \frac{\text{cm}^2}{\text{GeV}},$$

for ν_μ and

$$\sigma/E_{\bar{\nu}_\mu} = (0.28 \pm 0.01 \pm 0.03) \times 10^{-38} \frac{\text{cm}^2}{\text{GeV}},$$

per nucleon for $\bar{\nu}_\mu$ [141]. Other interesting results include the observation of back to back proton events[142][143] (Figure 3.7) and coherent charged pion production[144], all in a liquid argon TPC.

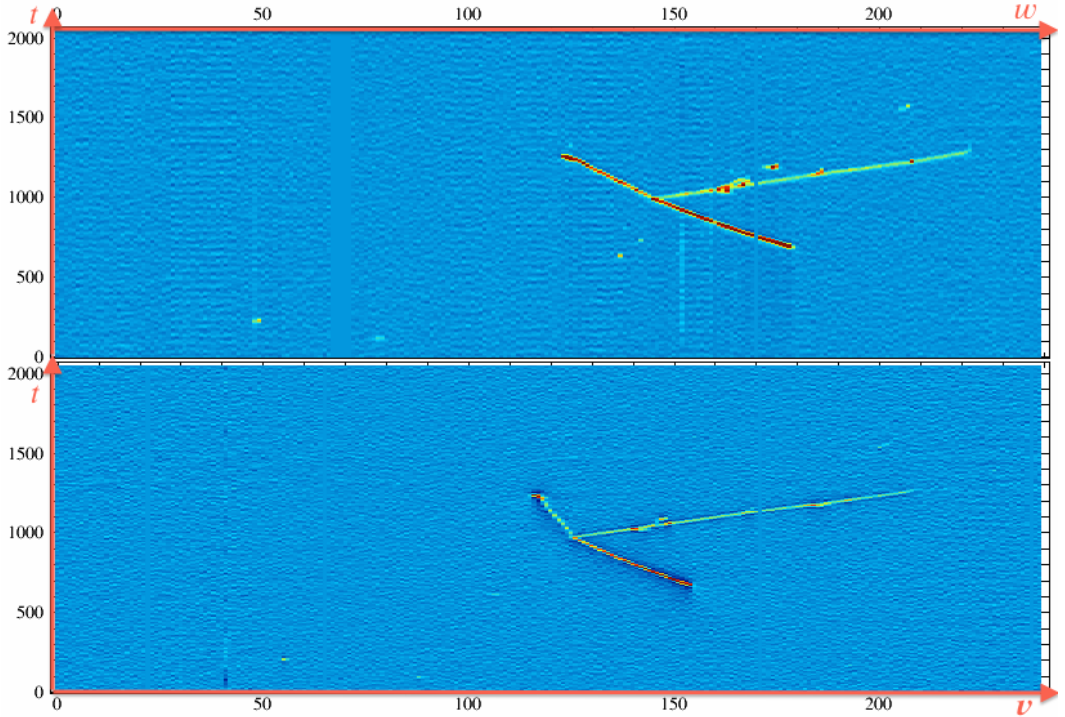


Figure 3.7: Two 2D views of a back to back proton event seen by ArgoNeuT, with protons ~ 500 MeV. Coordinates are the IDs of the collection planes picking up charge (w and v for collection and induction respectively) and time (t). From [142].

As an argon TPC with a large amount of data behind it ArgoNeuT invites comparison with the T2K TPC gas interaction analysis covered in this thesis. There are, though, substantial differences between the detectors. ArgoNeuT observes neutrinos at much higher energies than the ND280, with $\bar{\nu}_\mu$ energies averaging 9.6 GeV (3.6 GeV). This compares with T2K's average around 600 MeV. ArgoNeuT is also restricted to a fairly narrow range of muon angles by a requirement for tracks matching the downstream MINOS near detector whereas the ND280 reconstruction for gas interactions has been designed to be as isotropic as possible. Finally ArgoNeuT uses liquid argon rather than the ND280 TPC argon gas, raising the threshold for the detection of low energy protons.

Chapter 4

The T2K Detectors

The T2K experiment[145] is designed to observe neutrino oscillations over a baseline of 295 km. Specifically, a beam of either mostly ν_μ or mostly $\bar{\nu}_\mu$ is produced at the J-PARC facility in Tōkai-mura on the East coast of Japan by high energy protons striking a target. The beam is observed at the Super-Kamiokande facility near the West coast where the detector's 50 kton volume of water ensures an appreciable number of neutrino interactions. A near detector facility located 280 m downstream of the proton beam target target allows the fluxes of ν_μ , ν_e , $\bar{\nu}_\mu$ and $\bar{\nu}_e$ in the beam to be determined and also allows for measurements of neutrino cross sections. An overview of the baseline is shown in Figure 4.1.

T2K was designed to measure the mixing angle θ_{13} through electron neutrino appearance

$$\langle \bar{\nu} \rangle_\mu \rightarrow \langle \bar{\nu} \rangle_e, \quad (4.1)$$

and to precisely measure the atmospheric mixing parameters θ_{23} and Δm_{23}^2 through muon neutrino survival

$$\langle \bar{\nu} \rangle_\mu \rightarrow \langle \bar{\nu} \rangle_\mu, \quad (4.2)$$

In 2011 T2K published first indications of ν_e appearance in a ν_μ beam, at 2.5σ significance[146]. This was confirmed in 2013 with a significance of 7.3σ [147]. This evidence, along with reactor data from experiments such as RENO[90] and Daya Bay[89], confirm non-zero θ_{13} at high significance. Despite being unknown a few years ago, θ_{13} is now the most precisely measured mixing angle.




Figure 4.1: The T2K experiment layout, showing the neutrino beam’s journey across Honshū, from neutrino production at J-PARC to the near detector 280 m downstream to the Super-Kamiokande far detector 295 km away.

4.1 Beam

The setup for producing the T2K neutrino beam is illustrated in Figure 4.2. Neutrinos are produced at the J-PARC neutrino experiment facility[148] where accelerated protons from the main ring strike a graphite target. The secondary particles from these collisions are selected and steered into a decay pipe, where they decay into neutrinos.

T2K utilises an off-axis neutrino source produced by directing a wide band beam in a slightly offset direction relative to the detectors. In this case the beam is aimed 2.5° below the direction of the far detector. For a far detector at 295 km, optimum ν_e appearance probability occurs when $E_\nu \sim 600$ MeV (see Figure 4.3). Although aiming the beam 2.5° away from the target reduces overall neutrino flux it actually increases flux at this maximum, giving a sharp peak near 600 MeV as can be seen in Figure 4.4. Background is also reduced.

4.1.1 Off-axis configuration

T2K spearheaded this off-axis configuration. Its advantages stem simply from the two body π decay producing the majority of neutrinos. In the decay’s centre of mass frame, with negligible neutrino mass,

$$E_{\nu\text{CM}} = \frac{m_\pi}{2} \left(1 - \frac{m_\mu^2}{m_\pi^2} \right), \quad (4.3)$$

which is boosted to the laboratory frame along the z axis. Assuming a small angle θ between π and ν direction and that $\beta \sim 1$, we find

$$E_\nu = \left(1 - \frac{m_\mu^2}{m_\pi^2} \right) \left(\frac{E_\pi m_\pi^2}{m_\pi^2 + E_\pi^2 \theta^2} \right). \quad (4.4)$$




Figure 4.2: The T2K setup for neutrino beam production.




Figure 4.3: Neutrino oscillation probability as a function of energy, for neutrinos starting as ν_μ , at 295 km, assuming $\delta_{CP} = 0$ and using recent oscillation parameters from [51]. Figure 4.3a shows normal hierarchy and Figure 4.3b inverted.




Figure 4.4: Effect of off-axis angle on J-PARC fluxes. Neutrino energy as a function of pion energy at different angles is shown in Figure 4.4a, where the dotted lines give the relationship between E_ν and E_π at different angles, the solid line gives the 2.5° angle used by T2K and the blue line shows our desired peak energy. Neutrino fluxes at different off-axis angles are shown in Figure 4.4b, from [149].

As $\theta \rightarrow 0$ the left side of this equation becomes directly proportional to E_π , and neutrino energy gains the same spread as pion energy. As θ increases, though, the denominator's dependence on E_π increases faster than the numerator's leading to its overall suppression for a range of off-axis angles. Since E_π is the only variable quantity in Equation 4.4 this greatly reduces the spread of E_ν . The effect of θ on neutrino energy distribution is shown by Figure 4.4a.

This off-axis configuration has since been taken up by experiments such as NOvA [121]. The Hyper-Kamiokande successor to Super-Kamiokande plans to use the same off-axis angle as T2K[124].

4.1.2 Accelerator

In order to produce a beam of accelerated protons the J-PARC facility contains three accelerators; the linear accelerator (LINAC), the rapid cycling synchrotron (RCS) and the main ring (MR). The direction of our beam and location detectors relative to this main ring are shown in Figure 4.5.

The LINAC accelerates H^- anions before they are converted to H^+ by charge stripping foils. The protons are injected into the RCS where they are accelerated




Figure 4.5: The geographical location of the J-PARC neutrino beam production facility used by T2K relative to the main ring, from [124].

to 3 GeV. Most bunches are passed to other facilities but about 5% proceed to the main ring where they are accelerated in 8 bunches to 30 GeV for use by J-PARC's hadron and neutrino facilities.

For the neutrino beam a fast extraction takes all 8 bunches at once to produce a 30 GeV proton pulse every 2.1 s. These pulses are directed towards the T2K graphite target via the primary beamline. At this point the beam is directed downwards at 3.637° , providing the experiment's 2.5° off-axis angle relative to the near and far detectors.

4.1.3 Secondary beamline

The proton target, magnetic horns, decay volume and beam dump are all contained in an iron vessel filled with 1 atm of He. This reduces production of unwanted ^3H and N_2O and slows oxidation of the target.

The core of the target is a 91.4 cm (or ~ 2 interaction lengths) long, 2.6 cm diameter graphite rod. Secondary particles, mostly π s with some contamination from Ks, are produced when the proton pulses impinge on this target.

From these secondary particles, pions of the desired sign are selected and focussed using three magnetic horns[150]. These horns were designed for a peak current of 320 A, though in practice a peak of 250 A has been used for most of the T2K run periods. The horns contain two aluminium conductors in a coaxial arrangement. The first horn serves to collect the pions, while the second and third focus them.

For the majority of the T2K run periods π^+ have been selected, providing a beam of ν_μ . More recently the horn current has been reversed for some run periods to select π^- , providing a $\bar{\nu}_\mu$ beam. The gas interaction analysis described in this thesis uses the forward horn current configuration and a beam of mostly ν_μ .

The selected pions are directed to the decay volume. This is a steel tunnel around 96 m long, surrounded by 6 m thick reinforced concrete. The beam dump at the end of this volume is a core of graphite 3.174 m long, containing iron plates to a total thickness of 2.40 m. This large lump of solid material allows only neutrinos and muons above ~ 5 GeV to pass and reach the muon pit.

4.1.4 Muon flux

Since muons are primarily produced via two body π decay, measuring those escaping from the beam dump allows the intensity and direction of the neutrino beam to be inferred with high accuracy. The T2K neutrino beam direction is assumed to be the

same as the direction from the target to the centre of the muon flux.

A dedicated muon monitor (MUMON) setup is located just behind the beam dump featuring arrays of ionization chambers and arrays of silicon photodiodes, both ~ 120 m from the target. These arrays can resolve muon beam direction to ~ 2.5 mrad.

4.2 Super-Kamiokande

The Super-Kamiokande detector's first data taking period (phase I) ran from 1996 to 2001. Severe damage during upgrade work lead to the next phase (phase II) running with only half the original PMTs, from 2002 to 2005. A third data taking period (phase III) ran from 2006 to 2008 with almost the original number of PMTs, and in 2008 the detector was upgraded to its most recent phase (phase IV) with new electronics[152]. Its role in historic neutrino oscillation measurements has been described in Chapter 1.

Super-Kamiokande acts as the T2K far detector, doing the job of observing $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ appearance and $\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu$ survival in oscillated neutrinos. It is a large water Čerenkov detector located 1 km deep in a cavern in Mt. Ikeno, as shown in Figure 4.6, 295 km away from the neutrino source.




Figure 4.6: The Super-Kamiokande detector, showing the huge volume of the water Čerenkov inner detector, from[151]

4.2.1 The detector

The detector itself consists of 50 kton of highly purified water, split into the optically separated inner and outer detectors. The inner detector is surrounded by 11129 PMTs (originally 11146) each 50 cm in diameter, providing 40% coverage of the inner surface. This provides a high level of spatial resolution for observing the rings of photons produced by cones of Čerenkov light. The outer detector is a relatively sparsely instrumented region which serves to muons from cosmic rays. The total of 1885 PMTs arrayed around it are sufficient for cosmic muon rejection at almost 100% efficiency.

The PMTs are sensitive to magnetic fields, which alter the trajectories of photoelectrons within them producing an asymmetric reduction in hit collection efficiency. For this reason coils are arranged around the tank's inner surface to neutralise the Earth's natural magnetic field. The remaining field around the PMT dynode axis is around $2 \mu\text{T}$.

4.2.2 Particle identification

The primary requirement for the detector is an ability to distinguish e produced by $\bar{\nu}_e$ from μ produced by $\bar{\nu}_\mu$. In Super-Kamiokande this is done by differentiating between the Čerenkov rings produced by electrons and muons. This setup along with the lack of magnetic field makes it impossible to distinguish the positive leptons produced by $\bar{\nu}$ from the negative leptons produced by ν .

Muons are fairly massive in comparison to electrons and resist changes to their momentum. This produces a sharp, well defined Čerenkov cone leading to a sharp, well defined ring of PMT hits. Electrons by contrast are subject to scattering, almost invariably producing EM showers. This leads to a 'fuzzy' ring of PMT hits as separate cones of Čerenkov light from the shower's constituent e^+ and e^- overlap.

Unfortunately this leaves the Super-Kamiokande $\bar{\nu}_e$ appearance analyses sensitive to background from other showering particles. The significant backgrounds of this kind are π^0 from neutral current interactions. These are electrically neutral and thus invisible until they decay, predominantly into electromagnetically showering particles ($\gamma\gamma$ 99% of the time)[149]. From the detector's point of view the only difference between these pairs and e signal is that each γ can travel some distance before showering. The difficulty of distinguishing these decays from $\bar{\nu}_e$ interactions motivated the construction of the ND280 π^0 detector (POD, subsection 4.4.4) to get a handle on the background they are expected to contribute.

4.3 INGRID

The ND280 facility consists of detectors designed to measure the properties of the T2K neutrino beam 280 m downstream of the production target, before the neutrinos have a chance to significantly oscillate. One detector here is the ND280 off-axis near detector itself. The other is INGRID[153], the on-axis detector.

INGRID (the interactive neutrino grid) is a cross shaped array of 14 identical modules centred on the beam axis as illustrated in Figure 4.7, plus one slightly modified proton module placed in the centre of the INGRID cross, between the horizontal and vertical modules. The purpose of INGRID is to directly monitor the direction of the neutrino beam, with heavy iron targets providing a sufficient rate of neutrino interactions for day by day measurements during normal running. The module is able to measure the position of the beam centre to better than 0.4 mrad.




Figure 4.7: The INGRID on-axis detector as of the start of T2K running. Note that in early 2014, the shoulder module on the right was moved to a lower level at a 1.65° off-axis angle to measure Fe cross sections below 1 GeV. The proton module is located between the horizontal and vertical modules.

4.3.1 The detector array

The individual modules are illustrated in Figure 4.8a. They consist of 9 iron plate targets each sandwiched between two scintillator planes for tracking, all surrounded by veto planes of scintillator. The veto planes and fiducial volume are all enclosed within dark boxes formed from plastic plates held together in aluminium frames. In each module the majority (95%) of mass comes from the 7.1 ton of iron target. The two planes of scintillator in each layer each consist of two arrays of 24 bars; one in the horizontal and one in the vertical direction.

4.3.2 The proton module

The proton module illustrated in Figure 4.8b has a similar shape is slightly modified compared to the other modules. It was designed to detect the muons and protons produced by CCQE interactions with high efficiency. The module is fully active to allow for better tracking. It contains 34 tracking planes and 6 veto planes. The tracking planes again alternate between horizontal and vertical array of bars but in this case there are 32 with the middle 16 bars in the array being half the width and slightly thicker than the outer bars.

4.4 ND280

The ND280 off-axis near detector is designed to complement Super-Kamiokande by determining its expected signal and background rates. To accomplish this it measures neutrino flux, energy spectrum, flavour content and cross sections from the unoscillated J-PARC neutrino beam.

The obvious background to a $\bar{\nu}_e$ appearance measurement is $\bar{\nu}_e$ contami-




Figure 4.8: The INGRID modules, with standard modules shown in Figure 4.8a and the specialised proton module in Figure 4.8b.

nation in the beam, so the near detector is equipped to measure this. The beam's initial $\bar{\nu}_\mu$ content is also measured. Knowing to reasonable precision initial $\bar{\nu}_\mu$ content and $\bar{\nu}_e$ contamination is crucial for T2K's oscillation analysis.

Another significant background, as raised in subsection 4.2.2, is π^0 produced by neutral current neutrino interactions. Their decays into predominantly electromagnetically showering particles are difficult to distinguish from $\bar{\nu}_e$ charged current interactions producing electrons. In order to calculate the expected backgrounds from these interactions a dedicated π^0 detector as included in the ND280.

Practical considerations and the requirement for many different measurements led to the construction of a near detector which is quite different from the far detector. The ND280, as shown in Figure 4.9, consists of many sub-detectors surrounded by a large magnet. The POD (π^0 detector), TPC (time projection chamber), FGD (fine grained detector) and downstream ECal (electromagnetic calorimeter) sub-detectors are contained in a metal frame (the basket) surrounded by more




Figure 4.9: The ND280 off-axis near detector.

ECals (the P0D ECals and barrel ECals respectively), the UA1 magnet and the SMRD (side muon range detector).

Each sub-detector is specialised to perform a specific job. The P0D is a large target designed to examine interactions of π^0 as discussed. The FGDs provide dense neutrino interaction targets with the capacity for reconstruction of vertices and through-going tracks. The TPCs are sparse and designed to provide excellent tracking, reconstruction and particle identification for tracks originating in the P0D or FGDs. The ECals surrounding these detectors pick up escaping electromagnetic particles and also allow for particle identification. Finally, the SMRD detects escaping muons and provides triggers and vetoes for cosmic ray events.

4.4.1 The UA1 magnet

The ND280 dipole magnet was appropriated from the NOMAD experiment which acquired it from UA1. In order to allow the momentum and sign of charged particles to be measured it provides a field of 0.2 T. The dipole field is produced by a 2900 A current passed through aluminium coils. During running they are cooled by water.

The magnet's flux return yokes consist of low carbon steel plates. The two yokes are C shaped and enclose the ND280 sub-detectors but can be opened for access to the detector sub-modules. The SMRD modules are situated within the gaps of the yokes whilst the other detectors are located between the two symmetric halves.

The magnetic field effected by the magnet was measured to high precision during a dedicated mapping procedure in late 2009. The results of this mapping are utilised for magnetic field related corrections and systematic uncertainties during reconstruction and analysis.

4.4.2 SMRD

The ND280 side muon range detector[154] was designed for energy measurement of muons exiting the detector at high angles. It also functions as both a trigger and an active veto for cosmic rays entering the detector.

It consists of 3–6 layers of plastic scintillator embedded in the UA1 magnet's iron yokes. The sheets of scintillator are 875 mm \times 167 cm \times 7 mm or 875 mm \times 175 mm \times 7 mm depending on the gap they are inserted into, and the iron plates they come between are 4.8 cm. Scintillation light reaches a wavelength shifting (WLS) fibre embedded in an S-shaped groove in the scintillator and is read out by a multi-pixel photon counter (MPPC). In total, 440 of these modules are embedded

in the innermost gaps of the magnet. Readout is double ended, allowing hit position along the bar to be estimated by time difference between signal at each end.

4.4.3 ECals

The ND280 electromagnetic calorimeters[155] surround the P0D and the central tracker region (the TPCs and FGDs). There are 13 modules in total; 6 P0D ECals and 6 barrel ECals affixed to the magnet yoke surround the P0D and tracker region respectively and a single downstream ECal is mounted in the basket immediately downstream of the tracker region.

Each module contains sheets of optically isolated scintillator bars interleaved with lead target planes. In the tracker and downstream ECals these bars are oriented in alternating longitudinal and perpendicular layers whilst in the P0D ECals they are oriented in longitudinal layers only. Each bar has a cross section of $40 \text{ mm} \times 10 \text{ mm}$ (width \times depth). The scintillator bars have WLS fibres running through them to transport scintillation light to MPPC photo sensors at their edges. The downstream ECal bars and z -direction barrel ECal bars are read out at both ends, whilst the x - and y -direction barrel ECal bars and all P0D ECal bars are read out at one end only. The single ended bars are all mirrored at the uninstrumented end.

The barrel and downstream ECals both feature 1.75 mm layers of lead interleaved with scintillator layers. In the downstream ECal there are 34 layers with $2000 \text{ mm} \times 2000 \text{ mm}$ sheets of scintillator. In all tracker ECals there are 31 layers of scintillator with dimensions of $3840 \text{ mm} \times 2280 \text{ mm}$ for the side modules and 3840×1520 for the top and bottom modules.

The P0D ECals complement the P0D itself and have a modified design compared to the other ECals. Space constraints lead to the P0D being substantially shallower than other modules, with a total depth of 155 mm as opposed to 462 mm in the barrel. To compensate a shorter radiation length is obtained by using thicker lead layers. Six 4.0 mm thick layers of lead are interleaved with six layers of scintillator. All bars are longitudinal, with scintillator dimensions $2340 \text{ mm} \times 2760 \text{ mm}$ for the side modules and $2340 \text{ mm} \times 1520 \text{ mm}$ for the top and bottom modules.

The ECal modules were assembled at a number of institutions throughout the UK, including substantial involvement from Warwick on the P0D ECals and MPPC and WLS fibre quality assurance.

4.4.4 P0D

The π^0 detector[156] is primarily designed to measure processes with a π^0 in the final state. It features planes of horizontal and vertical scintillator bars with triangular cross section interleaved with planes of target material, as shown in Figure 4.10. The scintillator bars have WLS fibres running through them connected to MPPCs.

In the central region the target consists of brass sheets and fillable water bags. Either end of this are ECal regions where the targets are lead sheets. These ECal sections act as veto regions for entering particles, allowing for rejection of external backgrounds, and improve the P0D 's containment of electromagnetic particles.

4.4.5 FGDs

The fine grained detectors[157] in the main tracking region of the detector use plastic scintillator as both a target and a detector. This provides a large target mass for neutrino interactions as well as measurement of particles emerging from the vertices. Through-going particles can also be reconstructed.

The detectors feature bars of scintillator in planes of bars oriented in the x

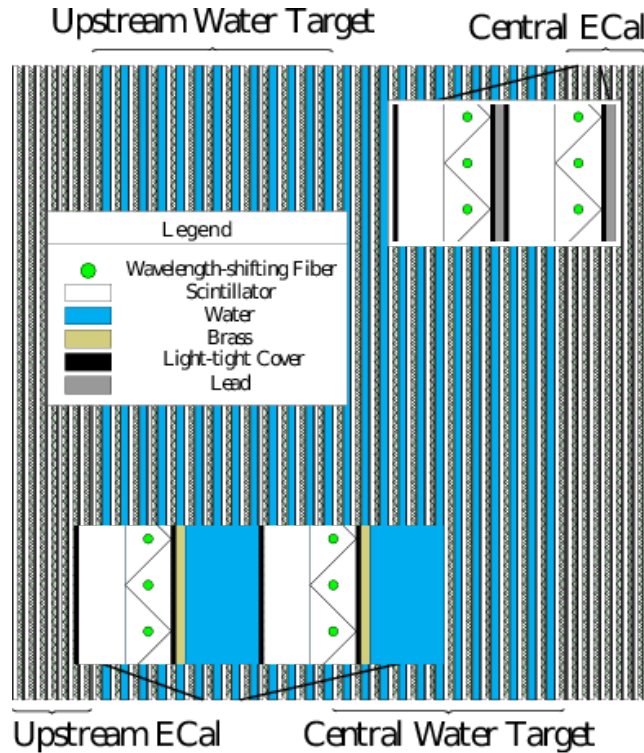


Figure 4.10: The ND280 P0D , showing the triangular layout of scintillator planes and the lead, brass and water targets.

or y direction. Wavelength shifting fibres passed through the bars transport light to MPPCs attached to one end, while the other end is mirrored. The most upstream FGD (FGD1) uses only plastic scintillator as a target while the downstream FGD (FGD2) also features six 2.5 cm thick layers of water targets. This allows for cross sections on water to be extracted by comparing interaction rates in the two FGDs.

4.4.6 TPCs

The time projection chambers[158][159] provide excellent three dimensional imaging of through-going tracks. The 0.2 T magnetic field applied to the ND280 and detailed in subsection 4.4.1 causes charged tracks to curve, allowing for their momentum and sign to be determined. This ensures that event rate as a function of neutrino energy can be calculated. The TPCs also collect information on energy deposited per unit length ($\frac{dE}{dx}$) which allows for particle identity to be determined. Momentum resolution in the TPCs is around 7%.

Each TPC contains a box of drift gas surrounded by an insulating box of CO_2 as illustrated in Figure 4.11. In the drift gas passing charged particles effect ionisation electrons. A central cathode creates a $\sim 2.5 \text{ kV/m}$ E field in the x




Figure 4.11: An illustration of a ND280 TPC detector, showing the drift volume and detector pads for one half.

direction which causes the electrons to drift towards the left or right sides of the TPCs. Here they pass a ~ 350 V micromesh before being picked up by micromegas pads. The location of the pad detecting the electrons provides the y and z position of the initial ionisation while the x position is determined from the time of the signal.

The total drift volume has dimensions of $897 \text{ mm} \times 1800 \text{ mm} \times 772 \text{ mm}$, with the co-ordinates corresponding to drift \times height \times beam. Each of the three TPCs has two drift regions separated by the central cathode, giving a total of 6 in the ND280.

The gas mixture contains 95% Ar. The remaining 5% comes from a CF_4 concentration of 3% and an iC_4H_{10} concentration of 2%. This mixture performs well with the micromegas units, gives fairly low diffusion and high drift speed, has a long attenuation length and is non-flammable. In total there are 3000 L of gas mixture per TPC, or 9000 L in total, and about 5 times this volume is cycled through the TPCs each day with most of the outgoing gas being purified and recycled. The outer gas volume's insulating CO_2 is kept at a slightly negative pressure relative to the TPC gas to avoid contaminating it. The gas system maintains the stability of both volumes and minimises O_2 contamination.

The micromegas pads are each $9.8 \text{ mm} \times 7.0 \text{ mm}$, largest in the z dimension, and grouped into modules containing 1728 pads each. The pads are grouped into 2 columns, each containing 6 pads, on each side of the central cathode in each TPC giving a total of 24 modules per TPC or 72 in total.

As shown in Figure 4.12 the ND280 TPCs allow for the tracks produced by charged particles to be resolved with high precision. Since their primary goal is as a tracker rather than a target they are not constrained by the need for very high density. This means that even short tracks from low energy protons can be resolved. The trade-off is a small total number of interactions.

Although the TPCs record detailed information on the tracks produced by ionising particles this information still needs to be processed by reconstruction algorithms. As will be discussed in the next two chapters the default ND280 reconstruction software was not designed with interactions within the TPCs in mind, motivating the development of new pattern recognition and reconstruction.




Figure 4.12: Hits left by a particles from the FGD passing through the TPCs, from [159].

Chapter 5

T2K Software

The T2K software, particularly ND280 simulation and reconstruction, is mostly written in C++ and based around the ROOT framework[160]. The Geant4[161] package for simulating the propagation of particles through matter is used extensively in Monte Carlo simulation. The related but older GEANT-3 package is used in more limited cases for beam simulation[162].

In this chapter we cover the steps involved in simulating (in the case of Monte Carlo), calibrating and reconstructing events. We do not discuss the TREx package for reconstructing TPC events — this is covered in detail in Chapter 6.

5.1 Simulating neutrino interactions

To simulate neutrino interactions accurate predictions of neutrino flux are required. Their derivation is discussed in detail in [163]. Hadronic interactions of protons with the baffle and target are simulated by the FLUKA package[164], tuned with both in situ measurements of the proton beam and with hadron production data primarily from the dedicated NA61/SHINE experiment[165].

From here the JNUBEAM package propagates particles through a simulated baffle and target, horns and magnetic fields, decay volume, beam dump and monitors. This simulation uses GEANT-3 and provides a prediction for the neutrino beam's flux, which is passed to neutrino interaction generators. For ND280 analyses we are concerned with fluxes at the ND280 detector and in the surrounding sand.

The JNUBEAM flux is passed to one of two primary generators for neutrino interactions, NEUT and GENIE, from which the rates and outgoing particle kinematics of neutrino interactions are calculated. Both follow a generally accepted

three pronged approach:

- Use a model for nucleon kinematics to determine the position and momentum of the target nucleon.
- Select a channel for the interaction and determine interaction products and their kinematics.
- Propagate these products out of the nucleus while accounting for final state interactions.

In theory the two generators are simulating the same processes. In practice each comes with its own specific models and implementations thereof which lead to different predictions, particularly in variables and areas of phase space which have not been precisely measured by experiment.

5.1.1 NEUT

NEUT[166] is one of the two interaction generators used by the T2K experiment. It was originally written for the Kamiokande experiment mentioned in Chapter 1 and was subsequently adopted by Super-Kamiokande and the accelerator experiments for which it served as a far detector. This analysis used data simulated with NEUT version 5.3.2.

In the present version of NEUT nuclear physics are modelled in one of two ways depending on the target nucleus. For C, O and Fe the simulation employs spectral functions as described by Benhar et al[167]. These provide momenta and removal energies for nucleons. All other targets are represented a relativistic Fermi gas model. The nominal model takes no account of correlations between nucleons, though an implementation of the Nieves model[168] for this purpose is available in the code.

For quasi-elastic and elastic interactions (charged and neutral current) the Llewellyn Smith model is used with dipole $F_A(Q^2)$ form as described in subsection 2.3.2. In its most recent NEUT implementation an axial mass of $M_A = 1.21$ GeV is used.

Resonant interactions are described using the Rein-Sehgal model noted in subsection 2.3.4 with a resonant axial mass of $M_A^{\text{res}} = 1.21$ GeV. Coherent pion production is simulated using the (unrelated) Rein-Sehgal model mentioned in subsection 2.3.5 with partially conserved axial currents. Deep inelastic scattering channels use the Bodek and Yang model (subsection 2.3.6) with PYTHIA/JETSET above 2 GeV and internal NEUT code at lower energies.

Final state interactions for all channels are simulated using a cascade model where pion scattering, charge exchange and absorption are considered. In steps determined by mean free path interaction products are propagated through the nucleus with the probability of an interaction considered each step. If such an interaction occurs the resulting particles are also propagated through the nucleus and their own potential interactions considered. This proceeds until no further interaction products remain within the nucleus.

5.1.2 GENIE

GENIE[169][170] was developed to be a general generator for neutrino interactions particularly in the range of a few GeV. This contrasts with NEUT, which was developed specifically for Kamiokande and updated for later experiments such as T2K. This analysis uses data simulated with GENIE version 2.8.0.

Nuclear physics are simulated with a relativistic Fermi gas model, with extensions as described by Bodek and Ritchie[171] to describe short range correlations between nuclei. The models used for individual interactions are analogous to those used by NEUT but with differing implementations.

Quasi-elastic scattering once again uses the Llewellyn-Smith model. The axial mass parameter is lower than NEUT's, at 0.99 GeV. Elastic scattering uses the model of Ahrens et al with axial vector form factor as in Equation 2.40.

Resonant production of baryons is represented with the Rein-Sehgal[65] model, as with NEUT, but with resonant axial mass $M_A^{\text{res}} = 1.12$ GeV. GENIE ignores interference between resonances and neglects lepton masses in calculating differential cross section. Coherent scattering is also modelled with the Rein-Sehgal model used by NEUT[66]. The cross section for deep inelastic scattering is calculated using an effective leading order model with Bodek and Yang's modifications[70]. For hadronisation GENIE used the AGKY model, in which PYTHIA is utilised above at energies above a few GeV whilst an empirical model is used at lower energies (a smooth transition is enforced between the two regions).

Final state interactions are simulated with a simplified, empirical model tuned with experimental data. The total cross section for each rescattering process for each particle is calculated as a function of energy and these cross sections are used to evaluate probabilities of a given final state leading to a given selection of particles exiting the nucleus.

5.1.3 Disagreements between generators

As shown in Figure 5.1 these generators do not produce identical kinematics. GENIE in general predicts interactions to produce a greater number of secondary particles, each with a lower share of the interaction’s total energy. NEUT on the other hand favours fewer secondary particles with a higher share of momentum. In addition, NEUT predicts a sharp drop-off in proton momentum below the Fermi-momentum at around 250 MeV whilst GENIE’s highest proton momentum increases steadily all the way down to around 100 MeV.

Disagreement in these distributions is unsurprising since there is a distinct lack of experimental data on low momentum protons and their multiplicity with which to tune generators. Our gas TPC, with its capability to measure protons down to very low energies and resolve vertices in sufficient detail for precise multiplicity measurements, affords us an excellent opportunity to provide such useful information.

5.2 Interactions outside the ND280

Cosmic rays are almost universally negligible in analyses as a result of beam timing cuts. They are used mainly for understanding detector performance and the computation of systematic errors.

Their simulation begins with the Fortran based CORSIKA package for show-

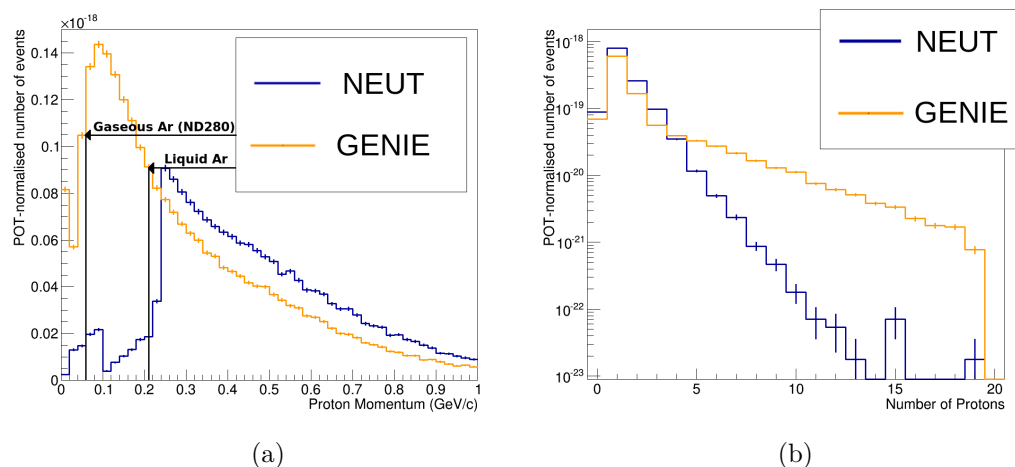


Figure 5.1: Differences in highest proton momentum (Figure 5.1a) and proton multiplicity (Figure 5.1b) between GENIE and NEUT. Also marked are the predicted minimum resolvable energy for our gas TPC and a liquid TPC. From [172].

ers produced by high energy cosmic particles. This software can provide a flux specific to the ND280 pit. From here a Geant4 based simulation propagates the flux, accounting for any interactions between muons and the pit's concrete walls or surrounding sand. From the edges of the ND280 detector the general ND280 simulation described in Section 5.3 takes over. An additional step after this simulates the responses of cosmic triggers.

Also simulated are sand muons. These result from interactions of neutrinos in the sand surrounding the ND280 detector. The simulation procedure uses NEUT as a generator and is identical to generation of interactions inside the detector but with the addition of an intermediate step for propagating interaction products from the sand to the ND280 detector. At this point general ND280 simulation takes over.

5.3 ND280 simulation

The Monte Carlo generators simulate interactions for individual neutrinos resulting in a selection of interaction products and their initial kinematics. For propagating these products through the ND280 detector Geant4 is used.

This simulation package covers any expected interactions and decays and, crucially for the ND280, the interactions of charged particles with electric and magnetic fields. Our simulated ND280 geometry covers all active detector modules plus dead regions such as our large surrounding magnets. Positions of individual modules can be modified in light of new measurements using alignment constants.

During the gas interaction analysis a bug was discovered in the ranges of low momentum protons. The default photoabsorption ionization model provided significantly lower than expected energy depositions when enabled for protons. Although this is fixed in later versions, our analysis requires corrections to the overly high proton ranges in Monte Carlo.

5.4 ND280 electronics simulation

Our electronics simulation covers the response of detector electronics to through-going particles. It includes production and transportation of photons in scintillator bars (for the FGDs, ECals, P0D and SMRD) and the production and drift of ionisation electrons (for TPCs). Electronics readout is also simulated. The package then handles electronics chains from readout up to the point of producing digits which can be calibrated in the same way as real data.

The first iteration of the gas interaction analysis uncovered electronics effects

in data which are not accounted for by simulation. These involve the creation of additional unphysical noise hits in channels surrounding large energy depositions. Because of their appearance the problematic tracks surrounded by such noise were dubbed ‘hairy tracks’. The way they are handled by **TREx** reconstruction and resulting systematic uncertainties are described in subsection 6.3.4 and subsection 8.6.3.

5.5 ND280 calibration

Calibration is the first step to apply to real data as well as Monte Carlo. It encapsulates two jobs. One is unpacking input data and saving it in a form useful to reconstruction. Both Monte Carlo and real data digits are unpacked. The other involves applying corrections to data accounting for both detector and electronics effects. In the case of Monte Carlo most of the individual corrections are irrelevant and thus skipped.

Tuning for the data calibration is driven by on site measurements including cosmic triggers, pedestal (measurements in the absence of signal) and charge injection data. TPC calibration uses a magnetic field map produced before their installation. Drift parameters are also regularly checked using measurements from a laser system which produces photoelectrons at known positions.

5.6 ND280 reconstruction

ND280 reconstruction[173] does the job of converting complicated hit level information into objects such as tracks and vertices which are of interest to analysers. Each module has its own reconstruction algorithms, though most at some stage use information from TPC level reconstruction.

Much of the reconstruction involves combining and refitting objects within or between detectors. This is achieved using a Kalman filter[†] for the following cases:

- Merging a track or seed and new hits or clusters with discrete position.
- Merging two TPC tracks, each with fitted position and momentum.
- Merging a track and a shower object in the POD or ECal.

The Kalman filter procedure provides a fast way of performing incremental matching, where a track object is matched with a single hit, refitted using the Kalman filter, then matched with another hit and so on.

[†]In general a Kalman filter is a tool for incrementally estimating the value of a quantity by combining previous estimates and new information, both of which have some known uncertainty[174].

5.6.1 TPC reconstruction

The previous TPC reconstruction used a pattern recognition algorithm which assumed through-going tracks in the forwards (z) direction. Consequently it was completely unsuitable for gas interaction analyses and performed poorly on topologies such as backwards curving tracks.

The algorithms have been replaced by **TREx**. Pattern recognition was completely replaced and fitting, particle identification (PID) and absolute time (t_0) determination have undergone general improvements in both structure and content. As a result of this extensive development **TREx** is discussed in great detail in Chapter 6.

5.6.2 FGDs

FGD reconstruction occurs immediately after TPC reconstruction. Hits are either matched with TPC tracks or reconstructed as isolated FGD tracks. After this PID and track time are calculated.

TPC matching used a TPC track as the starting point for creating an FGD object. FGD hits are added and fitted incrementally using a Kalman filter, in order of matching χ^2 with the extrapolated TPC track, until no hits remain with a χ^2 below threshold. Matching proceeds sequentially in the order TPC1 \rightarrow FGD1, TPC2 \rightarrow FGD1, TPC2 \rightarrow FGD2, TPC3 \rightarrow FGD2.

Stand alone reconstruction produces objects ending in an FGD and not continuing into a TPC. Previously unmatched hits undergo cellular automaton based pattern recognition (with the **sbcats** algorithm[†]) and are reconstructed within the FGD.

5.6.3 Tracker

Tracker reconstruction encapsulates both TPC and FGD reconstruction and concatenates tracks from separate modules to form tracker objects.

FGD-TPC matching creates tracks spanning a single TPC with hits in one or more FGD. After this step tracks in different TPCs are matched based on whether or not their matching χ^2 is below a threshold. Each TPC track end is matched to a maximum of one other track in an adjacent TPC. Once matching is done the tracks are refitted with a Kalman filter to update their kinematics.

[†]**sbcats** is a cellular automaton originally developed for pattern recognition in the SciBar detector used for K2K and SciBooNE. A three dimensional version was subsequently used for TPC reconstruction before its replacement by **TREx**.

5.6.4 P0D

P0D reconstruction involves both a track and a shower reconstruction algorithm, in that order.

Track reconstruction fits the P0D 's two two dimensional views ($x-z$ and $y-z$) separately and merges the two outputs where possible. Vertexing is performed after initial tracks have been identified, with vertices identified between each pair of tracks. The resultant vertices are subsequently merged into a smaller number of vertices where possible. A PID step separates minimum ionising particles from electromagnetic showers and tags tracks of awkward length and angle as of indeterminate identity.

Shower reconstruction proceeds using the results of track reconstruction. Vertices are used to construct showers from particles tagged as electromagnetic or indeterminate identity plus unused hits. The reconstruction assumes that all hits from a shower lie in a cone as seen by the vertex and adds hits based on this.

5.6.5 ECals

During my time at T2K I worked on ECal software, particularly the EM energy reconstruction and PID. These two algorithms are thus covered in more detail below.

As with the P0D and FGD ECal hits are divided between two orthogonal views[†]. Each two dimensional view is fit separately. A clustering algorithm seeds clusters from the highest charge hits and tries to add neighbouring hits. This runs recursively on each added hit until no more can be grouped. Clusters are passed onto the combining stage where small clusters (two hits) and unmatched hits are added into larger clusters to which they are spatially close.

Two dimensional clusters are then matched between views to form three dimensional objects. This matching is likelihood based, with a likelihood constructed from the ratio of cluster charges and the difference between starting layers. Tracker objects can be used as seeds for the matching. Since an ECal can contain many objects only the matches with the best likelihoods (above a threshold) are accepted. After this stage cluster position is recalculated for all three dimensional objects.

At this point the EM energy measurement described in Section 5.6.5 is made.

Three dimensional clusters are fitted as both tracks and showers. Shower fitting utilises three dimensional principal component analysis while track fitting uses a three dimensional linear fit on points derived from combined average hit

[†]Hits are in $x-z$ and $y-z$ views for the downstream modules, $y-z$ and $x-y$ for top and bottom modules and $x-z$ and $x-y$ for side modules.

positions in each layer.

After three dimensional fitting this point the particle identification procedure covered in Section 5.6.5 is run.

EM energy

Three dimensional clusters are passed to an energy reconstruction algorithm. The purpose of this is to determine the amount of energy deposited in an ECal assuming the cluster belongs to an electromagnetic shower. Uncontained minimum ionising particles have their energy underestimated by the algorithm.

Three variables are fed into an energy fit. They are

- **Charge sum**; the total charge of the cluster.
- **Charge RMS**; the spread of the cluster's charge distribution.
- **Charge skew**; the skewness of the cluster's charge distribution.

Of these charge sum is the most correlated with total energy. At a given energy the PDF for each variable is assumed to follow a Gaussian distribution. The mean and standard deviation for each variable's PDF and correlations with other variables are extracted at each energy. Fitting these with splines provides provides the PDFs as a continuous function of energy.

The reconstructed cluster's observed charge sum, RMS and skew variables used in conjunction with these PDFs to produce a combined likelihood which varies as a function of energy. The energy providing maximum likelihood is taken to be the object's uncorrected EM energy.

An additional tuning is applied both to correct for differences between ECal submodules and to scale the most probable value returned by the fit to the statistical mean instead. This is a simple scaling factor which varies as a function of ECal submodule and uncorrected energy. Final EM energy is the initial, uncorrected energy multiplied by this factor.

Particle identification

ECal particle identification[175] uses several variables to construct log likelihood ratios for each cluster between two hypotheses. These are

- **MIP-EM**; a low value for which favours the minimum ionizing particle hypothesis and a high value the electromagnetic shower hypothesis.

- **MIP-pion**; a low value for which favours the minimum ionizing particle hypothesis and a high value the showering pion hypothesis.
- **EM-HIP**; a low value for which favours the electromagnetic shower hypothesis and a high value the proton hypothesis.

These discriminators are each built from a number of low level variables. These are chosen for their systematic consistency between data and simulation and their low correlation with each other. The variables are

- **Circularity**; a measure of the extent to which the cluster has a large width in both views. This is near 1 for track-like objects and small for shower-like objects.
- **Truncated max ratio**; a measure of the extent to which the cluster's highest charge layers feature much more charge than it's lowest charge layers. This is formed by taking the ratio of the highest to lowest charge layer after removing the top and bottom 20% of layers.
- **Charge RMS** as used for EM energy; a measure of the spread of the cluster's charge distribution. This is higher for showering than non showering particles.
- **Front-back ratio**; a measure of the extent to which the track features higher energy deposition at the end of the track. This is substantially higher for stopping than non stopping particles.

PDFs are produced for each variable for each particle. For reconstructed object the measured values of these variables are used to compute total likelihoods for each hypothesis (MIP, EM, pion and HIP), the ratios of which give the log likelihood discriminators.

5.6.6 SMRD

Reconstruction for the ND280 SMRD takes place after the reconstruction for all other modules. It involves first evaluating hit position along SMRD bars and then combining hits to create tracks. There are two algorithms employed. One uses matching between tracks from other modules and SMRD hits. The other acts only on isolated SMRD hits.

The first algorithm uses a Kalman filter to match calibrated hits within the SMRD to objects within the other detector modules so long as matching χ^2 is below

a threshold. Once all SMRD hits which can be matched have been they are fitted together to form an SMRD track.

The second algorithm is performed on SMRD hits which are not matched to objects from other modules. A pattern recognition step groups hits depending on their location, time and position. This is followed by two dimensional linear fits which are combined to form 3D tracks for global reconstruction.

5.6.7 Global reconstruction

Global reconstruction involves stitching the results of individual reconstruction algorithms into final reconstructed objects. After an initial hit preparation stage the objects are combined together.

The process begins by matching objects from the tracker to those in adjacent modules by requiring a matching χ^2 below a threshold and small time differences between objects. Matched objects are combined and refit using a Kalman filter one pair at a time. This proceeds until no more matches are found. After this initial matching the process repeats with remaining unmatched POD objects and finally with remaining unmatched ECal objects.

5.7 Analysis

The last step in the ND280 software chain is to strip out useful information from reconstructed events and format them in a lightweight ‘analysis’ format. This results in files of much smaller size than the outputs of simulation, calibration and reconstruction which makes large scale processing feasible.

Our analysis, like many others, uses the HighLAND package to run our selection on these analysis files. This high level framework handles analysis using reconstructed objects and variables as well as standardising corrections, selections, and systematic uncertainty propagation.

Chapter 6

TREx Reconstruction

As mentioned in Chapter 5 the previous reconstruction for ND280 TPCs was insufficient for studies of interactions *inside* the gas volume. The old algorithms assumed forwards going tracks originating upstream of the TPC. There was no concept of a TPC vertex and the reconstruction had difficulty correctly reproducing backwards going tracks and other awkward topologies.

TREx was designed primarily with gas interactions in mind but also to overcome general limitations of the old TPC reconstruction. The underlying TREx philosophy was to produce an algorithm which could

- Identify both vertices (from primary interactions) and secondary interactions (for example δ -rays branching off from a main track).
- Be as isotropic as possible, producing tracks regardless of the angle at which they emerged from a vertex.
- Accommodate arbitrary shapes such as spiralling and backwards curving tracks[†].

A well known path finding algorithm, the A* algorithm[176], forms the foundation of TREx pattern recognition, allowing junctions and track ends to be connected as well as facilitating many other aspects of the algorithm. Given successful pattern recognition, paths can be fitted to obtain kinematic information and PID can be evaluated using $\frac{dE}{dx}$.

6.1 Framework

The basic structure employed by TREx is illustrated in Figure 6.1. Each event is divided into multiple patterns. These are defined as complete groups of hits which

[†]These are likely to be produced by low energy electrons

can be connected through the A* path finding described in subsection 6.3.1.

Hits are first fed through hit preparation and then handed to pattern recognition. At this stage in the algorithm only hit pad objects are available. Pattern recognition breaks the entire event into patterns of hits between which A* connections can be made. Within each pattern it assigns hits to either clusters (as part of paths) or to junctions.

These path and junction objects are passed through several stages of seeding, merging, fitting and PID determination. The full scope of TREx is detailed in the rest of this chapter.

6.1.1 TREx base classes

TREx was designed to be as agnostic as possible to the underlying physics, in contrast with the previous reconstruction (which assumed a specific topology and worked from there). For this reason, naming conventions have been chosen which do not imply assumptions about the true identities of reconstructed objects.

‘Track’-like objects are referred to as paths, reflecting the fact that this may represent a full true track or just some segment of it. ‘Vertex’-like objects are referred to as junctions, reflecting uncertainty over whether they are primary vertices, secondary interactions, or some artefact of an unusual topology (for example the rare case of two unconnected tracks passing through each other in the TPC).

The connections of TREx data classes to one another are illustrated in Figure 6.1. The base classes are:

- **Hit pads**; basic hit objects containing charge collected over a given time window and information such as the number of peaks and saturated bins.
- **HV clusters**; clusters of hit pads in the same micromegas row or column (for horizontal or vertical clusters respectively), arranged in such a way to facilitate likelihood fitting later in the algorithm.
- **Paths**; ordered lists of HV clusters, with the possibility of both horizontal and vertical clusters in the same path.
- **Junctions**; objects where two or more paths meet containing any hit pads which cannot reliably be assigned to a single path.
- **Patterns**; high level objects containing either connected paths and junctions or a single path.

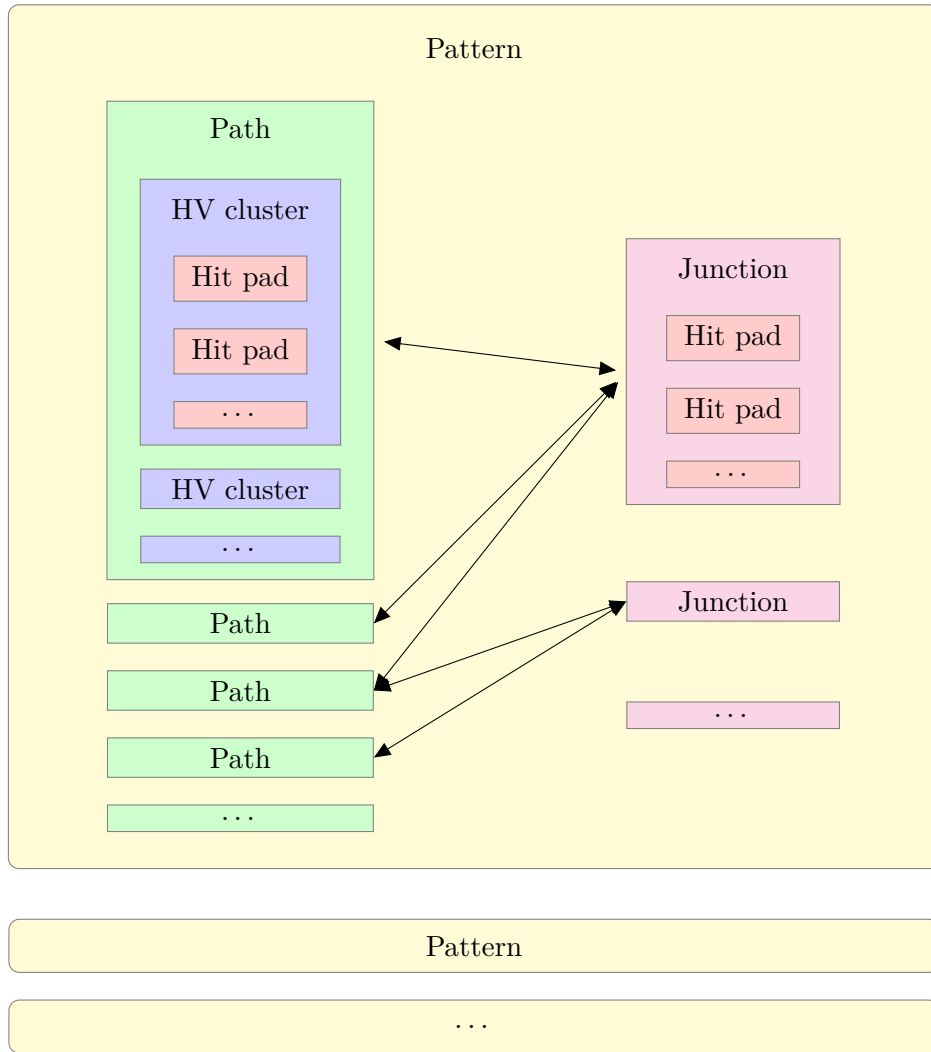


Figure 6.1: Data structure used for TReX reconstruction. Objects which are constituents of others are illustrated (hit pads inside HV clusters inside paths et c.). Arrows indicate objects which are aware of each other (paths are aware of which junctions they connect to and vice versa).

6.2 Hit preparation

The basic hit object from calibration described in Section 5.5 are ‘waveforms’ representing charge collected as a function of time. Each provides a time window for a single micromegas pad within which charge above the pedestal threshold was collected. At the hit preparation stage these waveforms are examined and, if possible, broken into smaller objects.

The procedure for analysing and breaking individual waveforms is illustrated

in Figure 6.2. They are first examined in order to find peaks by looking for local maxima in charge. In general a peak is defined as a bin with higher charge than the one either side, though if charge is low (and thus relative noise high — a 20 ADC threshold is used for this) five bins of lower charge either side are required. If a peak consists of multiple bins at the exact same charge it is marked as saturated, in which case its maximum charge and time are extrapolated from the saturated bins.

Ideally one hit pad object is created per peak. If two or more peaks are present the algorithm attempts to break the waveform into new hit pads. Peaks which are too close together are not separated; instead a single hit pad object uses the peak with highest charge and records the number of inseparable peaks it contains.

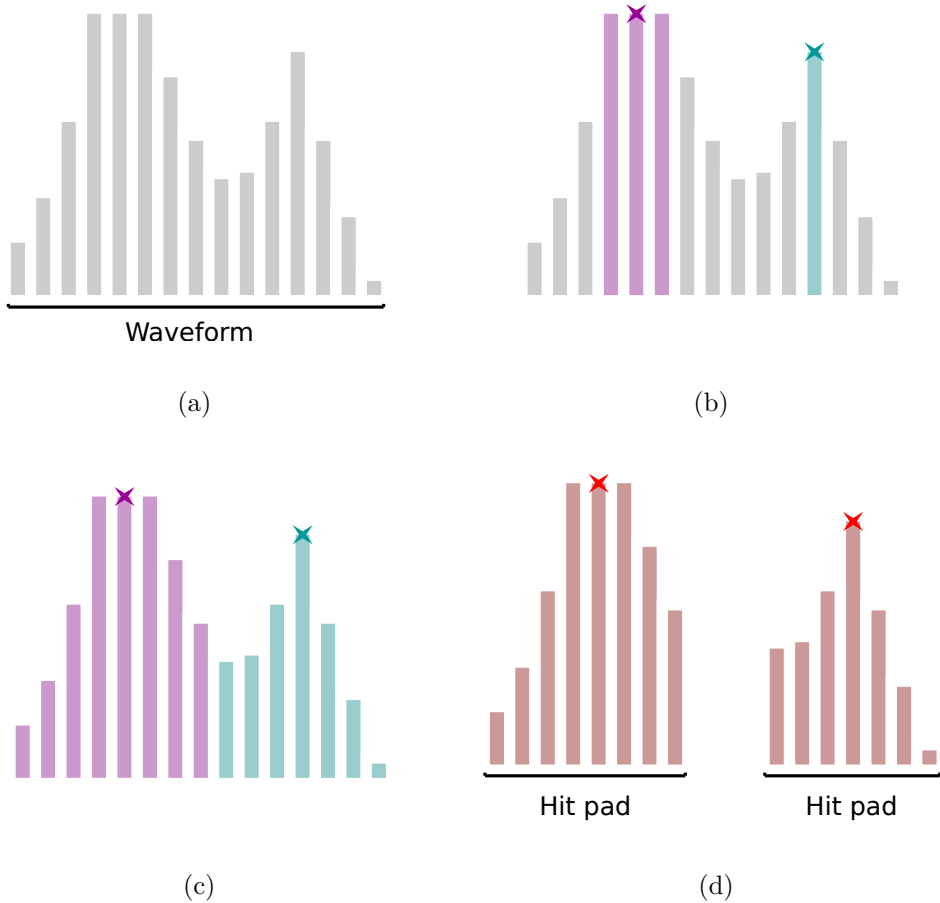


Figure 6.2: Illustration of TREx hit preparation. Blue and green stars represent separate peaks. Figure 6.2a shows the input waveform with two peaks, one saturated. Figure 6.2b demonstrates identification of the two peaks while Figure 6.2c shows bins assigned to each. Figure 6.2d shows the two output hit pads with time and charge of peaks saved.

6.3 Pattern recognition

TREx pattern recognition is completely unrelated to previous TPC software. The basic procedure is outlined in Figure 6.3.

Note that pattern recognition uses custom classes designed for efficiency and utility and distinct from TREx base classes. In practice there is a large degree of

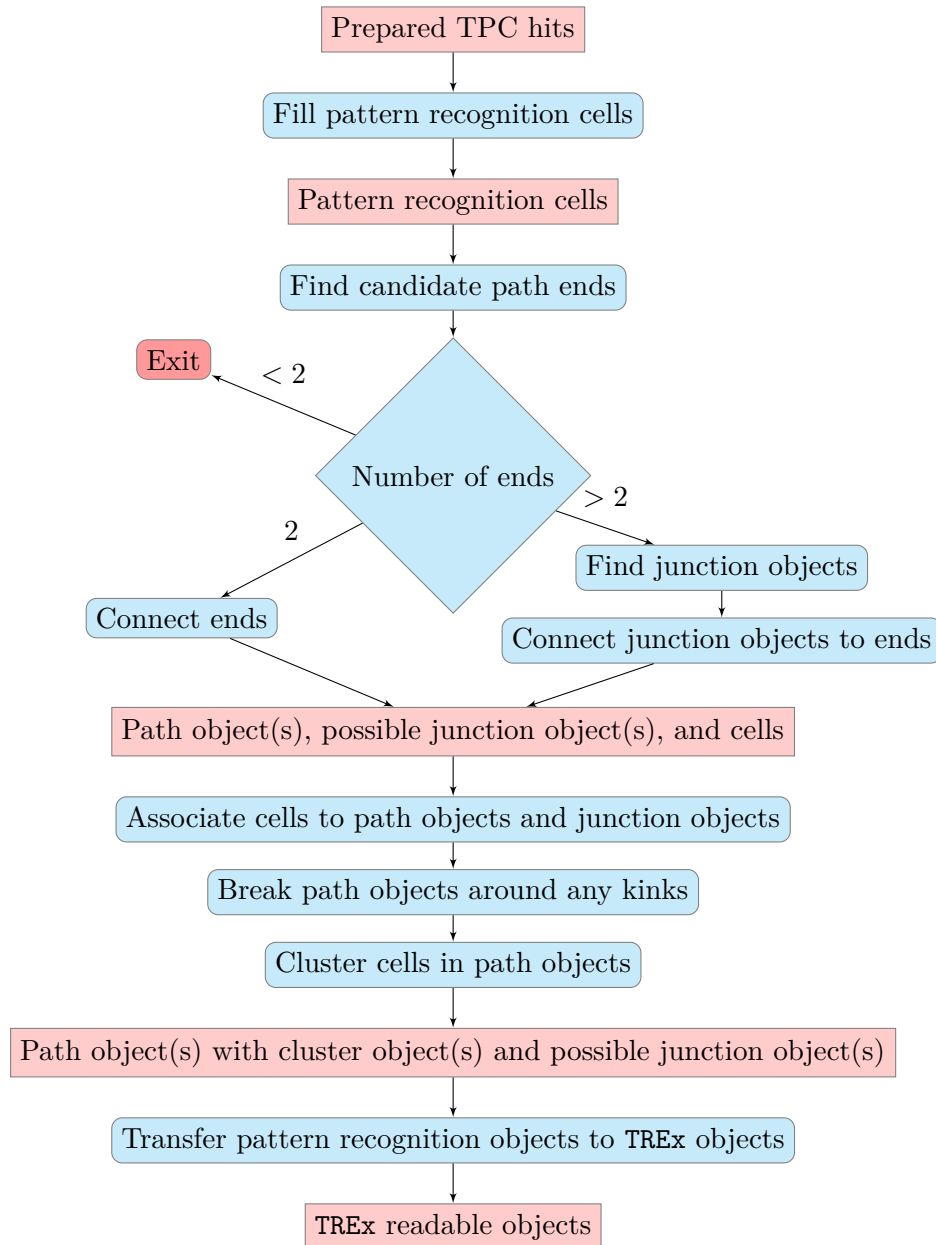


Figure 6.3: Overview of the main TREx pattern recognition steps. A number of more minor intermediate steps have been skipped for the purpose of this illustration.

overlap between pattern recognition classes and the **TREx** base classes discussed in subsection 6.1.1. This leads to potential for confusion over nomenclature which is clarified here.

The most basic pattern recognition element is a single unit of TPC volume. This differs from a **TREx** hit in that it is a cell at some position in a regular three dimensional grid. Each represents an area of physical space in x , y and z and may contain one or more hit pad. For clarity and brevity these voxels are referred to as cells.

Cluster, path and junction objects are set out similarly to the base classes. For pattern recognition cluster and junction objects hold cells. Path objects hold cluster objects. Pattern objects are also analogous to their **TREx** base class counterparts, containing path and junction objects.

In contrast with with **TREx** base classes each path object contains three additional groups of cells. Two represent groups of cells at the path object’s front and back which serve as junction candidates. The third represents all cells which can be associated with the path object — it is the set of hits used for building cluster objects for this particular path object.

To resolve ambiguity in this section the pattern recognition classes will be referred to as cells, cluster objects, path objects, junction objects and pattern objects. Their equivalents among **TREx** base classes will be referred to as **TREx** hit pads, **TREx** HV clusters, **TREx** paths, **TREx** junctions and **TREx** patterns. Distance in pattern recognition space is measured simply in number of cell lengths, hereafter referred to as ‘units’ (the distance between adjacent cells is 1 unit).

6.3.1 The A* algorithm

A* is a path finding algorithm for efficiently finding the minimum cost path through a graph[176]. This scenario has been covered by Dijkstra’s algorithm[177] but A* massively improves efficiency by employing a *heuristic* term to ensure that those paths most likely to lead towards the end node[†] are searched first. The algorithm terminates once the end node is reached regardless of how many potential paths remain unchecked.

The **TREx** specific implementation maintains a list of accessible nodes (the ‘free list’) with connection costs and heuristic costs saved. All nodes are initially flagged as ‘open’ (unchecked). The connection cost is the *actual* cost of connecting from the start node via a chain of intermediate nodes, computed by summing individual costs for each individual connection (for example, if the connection $\mathbf{A} \rightarrow \mathbf{B}$

[†]‘Node’ here is taken to have its simplest graph theory definition.

has a cost of 1 and $\mathbf{B} \rightarrow \mathbf{C}$ a cost of 4 then $\mathbf{A} \rightarrow \mathbf{C}$ is possible with a cost of 5). The heuristic cost is the *estimated* cost of connecting to the end node. Connection costs can be updated if more efficient routes are found whereas heuristic costs are constant for each node.

As well as costs, each node (apart from the first) stores its ‘parent’, or the node which connected directly to it. This allows the route to be traced back once the end node has been found, reconstructing the path as an ordered list of nodes.

Path finding is done through the following procedure:

- Ensure the free list is empty and add the start node with 0 connection cost.
- Set all other nodes to an arbitrarily high connection cost with undefined parents and flag them as open.
- Repeat the following steps until the current node is the same as the end node:
 - Find the node in the free list with the lowest total cost (connection cost plus heuristic cost) and set this to be the current node.
 - Examine all open nodes connected to current node.
 - For each calculate connection cost via the current node (sum the connection cost of the current node and the cost from the current node to the new node).
 - If this cost is lower than the new node’s previous connection cost:
 - * Set new node’s connection cost to this new cost.
 - * Set the new node’s parent to the current node.
 - * Add the new node to the free list if it isn’t already there.
 - Set the current node’s ‘open’ flag to **false** and remove it from the free list.
- At this point the end node has been reached.
- Form a new, empty list representing each node along the path.
- Starting with the end node work backwards through each node’s parent, adding each to the list, until the start node is reached.

In the **TREx** implementation nodes are only checked once — once connections from a particular node have been added to the list of accessible nodes the node’s ‘open’ flag is set to **false** so as not to check it again. This is to save processing time.

More rigorous A* implementations can re-check a node if a lower cost connection to it is found.

Figure 6.4 illustrates the algorithm in action on a sample event. A* is robust enough to efficiently find paths between start and end nodes through graphs of nodes with arbitrary shape so long as the heuristic term can be given a meaningful form which generally gets smaller as the end node is neared.

A* is a general algorithm so it is up to TREx to define specifically what is defined as a node, a valid connection and a connection or heuristic cost. The specifics of converting TPC hits to pattern recognition cells are described in subsection 6.3.4; these cells act as nodes.

In theory a connection could be formed between any two cells. In practice attempting this would result in a punishingly slow reconstruction with the added disadvantage, shown in Figure 6.5, of assigning unconnected but spatially close tracks to the same pattern. The opposite approach would be to only allow connections between adjacent cells. This is also problematic since valid tracks often feature

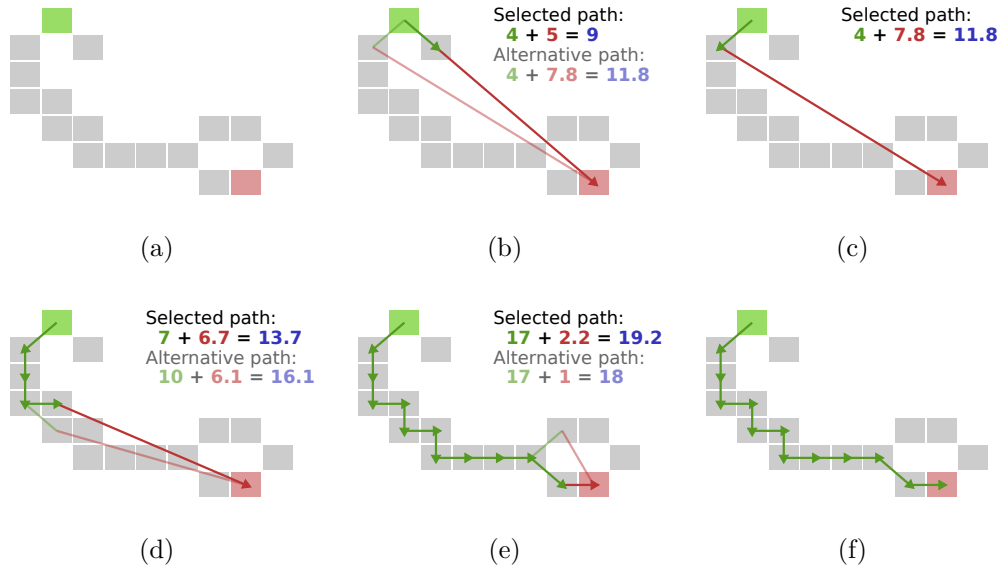


Figure 6.4: Illustration of the A* algorithm handling various features in a difficult event, with connections possible between adjacent elements at cost 1 and diagonal elements at cost 4. The green node represents the start, red the end and grey intermediate cells. Figure 6.4a shows a start and an end cells and Figure 6.4f the final path between them. Figure 6.4b demonstrates the algorithm attempting a correction that ends up being dropped due to a more efficient route in Figure 6.4c. In Figure 6.4d shows the effect of setting the cost of a diagonal connection $2\times$ as high as two adjacent connections in directing the algorithm away from cutting corners. Figure 6.4e shows the heuristic term directing the path towards the end node.




Figure 6.5: Illustration of the issues with different connect distances in the TREx implementation of A^* . Figure 6.5a demonstrates a connection being missed due to too low a connect distance whilst Figure 6.5b shows connections being found between unconnected objects due to too high a connect distance.

small gaps and would thus be broken by this condition, as also shown in Figure 6.5. TREx finds a middle ground by allowing connections between hits within 2 units of each other in the y and z directions and 6 units in the x direction. Larger connection distances are allowed for hits on the edges of horizontal micromegas gaps to compensate for the extra uninstrumented space between pads.

Connection costs have been defined to encourage connections, where possible, between adjacent cells before attempting connections over larger distances. Strictly they are defined by

$$\text{Cost}(\mathbf{A} \rightarrow \mathbf{B}) = \left(\left(\frac{x_{\mathbf{A}} - x_{\mathbf{B}}}{s_x} \right)^2 + \left(\frac{y_{\mathbf{A}} - y_{\mathbf{B}}}{s_y} \right)^2 + \left(\frac{z_{\mathbf{A}} - z_{\mathbf{B}}}{s_z} \right)^2 \right)^2, \quad (6.1)$$

where x_i , y_i and z_i correspond to integer cell position in x , y and z respectively and s_x , s_y and s_z are characteristic connection distances set to 3, 1 and 1 respectively (s_x is higher to reflect the larger gaps between TPC hits in the x direction). As can be seen in this equation, position differences are effectively placed to the power of four which favours several short connections over a smaller number of long connections[†].

[†]If \mathbf{A} , \mathbf{B} and \mathbf{C} are at positions $(1, 1, 1)$, $(1, 1, 2)$, and $(1, 1, 3)$ respectively then $\mathbf{A} \rightarrow \mathbf{B} \rightarrow \mathbf{C}$ will cost $1^1 + 1^1 = 2$ whilst $\mathbf{A} \rightarrow \mathbf{C}$ will cost $2^4 = 16$ and only be chosen in the absence of viable alternatives. Diagonal connections are similarly penalised though to not as great an extent.

Heuristic costs are defined by the distance from a given hit to the end node. They are given by

$$\text{Cost}(\mathbf{A} \rightarrow \text{end}) = f_{\text{heuristic}} \sqrt{(x_{\mathbf{A}} - x_{\text{end}})^2 + (y_{\mathbf{A}} - y_{\text{end}})^2 + (z_{\mathbf{A}} - z_{\text{end}})^2}, \quad (6.2)$$

where $f_{\text{heuristic}}$ is a weighting given to the heuristic cost relative to connection cost. Higher weighting reduces the average iterations required for the algorithm to reach the end point but also reduces the chance of an optimal route being found, particularly for complicated paths. Lower weighting increases average iterations but also makes the returning of a suboptimal path less likely. A value of 1.2 was chosen for `TREx`.

Path finding forms the basis for `TREx` pattern recognition so it is important that it can be performed efficiently and accurately. `A*` accomplishes this.

6.3.2 Cell definition

`TREx` pattern recognition is fairly agnostic to underlying detector geometry and physics. It operates on a three dimensional grid of cells based on the positions of TPC hits. Because of this its first step must be to convert ND280 TPC hits into pattern recognition cells at discrete positions.

The cell objects contain minimal information needed for pattern recognition processing. The three dimensional position of a hit is represented by integer indices of a cell, giving its position in a three dimensional grid. Additional flags are saved for information used by pattern recognition, for example if the cell represents hit pads the edge of a micromegas unit. Flags are also set for low charge hits and hits on highly occupied ASIC units as discussed in subsection 6.3.3. Finally each cell saves a list of contributing hit pads so that `TREx` classes can be easily filled and returned at the end of pattern recognition.

The position indices for each cell are strictly defined in the y and z directions, corresponding unambiguously to the discrete micromegas pads (see subsection 4.4.6). The only exception is horizontal pad gaps, which are defined as empty rows 1 cell wide.

The x dimension comes from hit time. Since this is not a discrete quantity some binning is required. The width of each time bin corresponds to the distance between micromegas pads in z divided by TPC drift speed giving roughly equal spatial resolution between cells in x , y and z . In addition hits are divided into two sets of bins — one for each readout plane. Charge collected at a given time will

thus be put into one of two cells depending on which side of the central cathode is was deposited on.

6.3.3 Charge and occupancy flags

One major issue identified during the gas interaction proto-analysis was the behaviour of TPC electronics in the presence of highly ionising particles (see Section 5.4). A number of non-physical effects were seen which played havoc with TREx pattern recognition. To remedy this two flags were developed to alter the behaviour of pattern recognition in their presence.

The effects take two broad forms, illustrated in Figure 6.6. One involves a physical track with many surrounding low charge hits. These can form many structures, for example a shadow reflecting the shape of the main track at delayed time, or repeating unphysical patterns in y and z . Another form of problem involves large blocks of hits, usually with a sharp rectangular shape reflecting the underlying




Figure 6.6: Illustration of TREx flags for non-physical effects. An event with many spurious low charge hits is shown in Figure 6.6a and these hits are separated from the others by their low (< 30 ADC) charge as shown by red cells in Figure 6.6b. Another event with large blocks of charge is shown in Figure 6.6c. These blocks are identified as shown by purple cells in Figure 6.6d.

electronics. Both of these are now addressed by pattern recognition.

Cells identified as being probably non-physical have a flag set to identify them as either low charge or part of a high occupancy region. Low charge hits are simply identified as those with charge below 30 ADC whilst hits in areas of high occupancy are identified by looking for ASIC modules with hits on more than 50 individual pads, or half ASICs with hits on more than 30.

Each of these flags was tuned to avoid rejecting muon tracks. In the case of the charge cut 30 ADC was identified as the cut value which rejected as many hairy hits as possible without damaging our efficiency in reconstructing minimum ionising particles. In the case of the ASIC occupancy cuts minimum ionising particles trivially pass. The only non-hairy hits which may be flagged are those near high multiplicity vertices and since neither flag affects vertex reconstruction this does not damage our efficiency.

6.3.4 Pattern object preparation

The first step in the pattern recognition proper is to place every hit into a pattern. A TREx pattern is defined as a complete selection of hit pads which can be connected to one another through the A* algorithm. As shown in Figure 6.7 two unconnected tracks in the same TPC will be grouped into different patterns as will hits in different TPCs and hits either side of the central cathode or a vertical micromegas gap.

The grouping is performed using an identical criterion to that for forming A* connections. This ensures:

- There exists an A* path between any two cells in the same pattern object.
- Any two cells connectable through A* are in the same pattern object.

To form the pattern objects a recursive algorithm is employed. A ‘free list’ is produced of all cells in the event. The following steps are iterated through until the free list is emptied:

- Create a new group and select a random cell from the free list as the current cell.
 - Remove the current cell from the free list.
 - Add the current cell to the new group.
 - Make a connection list of all cells which are both in the free list and within A* connection distance of the current cell.




Figure 6.7: Illustration of the formation of pattern objects. Differently coloured groups of hits represent different patterns whilst red boundaries correspond to edges of horizontal and vertical pad gaps.

- Repeat these steps, setting each cell in the connection list as the current cell, until no more connections are found.
- Repeat these steps until the free list is fully depopulated.

This procedure produces a list of groups of connected cells which will eventually form TREx patterns. From this point onwards the pattern recognition operates on each of these pattern objects individually, without knowledge of the state of other pattern objects.

6.3.5 Edge detection

The first step of pattern recognition on a pattern object is to search for edges. These are essential for pattern recognition as they act as seeds for finding junction objects and start and end points for forming A* paths. Throughout pattern recognition edges are defined by groups of cells with an average position weighted towards the side of the pattern object on which they are found.

The following steps apply only to cells not flagged as containing only low charge hits (subsection 6.3.3). Low charge hits are ignored for edge detection due to their appearance as non-physical hits surrounding highly ionising tracks in problematic events. This can create non-physical edges which are difficult to detect as false positives and cause problems throughout pattern recognition.

The steps for forming these edge groups are illustrated in Figure 6.8. Cells are first gathered into six groups representing those in both the highest and lowest 2 layers in x , y and z . Within each of these groups cells are further grouped using the recursive strategy described in subsection 6.3.4 into candidate edge groups. From this point on the algorithm seeks to reduce the number of candidate edge groups by removing redundant or obsolete ones.

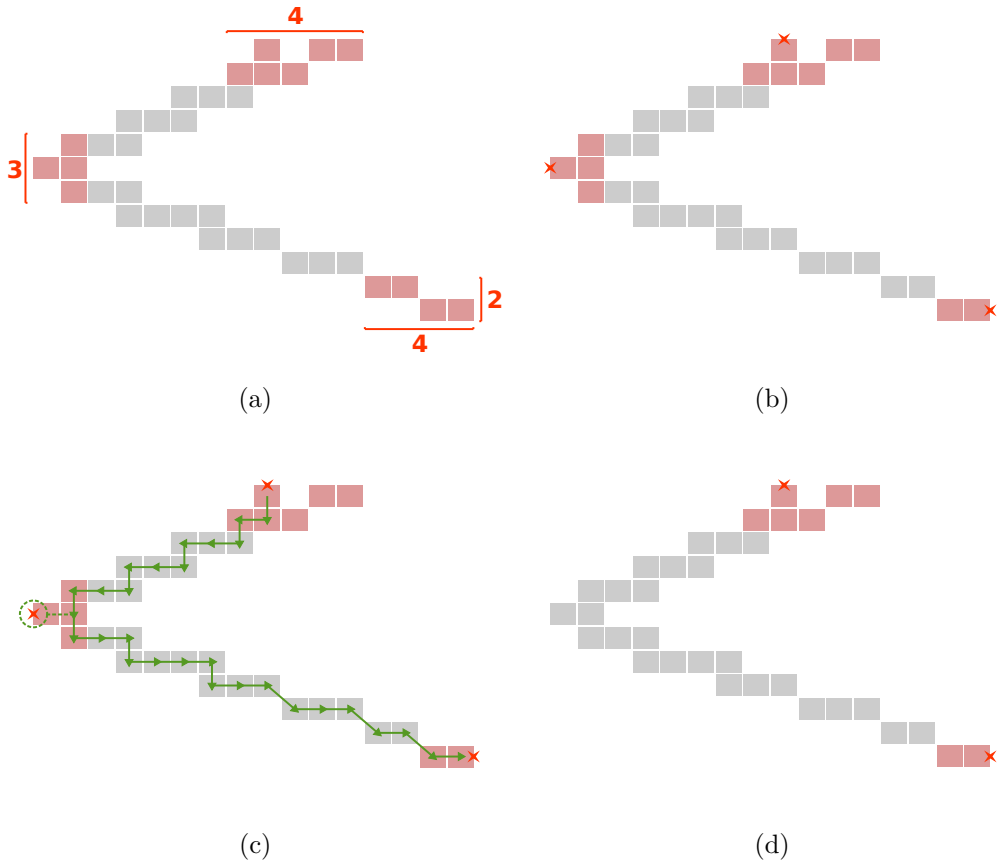


Figure 6.8: Illustration of the finding of edges in a pattern object with red squares indicating cells on edges and stars the average position of these edges weighted towards the appropriate edge. Figure 6.8a shows initial determination of these edge hits and Figure 6.8b their average positions once initial redundant groups have been removed. Figure 6.8c shows one false being removed leaving only two valid edges as shown in Figure 6.8d.

Redundant edge groups are those already expressed in other views. For example an edge group may be found at both maximum z and minimum y (see Figure 6.8a). To favour the most head on view all but the edge group of smallest extent in the two transverse directions is discarded in the case of such overlaps.

Many of the edge groups remaining at this point are false edges — they are actually just points on a path between two other edges. To remove them A* paths are formed between pairs of candidate edge groups. Any other group that lies on or near one of these paths is assumed to be a false positive and discarded, as shown in Figure 6.8c. This leaves the algorithm with a complete list of only ‘true’ edge groups in almost all circumstances.

The above step requires defining when an edge group lies ‘on’ a path. This is done through A* connection distance, since this can be computed quickly and provides some flexibility. If an edge group’s average cell has a connection distance of less than 5 from any cell in the path it is considered false.

The existence of the high charge blocks of hits discussed in subsection 6.3.3 presents an additional problem since these can register false edges at large distances from paths. To counteract this a modifier is applied to A* connections searched for the purposes of removing false edges. This modifier makes connections much cheaper between cells with an occupancy flag, by a factor of 0.3. In this way false edges from problematic events are cleared without effecting the reconstruction of standard events.

One remaining issue arises when considering tracks stopping within the TPC in a high multiplicity pattern. These do not have edges on extremes in x , y or z and so will not be picked up by edge detection. Since TREx was designed to find gas interactions with a goal of reconstructing low momentum protons this is a serious problem.

The solution is illustrated in Figure 6.9. Using A* paths again connections are made between existing edge groups. Any cells on or near these paths are removed from consideration. The edge detection process is repeated on any remaining cells, if there are any, followed by another pass of the clean up for false edges using every edge group found so far.

If only two edges are found then a simple path between them will suffice. In the event of three or more edges pattern recognition must seek out junction objects.

6.3.6 Junction object detection

If a pattern object contains more than two edges then paths between them must connect at one or more locations. Junction object detection pinpoints these loca-




Figure 6.9: Illustration of the second pass finding for extra edges in a pattern object. Figure 6.9a shows an edge which has been missed in the first pass. In Figure 6.9b a path is drawn between existing edges and nearby cells removed. Finally standard edge detection gives the new edges shown in Figure 6.9c and once false or edges have been cleared up we are left with the true edges shown in Figure 6.9d.

tions.

As shown in Figure 6.10 junction objects are detected by forming A* paths which start at the same edge but end at different edges. The junction object is taken to be the area within which these paths begin to diverge from each other.

In pattern objects with more than three edge groups a given start edge will provide multiple choices for pairs of end edges. To ensure the best estimate of true junction location only the pair most likely to lead to a clear point of divergence is chosen. This is computed as the pair of end edges that maximises the area of a triangle between the start and two ends.

Some leeway is required when choosing the cell where divergence occurs. Since paths left by charged particles have non-trivial girth two path objects starting in the same location may separate before reaching a true junction location. To counter this paths must pass a certain threshold before the point of divergence

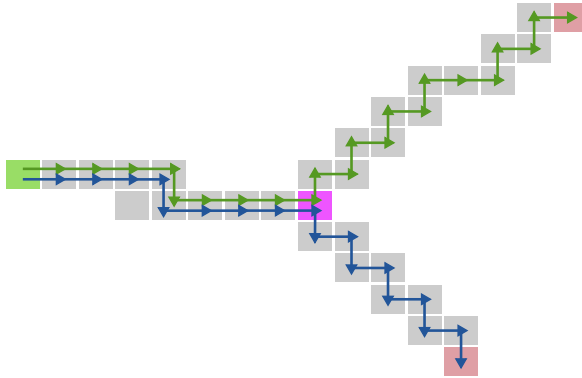


Figure 6.10: Illustration of junction object detection. Green and red boxes represent start and end edges respectively with green and blue arrows representing two different A* paths between the start and two ends. The pink box represents the junction object at which the two paths diverge.

registers. Once this threshold is reached the paths are extrapolated backwards slightly to gain the best estimate of true junction location and the cell at this location is stored.

The procedure gives one estimate of junction location per edge. The cells at these estimates are considered together and combined into junction objects using the same algorithm as pattern formation (subsection 6.3.4). Multiple junction objects can thus be formed when their candidate cells are clearly in different locations.

6.3.7 Path object formation

In the case of two edge groups producing path objects is simple. It is formed by making an A* connection between them.

For multiple edge groups with one or more junction objects the picture is more complex. The simplest case is one junction object and many edges — here path objects are formed between the junction object and every edge group.

The case of multiple junction objects is more complex still. In this case preliminary path objects are formed between every junction object junction object and every edge group and also between every junction object and every other junction object. Most of these are then eliminated on the grounds that they pass via a junction to which they're not supposed to connect (see Figure 6.11b) and are thus redundant.

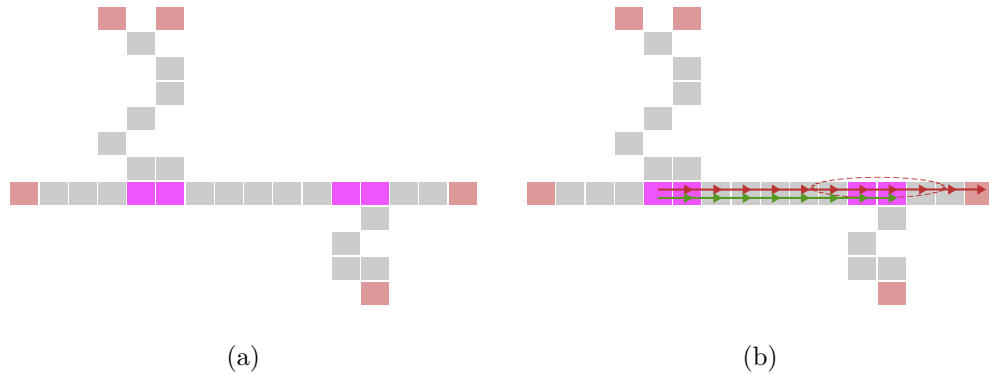


Figure 6.11: Illustration of path object construction in events with multiple junction objects. Edge groups are shown in red and junction objects in pink. Figure 6.11a shows the pattern object with edges and junctions highlighted. Figure 6.11 shows one accepted path object (green arrow) and one rejected path object (red arrow) where the rejection is due to proximity with the second junction.

6.3.8 Cell association

So far path objects have been formed providing estimates of the start, end and intermediate cells for a true path. A* does not return all cells belonging to a true path though, rather it returns a chain of cells representing the most efficient route from start to end. To form cluster objects the association of surrounding cells with the path objects is necessary.

In pattern objects containing single path objects this association is fairly straight forward but for multiple path objects care must be taken to associate cells with the correct one.

The process is illustrated in Figure 6.12. Using the A* connection costs calculated for path finding, cells are added simultaneously to all path objects.

Briefly, this is done by iteratively checking cells and adding them to the path with which their A* connection cost is lowest (unless the cost is above a threshold in which case the cells are not used). Through this method cells end up associated with their optimal path object.

If a cell has equal connection costs to two different path objects it ends up associated arbitrarily with the first path object — this circumstance arises fairly rarely.

6.3.9 Kink finding

Peculiar to gas interaction analyses is the need for identifying kinks. So far junctions have only been detected for pattern objects with three or more edge groups. Since

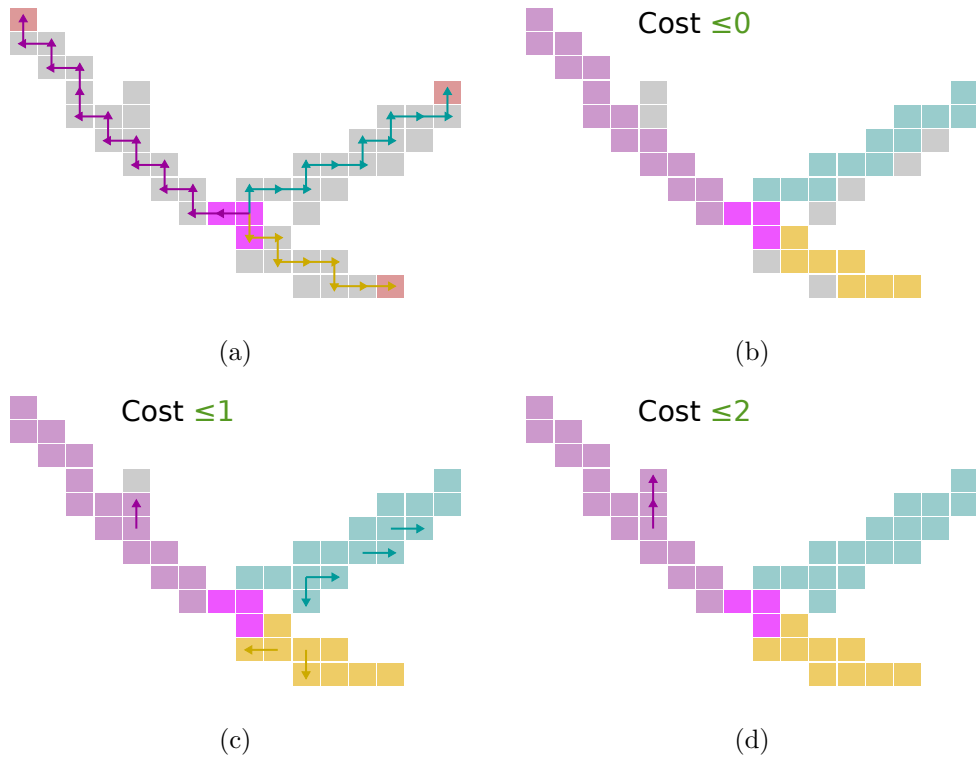


Figure 6.12: Illustration of cell association. Figure 6.12a shows a junction (pink) connected to edge groups (red) via distinct path objects (purple, blue and yellow). Figure 6.12b shows the association of cells with the path objects they are directly on (again, purple, blue and yellow). Figure 6.12c and Figure 6.12d demonstrate the addition of cells at greater A* costs.

a true vertex with two exiting tracks produces only two edge groups plus a sharp kink an extra step is needed to find it. This also applies to patterns with more than two edge groups since an event can have a junction object from a δ -ray *and* a kink from a true interaction.

Kink finding algorithms are illustrated in Figure 6.13. The central kink finding algorithm is the perpendicular distance check (Figure 6.13c and Figure 6.13d) which separates most path objects with sharp kinks from curving path objects. Other checks are needed for specific topologies where the central algorithm runs into difficulties. The local angle check (Figure 6.13a and Figure 6.13b) recovers true junctions from delta rays at the edges of instrumented regions. The parallel distance check (Figure 6.13e and Figure 6.13f) is for true junctions featuring one very short track.

For the perpendicular distance check kinks and curves are separated using the density of cells as a function of distance from a line between the two edge groups.

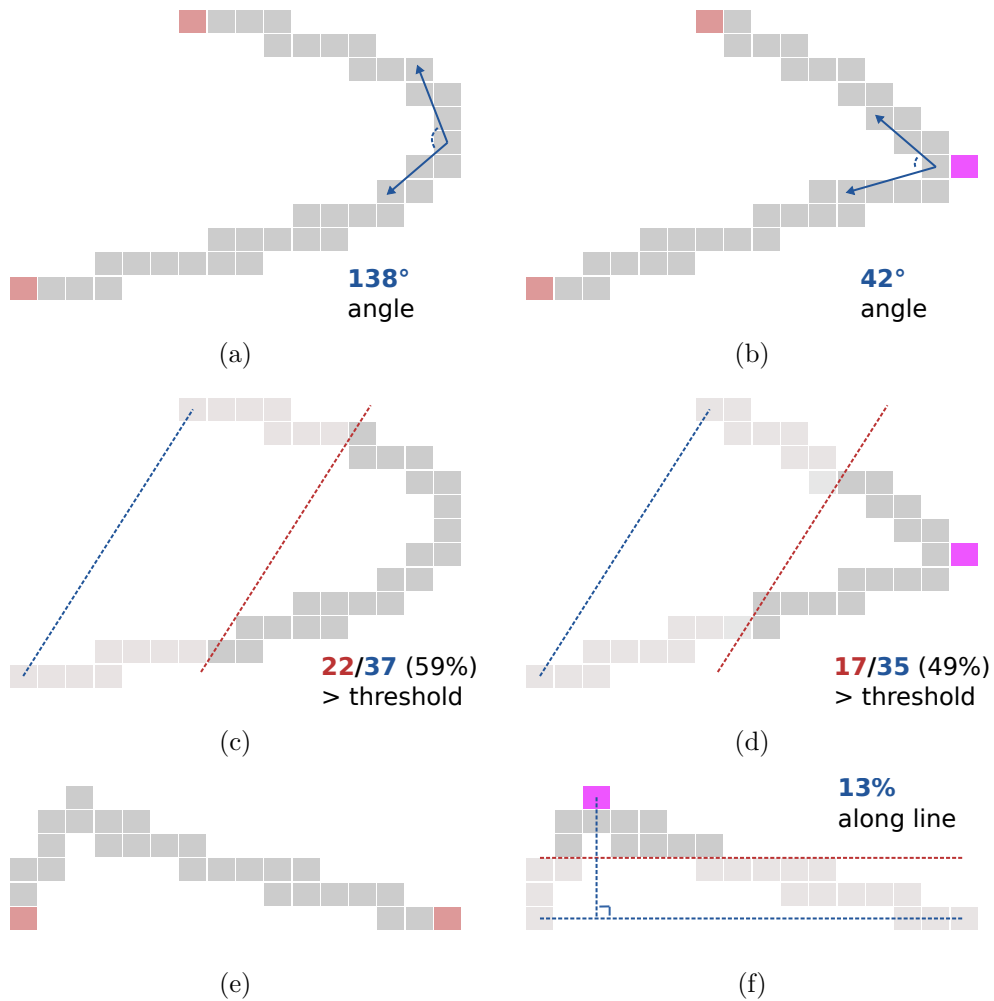


Figure 6.13: Illustration of TREx kink finding. The basic local angle check is illustrated in Figure 6.13a (curving track) and Figure 6.13b (true kink). Arrows show the points between which angles are calculated and the pink cell shows the candidate junction location. Figure 6.13c (curving track) and Figure 6.13d (true kink) show the perpendicular distance check counting cells above a certain distance (red line) from the line between start and end edges (blue line). Figure 6.13e and Figure 6.13f show the parallel distance check where position of the most distant cell from the line between start and end edges (blue line) is checked or its projected position provided its distance is above threshold (red line).

For true kinks this function is fairly flat but curves give more cells far from the line than near to it. Practically the procedure for each path is:

- Draw a line between the start and end cells.
- Find the cell at maximum perpendicular distance from this line and store this distance.

- Count the number of hits whose perpendicular distance from this line is above a fraction (0.56) of this maximum.
- If this number is below a fraction (0.52) of the number of cells in the pattern a kink is established.

The local angle check involves running through every cell in the path object and finding the angle between it and the two cells 8 units forwards and backwards from it. If this angle is lower than 80° a kink is formed. The conservative cut of 80° is used for two reasons. First this check is sensitive to local fluctuations, increasing false positives if a less conservative cut is used. Second the specific topology this check was designed for (activity on the edge of the instrumented region) characteristically provides kinks of very sharp angle, alleviating the need for a stronger cut.

The parallel distance cut looks for kinks missed because one of the two true paths had a very small length within the instrumented region. This is done by searching for kinks which are very close to the start or end of the path. The procedure is similar to the perpendicular distance check but somewhat simpler. For each path:

- Draw a line between the start and end cells.
- Find the cell at maximum perpendicular distance from this line.
- Make sure this cell is at least 2 units from the line to avoid false positives from fluctuations.
- Check the parallel distance of this cell along the line.
- If this parallel distance puts the cell less than 0.3 or more than 0.7 of the way along the line a kink is established.

Where a kink is found a junction object is placed at its location. This is taken to be the cell of maximum local angle if the local angle check found it or else the cell furthest from a line between the two edge groups. With the junction object in place the path object is split in two around it. The old path's associated cells are divided between the two new paths using the hit associating algorithm (subsection 6.3.8).

6.3.10 Determining cluster object orientation

We now have path objects containing both an ordered list of cells forming its backbone and an unordered list of associated cells to include in its cluster objects. The

procedure of forming these cluster objects is performed on a path by path basis. The first step is to go along each cell in the ordered list and determine whether it corresponds to a horizontal or vertical cluster.

Determining orientation is performed by checking the local angle at each cell along the path object. This is done using the dichotomy technique illustrated in Figure 6.16. Angles are found over progressively smaller intervals until local fluctuations due to detector geometry become more significant than the true shape of the track. At this point the angle at each cell is found by interpolation between the established angles.

In the case of a spiralling track, where direction changes often, a preliminary step inserts breaks for computing angle wherever the path object changes z direction. From here the check proceeds as normal.

Once local angles have been determined they are used to divide regions of

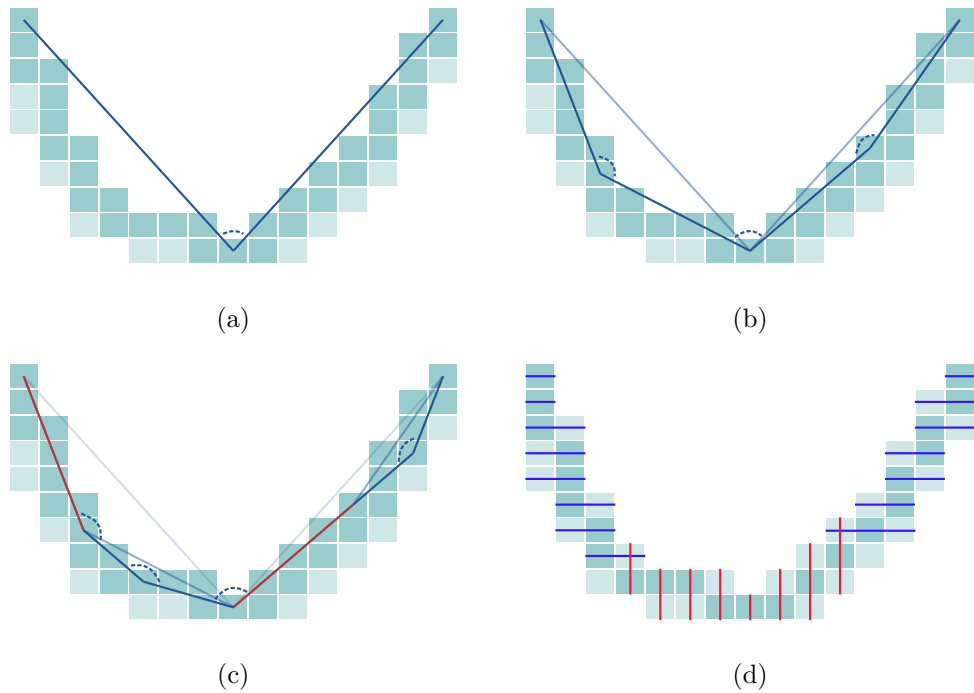


Figure 6.14: Illustration of cluster object angle determination via the dichotomy technique. Figure 6.14a shows the angle being determined between the edges of a path and the cell in its middle (blue lines). In Figure 6.14b this division is taken further with angles determined at points between them middle cell and the edges (solid blue lines). At further divisions as seen in Figure 6.14c some angles are ignored because they are no longer bigger than those at the previous step (red lines) while others can proceed further (solid blue lines). Figure 6.14d shows cells of horizontal (blue line) and vertical (red line) orientations determined from these angles.

horizontal and vertical orientation. Wherever the angle is more than 55° from the vertical a horizontal orientation is assigned. Otherwise a vertical orientation is. To prevent constantly changing cluster orientation in tracks near 55° runs of fewer than 8 cells of the same orientation are flipped so orientation matches their neighbours.

At this point in the algorithm each cluster object features a single, seed, cell and an orientation.

6.3.11 Interpolation and extrapolation

The ordered list of cluster objects within each path object may contain small numbers of missed cells in gaps or at the edges (usually because of charge cut procedures mentioned in subsection 6.3.3 and subsection 6.3.5). To correct for this some extrapolation and interpolation is performed.

This is a simple procedure which involves looking for gaps in or at the edges the list of cluster objects and filling them with any element from the path object's list of associated cells. It is illustrated in Figure 6.15.

Interpolation and extrapolation both use orientation information to determine direction for projection. Vertically oriented cluster objects give interpolation and extrapolation in the y direction and horizontally oriented cluster objects give interpolation and extrapolation in the z direction. New cluster objects are given the same orientation as the cluster objects they were interpolated or extrapolated from.




Figure 6.15: Illustration of path object extrapolation and interpolation. Figure 6.15a shows a gap (dotted arrow) within a path object (blue arrows) filled in Figure 6.15b by interpolation. Similarly Figure 6.15c shows missing cells at the end which are recovered in Figure 6.15d by extrapolation.

6.3.12 Horizontal and vertical clustering

Each path object should now feature a complete ordered list of cluster objects, each with established horizontal or vertical orientation, and an unordered list of associated cells. At the clustering stage the cluster objects are filled with elements from the associated cells.

Before clustering proper it is necessary to ensure there are no redundancies. From the start of the path object to the end any cluster objects in the same layer (y for vertical clusters and z for horizontal) are candidates for removal. If two in the same layer are spatially close (within 5 units) one is removed (since the cluster seed is irrelevant for final output it doesn't matter which).

Finally associated cells are added to cluster objects. Since redundancies have been removed all associated cells can be safely attached to their nearest cluster object in the same layer. After this step all associated cells that can be added to a path have been.

Tightly curving path objects may lead to small numbers of cells remaining unused at the transition between horizontal and vertical cluster objects (see Figure 6.16b). This is a known issue relating to the way A* paths tend to follow the inside of curves and can potentially be resolved in future code by forcing some extrapolation between horizontal and vertical cluster objects.

6.3.13 Cell association corrections

The algorithm has been tuned to ensure that the maximum number of cells end up associated with the correct path. Despite this there is some chance of cells from




Figure 6.16: Illustration of TReX clustering. The initial path object (blue arrow) and associated cells (blue cells) are shown in Figure 6.16a. In Figure 6.16b the associated cells are in horizontal (purple), vertical (red) or a horizontal and a vertical (pink) clusters. A small number of cells in tightly curving tracks may be unused in either.

one true path being associated with a different reconstructed path object or cells in a path object containing charge from multiple particles. To fix this an additional correction step is employed.

Two broad steps are involved, both illustrated in Figure 6.17. The first check looks for overlaps between path and junction objects while the second looks for cells associated with the wrong path object.

The first check considers cells in a path object which are in the same layer as those from a junction object (at the same y for vertical cluster objects or z for horizontal cluster objects). If these cells are spatially close (within 2 units in the other two dimensions) then that cluster object is removed from its path object and all of its cells are added to the junction object. If this creates a break in the path object any isolated clusters are also removed from and their cells added to the junction object.

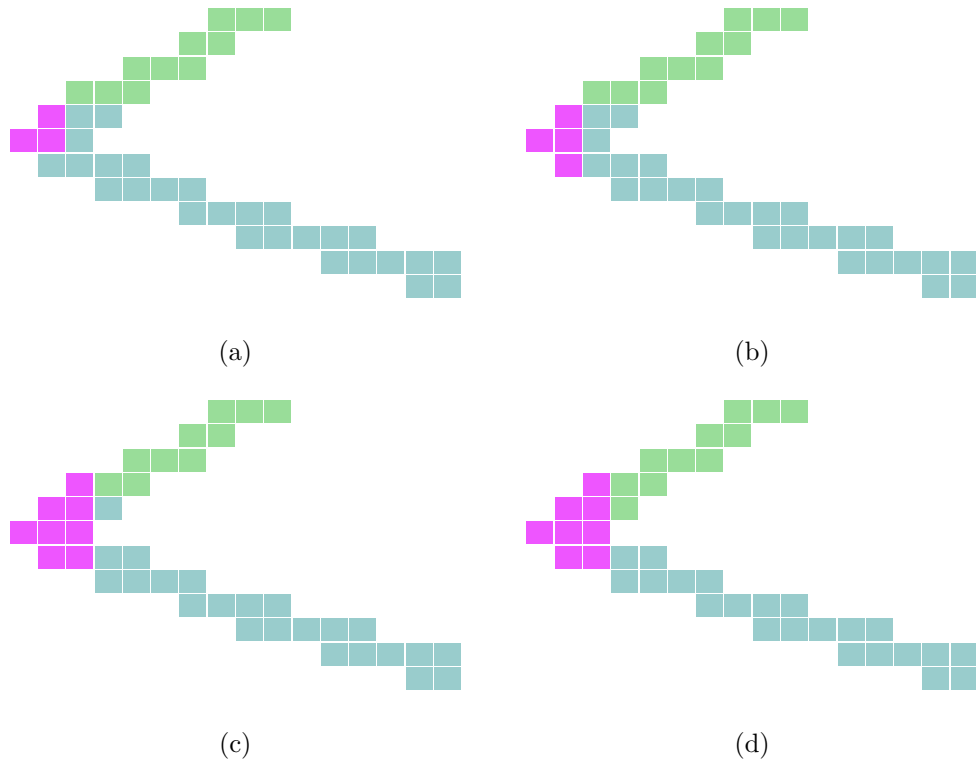


Figure 6.17: Illustration of cell association corrections. Figure 6.17a shows two path objects (green and blue) and one junction (pink). In Figure 6.17b blue cells adjacent to the junction object have been added to it. In Figure 6.17c and Figure 6.17d cells confused between adjacent path objects are added to the junction object and split between the path objects respectively.

The second check considers cluster object of the same orientation and the same layer but in different paths. If such a pair is spatially close together (again within 2 units) a potential conflict is established. If no spatial gap can be found between the cells of the cluster objects they are simply removed from both paths and their cells added to the nearest junction. If a spatial gap can be found all cells in both cluster objects are split so cells either side of the gap end up in the cluster object on the same side.

6.3.14 Junction object expansion

There are three remaining issues in the pattern recognition output. Unused hits, path objects too short to be of use and junction objects connected by them. Junction object expansion aims to resolve these issues.

The steps used for junction object expansion are illustrated in Figure 6.18. The procedure essentially involves deleting all junctions and unusable paths and reforming them from seeds.

Path objects which are connected to junction objects but too short to fit are simply removed (their cells are now unused). All junction objects at this stage are temporarily removed (again their cells flagged as unused) but the positions of their average cells are saved as seeds.

New junction objects are formed by expanding junction seeds into the list of unused hits using the pattern object formation algorithm (subsection 6.3.4). If two seeds form the same junction object they are merged into one new junction object. New junction objects take on the connected paths of any old junction object whose seeds they contain.




Figure 6.18: Illustration of junction object expansion. Figure 6.18a shows two junctions (pink) connected by a short path (green) with other paths around them (blue, yellow, red and purple) and some unused cells (grey). In Figure 6.18b the short path object and junctions are removed with only their seeds (pink stars) saved. By Figure 6.18c they and nearby unused cells are clustered into one junction object.

This procedure only assimilates unused cells near junction objects. Since the pattern recognition algorithm ensures that this is the only place where cells are likely to be unused it effectively eliminates unused hits within patterns themselves (though patterns too small to form paths or junctions from still have their hits assigned as unused).

6.3.15 Translation

The final step of TREx pattern recognition is translation between pattern recognition objects and TREx classes.

Junction objects and cluster objects at the pattern recognition level contain lists of cells, each of which contains a list of constituent hit pads. When converting both to TREx junctions and TREx HV clusters every hit pad in every cell is added to the new TREx object. The TREx HV clusters take on the orientation of their pattern recognition counterparts.

Path objects map almost directly onto their TREx equivalents. The new TREx path contains an ordered list of TREx HV clusters formed by translating the cluster objects within the path object. Their ordering is preserved.

Pattern objects are also simple to convert. Every pattern recognition pattern corresponds to one new TREx pattern. This contains the translated TREx paths and TREx junctions formed from the original pattern object's path and junction objects. Links between them (which paths connect to which junctions) are preserved.

Once translation is complete the tracking steps of the algorithm can proceed.

6.4 First pass seeding

Seeding is the first step in reconstruction post pattern recognition. Basic track seeds are used for matching between paths on different sides of TPC micromegas gaps, determining t_0 and as starting points for the fitting from which kinematics are extracted.

Seeds are formed from the TREx paths returned by pattern recognition. Provided four or more viable clusters are found an attempt is made to form one. This is done using one of two methods; a Riemann sphere or a simple three point algorithm. Both are attempted and the seed providing the lowest total discrepancy between seed trajectory and the y and z positions of clusters is accepted.

6.5 Matching and merging

Up to this point reconstruction has operated on individual patterns. This separates hits from the same track on either side of the vertical TPC micromegas gap as well as any low angle tracks broken around horizontal micromegas gaps. The matching and merging step joins these broken tracks.

6.5.1 Horizontal gap merge

Horizontal merging is done between paths in vertically adjacent micromegas units associated with the same TPC and readout plane and on the same side of the vertical micromegas gap (i.e. within a single column of units). For each path end by the horizontal micromegas gap the corresponding seed (Section 6.4) is propagated forwards in search of a matching path or junction in a separate pattern.

A match is established where the residual between the propagated seed and a target path or junction is sufficiently low (below 30 mm). If multiple matches are found the one with lowest residual is chosen. In the case of a match the two patterns containing matching objects are merged into a new single pattern with the combined constituents of each as illustrated in Figure 6.19.

If the matched object was a path the two paths are concatenated, preserving cluster ordering and orientation and any connections between the paths and junctions. This produces a new path object which replaces the two originals in the new pattern. Seeding is then redone for this path.

If the matched object was a junction the propagated path is simply added to its constituents.

6.5.2 Vertical gap merge

Vertical gap merging proceeds in a similar way to horizontal gap merging. The main practical differences are the substantially larger gap size and the fact that TREx pattern recognition avoids connections across the gap by design (for horizontal gaps connections are possible but may be missed due to the geometry of tracks crossing at low angles). Recovering connections across the vertical micromegas gap at the matching and merging level is mandatory.

As in horizontal matching the propagation is done from paths in one pattern to paths and junctions in another. Matches are allowed with larger residuals (50 mm), reflecting the larger size of the gap.

Again the patterns are merged in the case of a match. The procedure for this is identical to that in the case of horizontal merging.

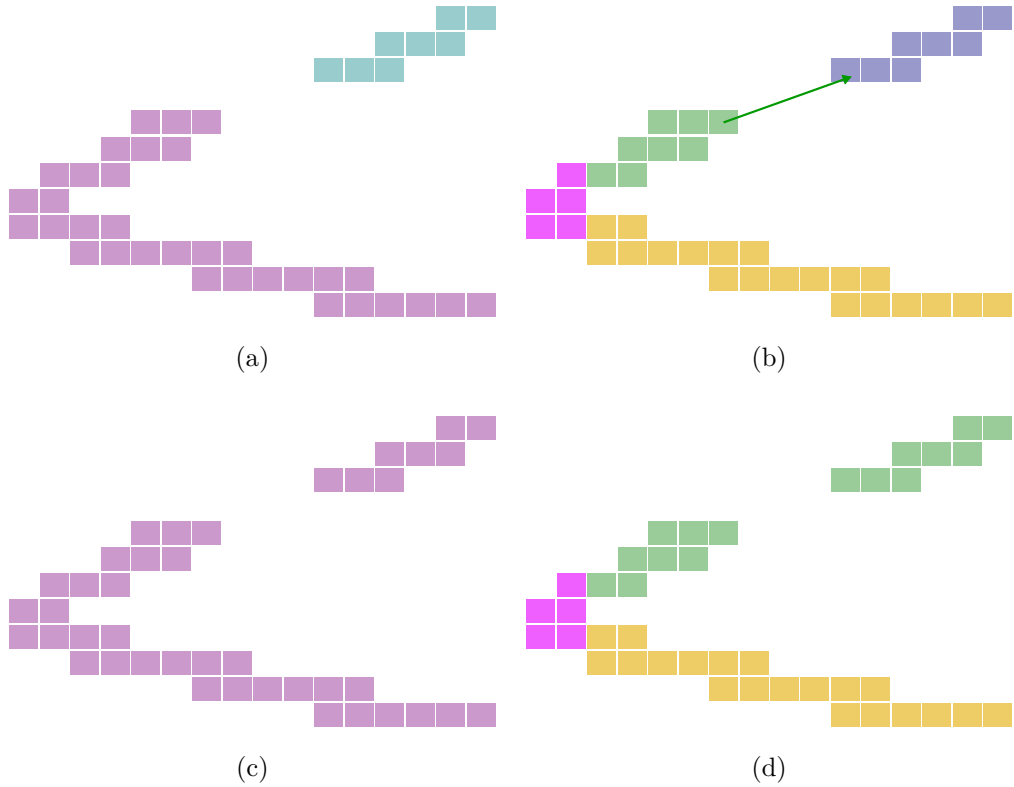


Figure 6.19: Illustration of TREx matching and merging, with different colours representing different paths and patterns. Figure 6.19a and Figure 6.19b demonstrate the state pre-merging, with two patterns containing a pair of paths which can be propagated onto each other (arrow in Figure 6.19b). Figure 6.19c and Figure 6.19d show the single merged pattern.

6.6 t_0 determination

In the ND280 TPCs y and z position for initial charge depositions comes from the physical positions of micromegas pads. The x position is calculated from the time at which ionisation electrons reach the pads. Unless either the central cathode or the whole volume of a TPC half is traversed absolute time (t_0) cannot be determined from the TPC alone. Matching with hits in other detectors is necessary.

6.6.1 Default t_0

A default t_0 for each pattern is selected based on TPC hits alone. This places the entire pattern in the centre of the TPC, with equal distance from the first hit to arrive to the readout plane and the last hit to arrive to the cathode. In lieu of any other information this provides a fair estimate of hit position and in fact provides the correct result in cases where the entire x extent of the TPC half is traversed.

6.6.2 t_0 from constituents

t_0 is determined on a pattern by pattern basis. This requires a valid result from at least one of its constituent paths.

The algorithm loops through every path in the pattern from longest to shortest looking for a t_0 . Once one is found the result is assigned to every constituent. If no t_0 result can be found the value is left as the default.

This preferentially selects t_0 from modules known to give fewer false positives, moving on to other modules only if no matches are found. A separate container is used for hits from each source, passing hits into either the x - z or y - z matching procedures.

6.6.3 Fitting from x - z hits

Working out t_0 from a vertically oriented scintillator bar requires matching the projected x - z position of the TPC path with the actual x - z position of the bar without. This is done by assigning the path a temporary t_0 from the time of the hit being examined, as illustrated in Figure 6.21 for the majority of detectors Figure 6.20 for side ECals.

Every hit in the current container is looped through in search of the best match. Each hit provides a temporary t_0 , provided it will not move the pattern outside the TPC volume. The path, with temporary t_0 assigned, is propagated to the bar's x - z position in an attempt to find a match. The candidate hit with the lowest matching χ^2 is chosen as t_0 source, so long as the χ^2 is below a threshold. Otherwise the search moves on to the next t_0 source.

6.6.4 Fitting from y - z hits

The algorithm for determining t_0 from y - z hits was largely recycled from the previous TPC reconstruction and is illustrated in Figure 6.22. Since these bars are parallel to the x axis matching is agnostic towards the TPC path seed's x position.

The path seed is propagated towards each hit in the current container in order to find the best match. As with the x - z matching described in subsection 6.6.3 the hit providing the lowest matching χ^2 is sought. If the resulting χ^2 is small enough this hit is chosen to provide the pattern with a t_0 . Otherwise the algorithm moves on to the next t_0 source.

6.6.5 Cathode crossers

Although paths crossing the central cathode are not recombined until later in **TREx** reconstruction these also require a t_0 . How this cathode t_0 is used depends on the previous t_0 values assigned to each path.

The cluster of maximum time is assumed to be on the edge of the cathode, giving t_0 equal to the hit's time minus time taken for ionisation electrons to drift across the TPC half's x extent. How this t_0 is used depends on the previous values assigned to the two input paths:

- **No t_0 values from input paths:** Assign the value from the cathode to both
- **One t_0 value from input paths:** Copy this value to both paths, unless it differs wildly from the cathode t_0 in which case use the cathode value
- **Two t_0 values from input paths:** If the two are close enough take the one from the longest path, otherwise take the value closest to cathode t_0 ; in either case copy it to both paths.

This ensures that both segments of a path matched across the cathode have the same final t_0 .

For each path the matching searches through possible t_0 sources. Hits are checked in the following order:

- x - z hits
 - Hits from FGD bars
 - Hits from (non-side) ECal bars
 - Hits from POD bars
- y - z hits
 - Hits from FGD bars
 - Hits from ECal bars
 - Hits from POD bars
- x - z hits
 - Hits from side ECal bars




Figure 6.20: Illustration of t_0 determination for TPC hits (grey) using x - z hits from an upstream detector (blue-grey), with x represented by vertical offset. In Figure 6.20a the TPC path's t_0 is indeterminate. Figure 6.20b shows the path checking an upstream hit (blue) which provides a temporary t_0 . In Figure 6.20c the TPC path has been assigned a permanent t_0 from the inspected hit.




Figure 6.21: Illustration of t_0 determination for TPC hits (grey) using hits from side ECal bars (blue-grey), with x represented by vertical offset. In Figure 6.21a the TPC path has indeterminate t_0 . Figure 6.21b shows the path checking an ECal hit (blue) which provides a temporary t_0 . In Figure 6.21c the TPC path has been assigned a permanent t_0 from the inspected hit.




Figure 6.22: Illustration of t_0 determination for TPC hits (grey) using y - z hits from an upstream detector (blue-grey), with x represented by vertical offset. In Figure 6.22a the TPC path has indeterminate t_0 . Figure 6.22b shows the path checking an upstream hit (blue). In Figure 6.22c the TPC path has been assigned a permanent t_0 from the inspected hit.

6.7 Tracking

A tracking stage computes kinematics on the current path objects (merged across micromegas gaps with valid seeds and t_0 defined). This requires a fit to be run on every path in every pattern, using the paths' horizontal and vertical clusters.

6.7.1 Cluster correction

Corrections are applied to cluster positions in data accounting for deviations from expectation due to the inhomogeneous nature of the TPC \vec{B} and \vec{E} fields (the nominal field is modelled as uniform). They use nominal field maps for the TPCs.

There is also an alternate pass of the likelihood fit with alternate corrections, used for calculation of systematic errors. The pass employs an empirical correction in Monte Carlo. Additionally an \vec{E} field correction is employed in both data and Monte Carlo.

6.7.2 Likelihood fit

A log likelihood minimiser is used to fit the path. Clusters of acceptable quality (with sensible extents and no suspicious pads) are used for the fit.

A function for computing the log likelihood is fed into the cluster. Likelihoods are computed separately in y - z and x , but both cases rely on the residual between the fitted helix and each hit pad in a given cluster as compared to expectation to compute a log-likelihood. The total log-likelihood is the sum of y - z and x results.

The likelihood fit searches for the optimal log-likelihood using the MINUIT minimiser[178]. A full fit in x , y and z is attempted first. If this fails an attempt is made to perform separate fits for the y - z and x parameters of the helix model. In this case a final x , y , z fit is attempted once to compute errors.

The full fit is performed once with nominal corrections and again with refit parameters, in the second case providing a 'refit momentum' used in the propagation of systematic errors at the analysis level (see subsection 8.4.3).

6.8 PID

TPC PID computes $\frac{dE}{dx}$ in order to infer particle identity. This is done by calculating energy deposited adjacent clusters and comparing it to the distance between them according to the fitted track.

As shown in Figure 6.23 $\frac{dE}{dx}$ is a function of particle momentum as well as type. In order to determine which true particle hypothesis most resembles the

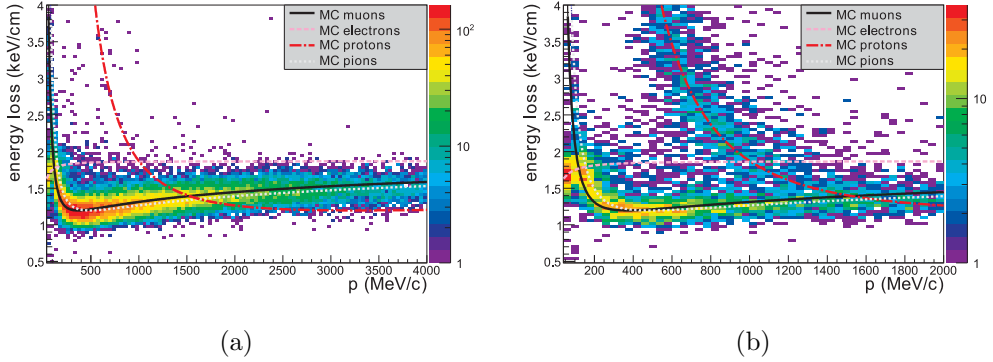


Figure 6.23: Illustration of $\frac{dE}{dx}$ for negatively charged (Figure 6.23a) and positively charged (Figure 6.23b) particles produced by neutrino interactions in the ND280 beam, from [158].

reconstructed path the expected $\frac{dE}{dx}$ is computed for electrons, muons, pions, kaons and protons at the reconstructed particle’s momentum.

These variables for measured and expected $\frac{dE}{dx}$ are saved and subsequently used at the selection stage to compute likelihoods for each particle type. See Section 7.2.1 for a detailed description of their implementation in analysis.

6.9 Likelihood matching

One of the largest backgrounds in the gas interaction analysis comes from through-going muon tracks which have been broken within the TPC, leading to reconstructed objects which appear to start within our fiducial volume. The likelihood matching step at reconstruction is a track–track matching quantifying the extent to which separate reconstructed tracks appear to correspond to the same true track. This matching depends on connections between different paths in either the same or different patterns.

The general procedure, illustrated in Figure 6.24, requires examining pairs of candidate paths (say path 1 and path 2).

Individual (unmatched) likelihoods are already available from the fits to each path (L_{11} and L_{22}) — these can be viewed as the matching likelihood for the path with itself. The matching likelihoods for each path with *the other path* (L_{ij} , where i is the path whose fit is used j the path whose hits are matched with this fit) are also extracted. This is done by computing the likelihood of path j ’s hits corresponding to path i ’s fit.

The total log likelihood of each matching hypothesis (path 1 to path 2 and




Figure 6.24: Illustration of TREx likelihood matching. Blue and purple hits belong to separate paths (1 and 2) separated by a pink junction caused by a δ ray (Figure 6.24a). Figure 6.24b shows individual likelihoods for each track fit with no extrapolation (L_{11} and L_{12}). Figure 6.24c shows the likelihood for path 1 plus the likelihood for path 1 propagated onto path 2 ($L_{11} + L_{12}$) and Figure 6.24d shows the reverse, with path 2 propagated onto path 1 ($L_{22} + L_{21}$).

path 2 to path 1) is given by the sums $L_{11} + L_{12}$ and $L_{22} + L_{21}$ while the likelihood for each path separately is given by $L_{11} + L_{22}$. These sums will eventually be used for both merging and analysis but for now the raw likelihoods of L_{12} and L_{21} are saved for each pair of candidate paths for later use.

All of these likelihood values are stored inside path objects which in turn are stored as constituents of the global tracks formed in later stages of the TREx reconstruction.

6.9.1 Intra-pattern matching

Many background events in the gas interaction analysis are through-going muons broken by a δ -ray. By design the TREx pattern recognition breaks paths around sufficiently large δ s to avoid the risk of discarding true neutrino interactions, leaving it to analysers to distinguish the two. Intra-pattern matching provides information to aid this by looking for matches between paths on either side of the same junction and saving their matching likelihood.

Each junction within each pattern is examined. All pairs of paths connected to the junction are compared (say 1 and 2) and for each pair the two likelihoods (1 onto 2 and 2 onto 1) are stored.

6.9.2 Inter-pattern matching

Another form of gas interaction background comes from through-going tracks with no δ -ray. Instead, due to either the geometry of pad gaps or missing hits a path is simply broken into constituents in two separate patterns.

Inter-pattern matching operates between patterns in the same drift volume (the same side of the central cathode in the same TPC). The procedure is the same as the intra-pattern case but now pairs of candidate paths are required to be in different patterns and with their ends not already connected to junctions.

6.10 Likelihood merging

The merging algorithm attempts to merge tracks broken by a junction or gap. Pairs of tracks are checked using the track–track matching likelihood discussed in Section 6.9.

For inter-pattern matching between two paths, 1 and 2, the likelihood of the two tracks being unmatched ($L_{11} + L_{22}$) is compared with the likelihoods for the two matching hypotheses ($L_{11} + L_{12}$ and $L_{22} + L_{21}$). A cut on the difference between log likelihoods is used here.

The intra-pattern matching is slightly more involved but uses the same basic logic. In this case there are potentially a large number of matching candidates, so all pairs of matches are checked at the same time and only the highest likelihood match is used. This prevents one track being matched to a suboptimal candidate before its true match has been checked.

Once matches have been found matched paths are merged, preserving the correct order of their clusters. The original paths are stored as constituents in the larger, merged path. The larger path is passed back through the seeding, tracking and PID algorithms to have its own PID and kinematics evaluated (see Section 6.4, Section 6.8, Section 6.7).

6.11 Cathode merging

At this stage most broken paths have been merged. The one remaining task is to recombine tracks broken by crossing the central cathode.

The track–track likelihood matching discussed in Section 6.9 and Section 6.10 is inadequate for this job since it doesn't take account energy loss and scattering as a particle crosses the 13 mm thick cathode. Instead a track–cluster χ^2 matching is used.

The matching involves projecting paths directly onto clusters on the opposite side of the cathode. This matching does not use path x co-ordinate since t_0 may not be correctly assigned for one or more candidate path. If a match is found then the output is in identical format to that described in Section 6.10, with a Kalman filter fit between two track segments (see Section 5.6).

6.12 Tracker and global reconstruction

An extra step is needed to convert TPC patterns to lists of TPC tracks which can be passed to tracker reconstruction in the same manner as in the old reconstruction. Beyond this both tracker and global reconstruction proceeds identically to as they were for previous TPC reconstruction, discussed in subsection 5.6.3 and subsection 5.6.7.

6.13 Validation

Simple scans of gas interaction Monte Carlo illustrate huge improvements in TREx as compared to the previous reconstruction. Some of these are shown in Figure 6.25. In

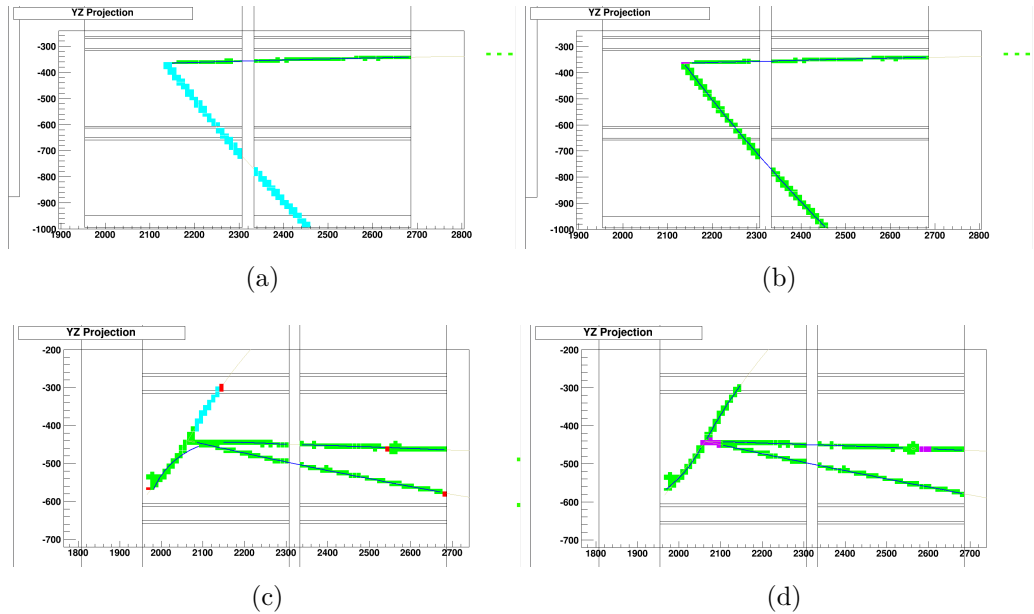


Figure 6.25: Illustration of TREx performance compared to old TPC reconstruction on gas interaction Monte Carlo. Old TPC reconstruction and TREx are illustrated for a two track event in Figure 6.25a and Figure 6.25b respectively and a four track event in Figure 6.25c and Figure 6.25d respectively.

general the performance of **TREx** compares very favourably with the inability of the old reconstruction to recognise vertices in a useful way. Even on non-gas interaction events performance is as good as (for most events) or better than (for topologies such as backwards curving tracks) the previous reconstruction.

Performance has also been evaluated through extensive comparisons between **TREx** and old TPC reconstruction on data. Two instructive variables, efficiency and purity, are shown in Figure 6.26. Purity represents the fraction of hits in a given reconstructed track which actually come from the corresponding particle, and this has seen a great improvement in **TREx** as hits from δ -rays and nearby particles are effectively filtered out of tracks. Efficiency represents the fraction of hits with charge from a given true track which end up associated with its corresponding reconstructed track. A slight increase in tracks with just below 100% efficiency is seen. This is expected since hits near junctions, for example δ -rays, are now being assigned to junction objects rather than paths.

It is unfeasible to compare vertex reconstruction performance since old TPC reconstruction had no concept of a vertex and was completely unsuited for detecting them. **TREx** performance for vertex reconstruction is shown in Figure 6.27. The x position of vertices is recorded with high accuracy, mostly accurate to within 10 mm. For y - z position resolution there is a hard limit of around 10 mm given by micromegas pad size[†]. Most vertices have their position reconstructed accurately to within one or two pads. A tail in this distribution is contributed by highly collinear tracks and higher multiplicity events, both of which make precise vertex position difficult to reconstruct.

Vertex track multiplicity is also checked, as shown in Figure 6.28. In general there is a diagonal trend with events showing n particles in truth most likely to be reconstructed with n particles. Underestimates of multiplicity by reconstruction mostly stem from tracks whose length within the TPC is too short to measure, either because the vertex itself is near an edge or the path is particularly short. A lower efficiency for reconstructing tracks as above 100 MeV (Figure 6.28b indicates that tracks are often found but lack sufficient hits for a reliable extraction of kinematics. Improving reconstruction for short tracks is a priority for future **TREx** development. We can also underestimate multiplicity where there are pairs of highly collinear tracks sharing hits. It is rare for reconstructed multiplicity to overshoot true multiplicity.

The distributions of general kinematic and PID variables are largely un-

[†]This could be overcome in future versions of the **TREx** code by fitting vertex position along with track kinematics. At the moment vertex position is simply given by the weighted average of the positions of all hits in a junction.

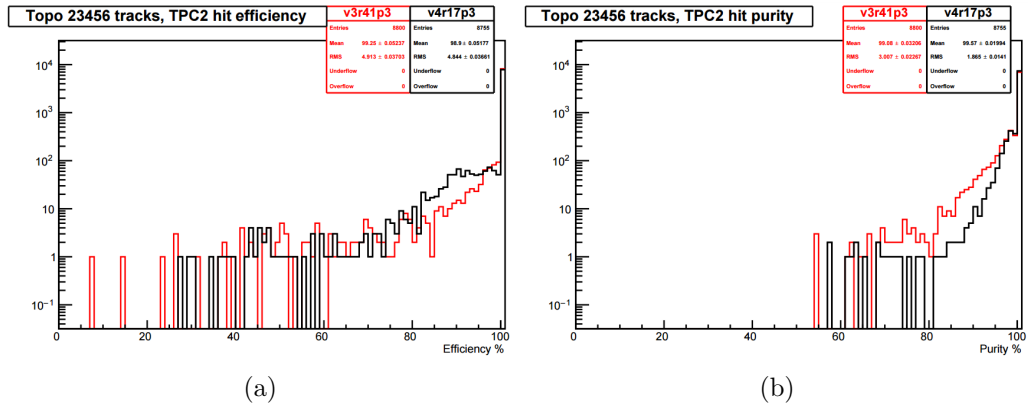


Figure 6.26: Illustration of TREx efficiency (Figure 6.26a) and purity (Figure 6.26b) compared to old TPC reconstruction on beam Monte Carlo. Efficiency for the old reconstruction version (v3r41p3) is in red while TREx software version used for analysis (v4r17p3) is in black.

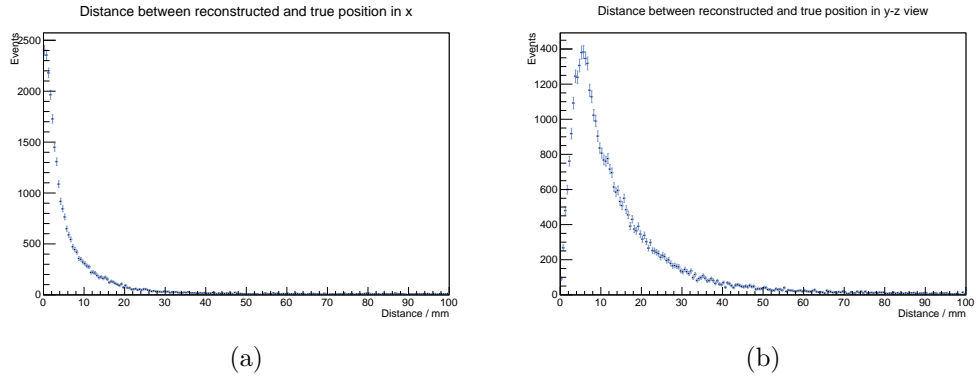


Figure 6.27: Illustration of TREx vertex position accuracy, showing the difference between reconstructed and true vertex using NEUT gas interaction Monte Carlo for events within the TPC instrumented region. Difference in x position (Figure 6.27a), which comes from t , and $y-z$ position (Figure 6.27b), which come from micromegas pad position, are shown separately.

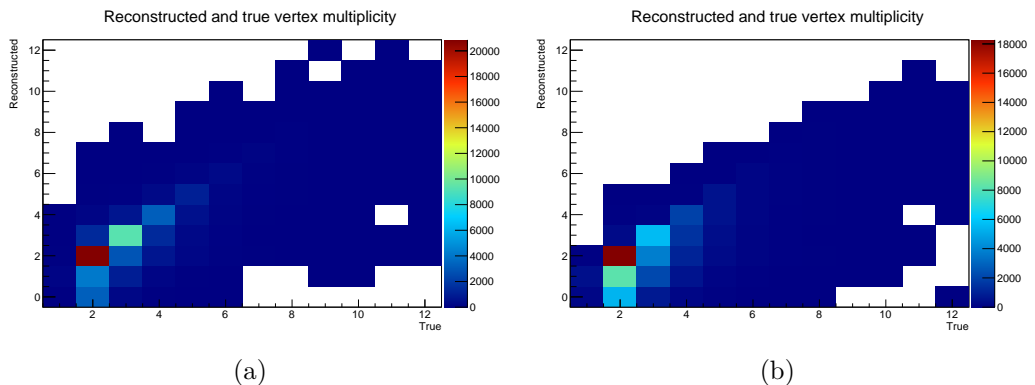


Figure 6.28: Illustration of TREx vertex multiplicity, showing reconstructed against true multiplicity for all particles (Figure 6.28a) and for particles reconstructed with over 100 MeV momentum against true particles with over 100 MeV momentum (Figure 6.28b) within the TPC instrumented region using NEUT gas interaction Monte Carlo.

changed in value since the old TPC reconstruction. This is also expected since the algorithms producing these variables have not been extensively redeveloped for TREx. Differences in momentum resolution are shown in Figure 6.29 and agree well between the two.

All of this demonstrates that the greatly improved flexibility of TREx with regards to handling interactions in the TPC and other awkward topologies has not come at a cost to general reconstruction performance. This combined with comparable processing times between TREx and the old reconstruction have led to TREx being adopted as the official TPC reconstruction for future data processing at T2K.

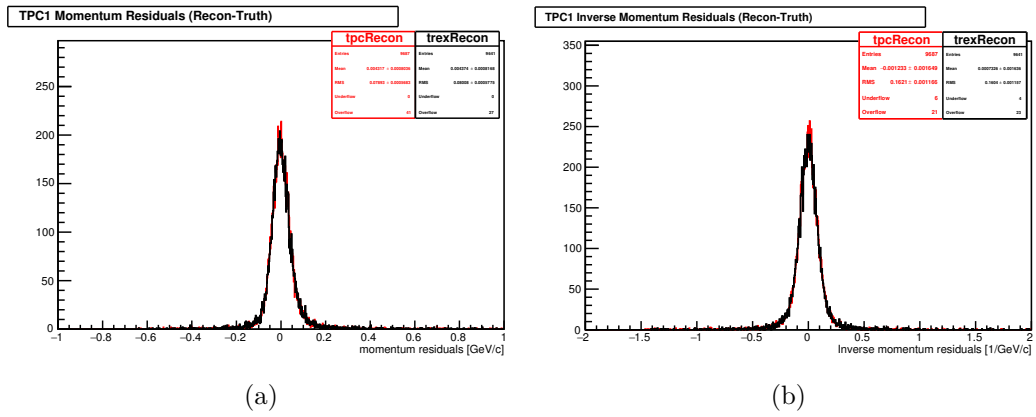


Figure 6.29: Momentum residual distributions for old reconstruction and TREx covering both momentum Figure 6.29a and inverse momentum Figure 6.29b.

Chapter 7

Selection

The gas interaction proto-analysis[172] successfully recovered many instances of neutrino interactions on argon gas in real data.

This analysis goes a step further. Its purpose is two-fold; to establish a reliable cross section for neutrino interactions on argon in our energy range and to gain understanding of their kinematics and the systematic uncertainties thereof so that future generations of the analysis can produce essential inputs for neutrino interaction models.

These goals require a more reliable selection than previously used. The selection itself is discussed here and systematic uncertainties involved in Chapter 8. The results are presented and discussed in Chapter 9.

7.1 Proto-analysis

A brief overview description of the proto-analysis is helpful for the purposes of understanding decisions made while developing the present analysis. Many of our decisions were informed by difficulties encountered over its course.

7.1.1 Preselection

Until recently **TREx** was not a part of general ND280 reconstruction. To properly reconstruct interactions in the TPC gas, events had to be processed locally. Doing this with the full ND280 dataset was impractical, so a preselection (or ‘stage one selection’) was employed to reduce the number of events to a manageable level. This used only hit-level information (prior to reconstruction). It demanded

- activity in at least one TPC

- a lack of activity upstream of the TPC in question

and thus cut down the number of events to reprocess by a factor of around twenty.

After the preselection, events were processed with TREx and passed to the main analysis.

7.1.2 Main selection

The main selection (or ‘stage two selection’), for charged current ν_μ interactions on argon in our fiducial volume, used the following cuts:

1. **Data quality** cuts requiring no problems with data taking during the time of the event (applied to real data only).
2. **Stage one selection** cut requiring the TPC being checked to correspond to the TPC which passed the preselection.
3. **Bunch timing** cuts requiring the t_0 of the pattern being checked to fall within a window of 60 ns of central beam bunch time.
4. **Muon identification** cuts requiring good TPC muon PID for the path being checked.
5. **Fiducial volume** cuts requiring the path being checked to start well within the TPC instrumented region.
6. Various δ **veto** cuts to remove background caused by delta-rays on a through-going path.
7. Various **broken track veto** cuts to remove background caused by reconstruction failures on through-going tracks.

Of these cuts item 6 and item 7 proved highly sensitive to the physics we were trying to observe, but at the time they were needed to achieve good purity.

7.1.3 Issues raised

The proto-analysis successfully identified a substantial number of interactions on argon gas in data. It achieved a purity of 77% with an efficiency of 17% on NEUT Monte Carlo events. However, substantial uncertainty over the true efficiency in data made a reliable cross section impossible to derive.

Much of the analysis, particularly the preselection and the delta-ray and broken track vetoes, were determined to be highly sensitive to the physics we were trying

to measure and caused substantial disagreement between Monte Carlo generators and data over our actual selection efficiency and purity.

In addition our ‘hairy events’ were identified for the first time. These have been discussed in Section 5.4 and also subsection 6.3.3 and caused unexpected differences between data and simulation.

7.1.4 Improvements since the proto-analysis

A huge practical difference is the availability of large amounts of TREx processed data. We currently have access to all ND280 data in this form can and run the analysis directly on it. This eliminates the need for a preselection and not only reduces our uncertainties but also allows for an almost 4π analysis rather than one restricted to forwards going tracks.

Improvements to reconstruction have also been implemented. Hairy tracks are now handled at the reconstruction level (see subsection 6.3.3) and have associated systematic uncertainties evaluated. t_0 determination has been greatly improved particularly for tracks entering from side ECals, greatly reducing our single track background.

Many of these reconstruction changes have fed into improvements at the analysis level. Likelihood matching information between tracks is available which allows much background from broken tracks and delta-rays to be removed. TREx is now fully integrated into global reconstruction allowing both TPC pattern level information and global tracks to be used in the analysis. Significantly, all of these changes allow us to replace any physics sensitive cuts without severely limiting our selection purity.

7.2 This analysis

The current analysis has been developed completely from scratch and without any cuts thought to bias physics results. There is no preselection and the cuts are designed to be as agnostic as practical to the precise kinematics of interactions in our fiducial volume.

7.2.1 Cut flow

The cuts we employ are:

1. **Event quality** cuts, requiring the standard data quality flags to be passed.

2. **Total multiplicity** cuts, requiring at least one TREx path.
3. **Track fiducial volume** cut, requiring the global track being checked to start in the fiducial volume.
4. **Track muon PID** cut, requiring the global track being checked to have muon PID.
5. **Track negative charge** cut, requiring the global track being checked to have negative charge.
6. **Track momentum** cut, requiring the global track being checked to have momentum over 100 MeV.
7. **x -angle** cut, requiring the global track being checked to have x -angle such that $|\cos \theta_x| < 0.9$.
8. **Vertex fiducial volume** cuts, requiring the projected positions of vertex candidates to lie within a strict fiducial volume.
9. t_0 **quality** cut, requiring a sensible t_0 source for the path being checked if necessary.
10. **Likelihood matching** cuts, removing backgrounds from delta-rays and broken through-going tracks.
11. **TPC cleanliness** cuts, removing backgrounds from misreconstructed through-going tracks in TPCs with large amounts of activity.

Since we're searching for any ν_μ interaction within our fiducial volume we may have a large number of tracks per event flagged as lepton candidates. This is particularly true at earlier stages in the cut flow. The selection maintains a list of potential muons starting in the fiducial volume which is updated at each cut. For the purposes of systematic error propagation and event categorisation the muon candidate is considered to be the highest momentum track in this list. The event as a whole is considered to fail a cut when no candidate tracks remain.

Event quality

Event quality cuts are standard for ND280 analyses. They apply only to data (bad quality Monte Carlo events are not simulated) and simply rely on flags for good data quality and good beam quality for the data taking period in question.

These flags are set manually by the ND280 data quality and T2K beam groups based on diagnostic studies.

Total multiplicity

This is another low level cut used to verify that an event is worth considering before moving on to more intensive checks. The event is examined to make sure it contains at least one track with a TPC component.

If this cut is passed a list of all global muon candidates is formed which initially contains any global reconstructed track with a segment in the instrumented region of any TPC.

Track volume

The general track volume is defined as the instrumented region of any TPC. This corresponds to six cuboids; one for each half of each TPC. Later in the selection the strict fiducial volume cuts will use harsher criteria.

Every global muon candidate is checked to make sure they meet the fiducial volume criteria. The exact condition is that one or the other but not both of the global track's ends is within the fiducial volume. In addition tracks which end but do not start in the fiducial volume have their direction flipped so that all candidate tracks are defined as starting in the fiducial volume from this point on.

Any global track which does not meet this fiducial volume condition is removed from the list of muon candidates.

Track muon PID

Particle identification is a powerful cut for ensuring that a given track corresponds to a muon rather than some other particle. The $\frac{dE}{dx}$ variables discussed in Section 6.8 are used to compute pulls as

$$\delta_i = \frac{\left(\frac{dE}{dx}\right)^{\text{measured}} - \left(\frac{dE}{dx}\right)_i^{\text{expected}}}{\sigma_i^{\text{expected}}}, \quad (7.1)$$

where i is the particle hypothesis and $\left(\frac{dE}{dx}\right)_i^{\text{expected}}$ and $\sigma_i^{\text{expected}}$ the mean and standard deviation of expected $\frac{dE}{dx}$ for that hypothesis at the particle's reconstructed momentum.

Each surviving candidate global track is checked for PID. To be counted as a muon candidate the track needs either a segment in the SMRD or for its longest well reconstructed TPC segment to pass the PID cut.

The PID cut uses probabilities defined from the pulls described in Equation 7.1:

These are normalised into likelihoods for each particle type:

$$\mathcal{L}_i = \frac{P_i}{\sum_j P_j}, \quad (7.2)$$

which are used to evaluate whether or not the path meets the following PID conditions:

- Muon likelihood $\mathcal{L}_\mu > 0.05$
- One of the following:
 - Momentum $p > 500$ MeV
 - MIP likelihood over proton unlikelihood $(\mathcal{L}_\mu + \mathcal{L}_\pi) / (1 - \mathcal{L}_p) > 0.8$

Any candidate global track failing these cuts is removed from the list of muon candidates.

The effects of these cuts on our sample at this stage in the cut flow are illustrated in Figure 7.1 and Figure 7.2. Whilst our large background of out of fiducial volume muons remains the PID cuts remove a large number of non-muon background which isn't handled anywhere else.

$$P_i \equiv e^{-\left(\frac{\delta_i^2}{2}\right)}. \quad (7.3)$$

Track negative charge

So far our cuts have been designed to pick out muons starting in the TPC fiducial volume with no distinction between μ^- and μ^+ . The negative charge cut is designed to discard any μ^+ which have made it this far. The list of global muon candidate tracks is filtered for charge with any tracks of positive charge removed.

Track momentum

Although we have ensured that our selection includes as little kinematic information as possible there are some unavoidable practical constraints. PID is generally not reliable for low momentum particles. This track momentum cut is necessary to account for this.

Global muon candidate tracks are checked and any where momentum $p < 100$ MeV are removed. Since our true muons are expected to have momenta around 500 MeV with relatively few below 100 MeV this cut has limited effect on our signal. The effect of this cut at this stage in cut flow is shown in Figure 7.3.

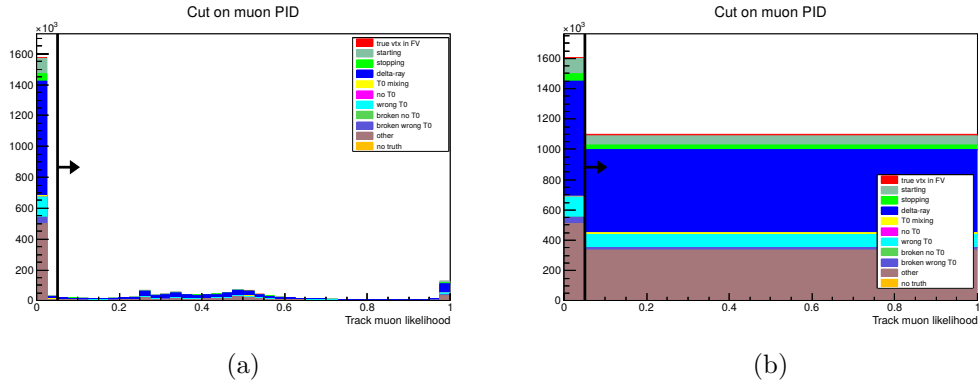


Figure 7.1: Muon PID cut in selection showing cut value and distribution of signal and background muon likelihood (Figure 7.1a) and the same distribution with binning to better illustrate signal survival (Figure 7.1b).

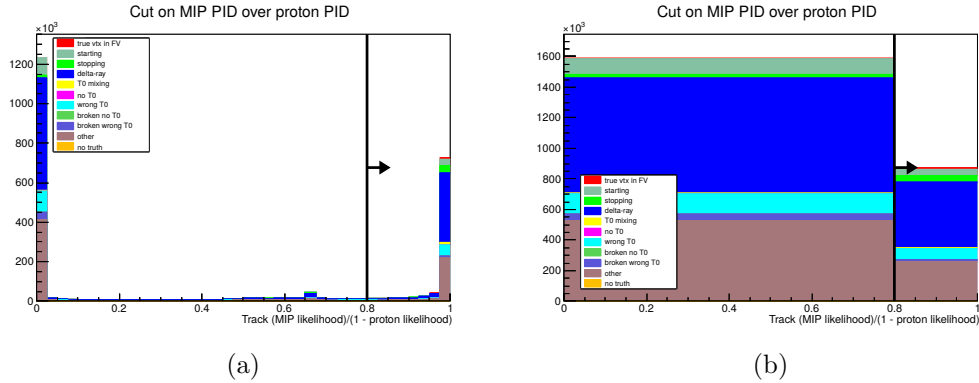


Figure 7.2: MIP PID cut in selection showing cut value and distribution of signal and background MIP likelihood values (Figure 7.2a) and the same distribution with binning to better illustrate signal survival (Figure 7.2b).

x-angle

Our other kinematic cut is motivated by detector geometry. We cut on candidate muon *x*-angle. This is motivated by our detector geometry; we are unable to reliably reconstruct tracks at high angle in *x* since they cross few micromegas pads for their length. The cut removes such tracks.

The cut removes tracks where $\cos \theta_x > 0.9$, where $\cos \theta_x$ is the *x* component of the unit vector giving the track's starting direction.

This cut is shown in Figure 7.4 at the current stage in cut flow. As with the momentum cut the region removed does feature lower signal to background ratios than shallower angles but our main motivation for the cut is based on the performance of our reconstruction.

Cut on momentum




Figure 7.3: Momentum cut in selection showing cut values.

Cut on x angle




Figure 7.4: x -angle cut in selection showing cut values.

Vertex fiducial volume

Our previous fiducial volume cut didn't account for expected vertex position. Nor did it consider pad gaps or for the difficulties in reconstructing tracks at the edge of an instrumented region. The stricter vertex fiducial volume cut accounts for all of these factors.

Projected vertex position for each candidate track takes account of the other tracks in its pattern. For single track candidates this position is simply the candidate's starting position. For multiple track events the position is based on extrapolating back from the junction's two highest momentum tracks as shown in Figure 7.5.

Reconstruction difficulties at the edge of a micromegas pad are solved by requiring projected vertex position to be at least 10 cm away from the edge of an instrumented region in x , y or z . Vertex positions are also required to be at least 3 cm away from a vertical micromegas gap since tracks near this gap are highly prone to reconstruction failures. Any global tracks whose projected vertex positions fail either of these conditions are removed.

t_0 quality

Tracks with bad t_0 are a large source of single track background. If a track enters from a side ECal but has faulty t_0 it can appear to be starting within our fiducial volume.

A sanity check removes much of this background. The requirement is relevant to tracks where the candidate vertex position is on the x edge of the pattern




Figure 7.5: Projected vertex position calculation, with tracks illustrated by blue lines and vertex position by a pink star. Default vertex position is shown in Figure 7.5a and projected position is shown in Figure 7.5b.

containing them furthest from the cathode. The pattern's t_0 source must correspond to a detector through which one of the pattern's paths passes.

Likelihood matching

Likelihood matching cuts are designed to remove backgrounds from reconstruction failures or delta rays. These can cause a through-going track to appear broken at some point along its length. The track's two halves can be matched to verify that they come from the same particle and remove the background.

The calculation of likelihood values is described in detail in ???. These are used at the analysis level to apply a stringent cut on backgrounds. If the path at the start of a candidate global track can be matched to some other path in the same TPC with a log likelihood difference of less than 1000 then the track is discarded. The log likelihood difference is defined between the one- and two-track hypotheses.

This effect of this cut at this point in the cut flow is illustrated in Figure 7.6.

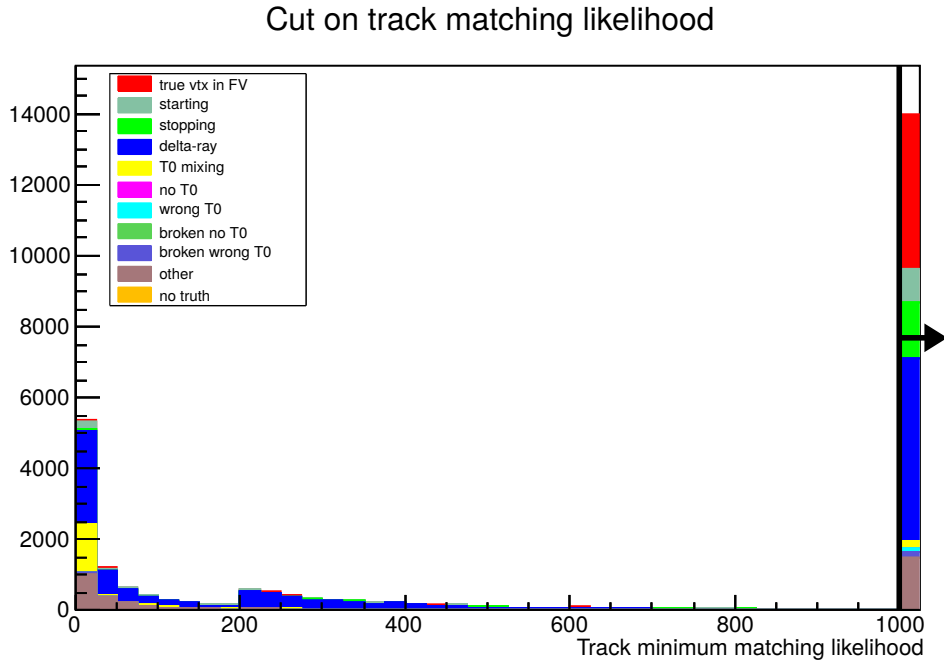


Figure 7.6: Likelihood matching cut in selection showing cut values and signal and background log likelihood difference values. The final bin contains all overflow and accounts for many orders of magnitude in difference between the likelihoods of most signal and the removed background.

TPC cleanliness

TPCs containing multiple nearby tracks from outside the TPC can create phantom vertices which are difficult to resolve through likelihood matching. To remove this background, tracks which pass by the vertex candidate are considered. These tracks are required to

- contain hits in the same TPC as the vertex
- have momentum of at least 100 MeV
- ‘pass’ the vertex (the vertex must lie between planes defined by the track’s start and end positions and directions)

in order to reject a candidate vertex from the selection.

Figure 7.7 illustrates the cut, showing the regions in which potential vertices are excluded for being alongside another track. These regions covers planes between the second track’s start and end, less some safety margin.




Figure 7.7: Illustration of passing by tracks in the selection. Figure 7.7a shows a phantom vertex candidate produced by two collinear tracks, with the misreconstructed tracks attached to the junction shown in blue and another passing by track in red. Figure 7.7b shows the area excluded for being alongside the second track.

7.3 Selection performance

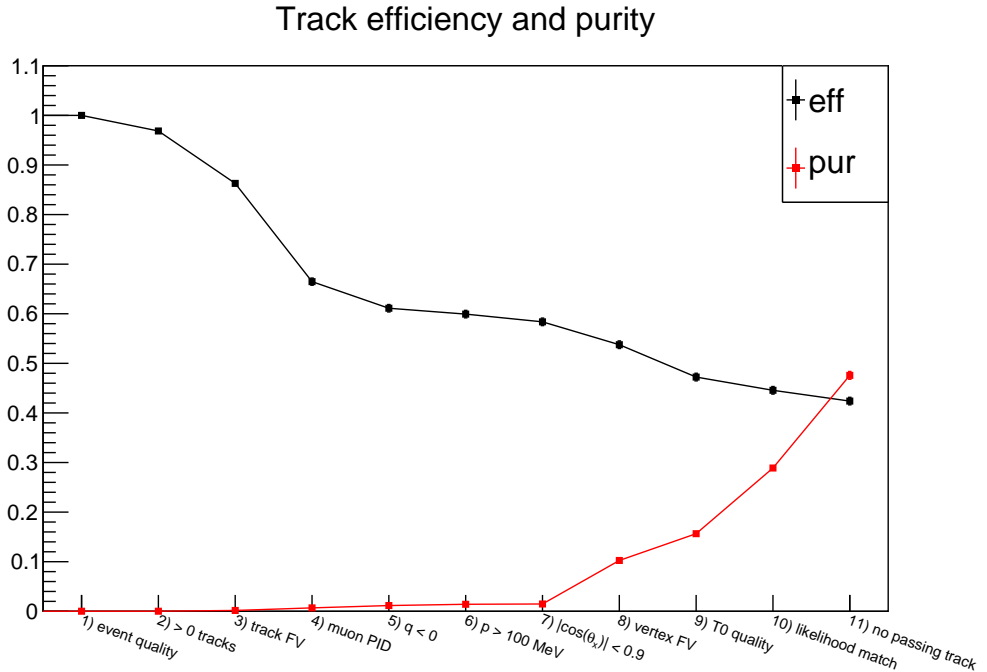


Figure 7.8: Efficiency and purity of selection cuts.

The performance of our current selection on GENIE Monte Carlo is shown in Figure 7.8. It is worth noting several points of signal loss:

- Both fiducial volume cuts and the ‘tracks exist’ cut remove signal. Since a large amount of the TPC gas volume is either uninstrumented or at the edge of a module, where vertices cannot be reliably reconstructed, such signal loss is inevitable and largely irrecoverable.
- The momentum and $\cos \theta_x$ cuts remove some signal in kinematic regions where our reconstruction is known to be unreliable.
- A true signal event may still be poorly reconstructed and hence removed due to failing PID cuts, the charge cut or the t_0 sanity check.

Overall we achieve an efficiency of $(43.4 \pm 0.8)\%$ and a purity of $(46.7 \pm 0.8)\%$, before corrections and systematic uncertainties are applied, for ν_μ charged current interactions within the TPC fiducial volume. The huge amount of out of fiducial volume background makes a higher purity impossible without either further

development on the reconstruction algorithm[†] or cuts risky cuts for which efficiency and purity on data cannot be reliably predicted. To reach our current purity we have removed about 4999 out of every 5000 background events.

In the following chapters we examine the systematic uncertainties associated with these numbers and the results which can be derived from examining data.

[†]The results of our current analysis and particularly an analysis of surviving background events in data can be used to inform which specific aspects of reconstruction need further improvement.

Chapter 8

Systematic uncertainties

Systematic uncertainties used in the gas interaction analysis fall into three broad categories:

- Those which are taken to be unchanged since pre-TREx analyses
- Those which are in principle unchanged but require recomputation
- Those which are completely new for the gas interaction analysis

In the first case some cross checks are employed to verify the validity of the old values for uncertainties. In the other cases the values are computed anew for the gas interaction analysis.

In addition to the propagation of systematic uncertainties several corrections are employed which shift the values of analysis level variables. For convenience these are also discussed in this chapter.

8.1 Corrections

Before systematic uncertainties are propagated there are several corrections applied to either data or Monte Carlo to account for known and understood differences between real data and simulation. The gas Monte Carlo and proton range corrections are newly developed for this analysis. All other corrections are implemented in other ND280 tracker analyses[179] and are unchanged in implementation.

8.1.1 Gas Monte Carlo correction

During the course of the gas interaction analysis a bug in the simulation of the gas mixture was discovered causing incorrect relative proportions of argon, isobutane

and tetrafluoroethane leading to incorrectly simulated total target masses in the TPC gas volume. To correct this each simulated event with an interaction in the TPC gas volume is given a weight according to the true identity of the target nucleus in simulation.

Nucleus	Correction
Ar	0.9555
H	1.4450
C	1.5787
F	2.1133

Table 8.1: Weights applied as corrections for relative masses of gasses in the TPC gas mixture.

The corrections are shown in Table 8.1. There are substantial corrections to the relative masses of hydrogen, carbon and fluorine and a 5% reduction in the simulated mass of argon.

8.1.2 Proton range correction

As discussed in Section 5.3 a bug in the default photoabsorption ionization model causes low momentum protons to deposit energy much slower in Monte Carlo than in reality. A correction is applied to proton ranges at the analysis level to resolve this.

The correction requires reducing the physical length of tracks by moving their end points closer to their starts, reducing the number of clusters the tracks are recorded as containing and correcting for the lower reconstruction efficiency for short tracks compared to long tracks.

Figure 8.1 shows the differences between simulated and measured ranges for low energy protons in argon. Also presented are the corrections which result from them.

8.1.3 PID corrections

Since PID determination is mostly unchanged from its pre-TREx implementation the $\frac{dE}{dx}$ corrections used for previous TPC reconstruction are still necessary. These are standard corrections implemented in general T2K analyses and stem from the need to correct for well understood limitations in the hardware, software and reconstruction.




Figure 8.1: Proton ranges in data and Geant4 simulation (Figure 8.1a) and the corrections which result from them (Figure 8.1b), derived by taking the ratio between version 4.9 simulation with the photoionisation absorption model (PAI) and Stopping Range of Ions in Matter program (SRIM)[180] data and enforcing no correction at 0 MeV or above 100 MeV.

PID corrections are as follows:

- Slight corrections for data (no more than a few percent) are applied to the measured $\frac{dE}{dx}$ in each TPC. These depend on the specific run and subrun range.
- Monte Carlo corrections of around 1% are applied to electrons. This is to account for a known overestimate of electron energy loss in simulation.
- Finally the expected $\frac{dE}{dx}$ values used to calculate pulls are corrected to account for unresolved inaccuracies introduced at the reconstruction stage.

8.1.4 Momentum resolution corrections

As with the $\frac{dE}{dx}$ corrections the momentum resolution corrections used with previous TPC reconstruction are still used in this analysis.

This correction is quite large, providing an almost 40% smearing in inverse transverse momentum. It accounts for known differences in resolution between data and Monte Carlo[181].

8.2 Propagation of systematic uncertainties

All systematic uncertainties are uniquely implemented in the analysis. The general procedure for doing so is outlined here and covers the two broad classes: variations and weights (of which efficiency-like uncertainties merit special mention).

Uncertainties based on weights only require the selection to be run once. A set of random weights within the systematic error in question is then calculated. Variation based uncertainties require multiple passes of the selection with the analysis level variable in question varied many times within its systematic error. The variation in the final number of signal and background events in a given bin post-selection gives our total systematic uncertainty for that bin.

Unless otherwise stated all sources of uncertainty are assumed to be gaussian.

8.2.1 Variations

Variation uncertainties account for uncertainty in some variable v , for example track momentum or $\frac{dE}{dx}$. The selection is repeated multiple times with this variable modified as

$$v = v_0(1 + \delta \cdot \sigma_v), \quad (8.1)$$

where v_0 is the uncorrected value and $\sigma_v \cdot \delta$ gives a random variation from the uncertainty on the systematic source in question[†].

8.2.2 Weights

Uncertainties such as target mass affect the weight for the final event but not the selection. These and efficiency systematic uncertainties are propagated by simply adding a final weight to the event,

$$w = 1 + \delta \cdot \sigma_w, \quad (8.2)$$

where $\sigma_w \cdot \delta$ gives a random variation from the uncertainty on the systematic source in question[‡].

8.2.3 Efficiencies

Many uncertainties, such as in the probability of reconstructing a true particle, correspond to a chance for a particular event to be either completely accepted or completely rejected. These contribute a total weight to each Monte Carlo event depending on whether reconstruction is successful (efficiency) or not (inefficiency) and are a special case of weight based systematic uncertainties.

[†] σ_v is the specific value of our uncertainty on the variable as a result of this systematic error and δ is a random number taken from a gaussian distribution of mean 0 and width 1.

[‡] σ_w is the specific value of our uncertainty on the weight as a result of this systematic error and once again δ is a random number taken from a gaussian distribution of mean 0 and width 1.

In either case the weight is computed from the predicted value in data $\epsilon_{\text{data}}^{\text{var}}$ and the analysis sample's Monte Carlo efficiency ϵ_{MC} . For efficiencies it is given by

$$w = \frac{\epsilon_{\text{data}}^{\text{var}}}{\epsilon_{\text{MC}}}, \quad (8.3)$$

while for inefficiencies it is given by

$$w = \frac{1 - \epsilon_{\text{data}}^{\text{var}}}{1 - \epsilon_{\text{MC}}}. \quad (8.4)$$

Specifically $\epsilon_{\text{data}}^{\text{var}}$ corresponds to the mean prediction for data plus a random value thrown from within our uncertainty on the systematic source in question which is varied between toy experiments.

The uncertainty is ideally given by the difference between efficiency (or inefficiency) in data and Monte Carlo in well understood control samples. For some gas interaction specific systematic uncertainties control samples are unavailable and we have used comparisons between nominal and modified Monte Carlo instead, where the modified Monte Carlo is set up so that true data performance in the relevant aspect of reconstruction will lie between the two.

8.3 Event level systematic uncertainties

8.3.1 Sand and cosmic muon background

At this time it is not feasible to perform a robust evaluation of sand and cosmic muon background due to practical constraints of time and data availability. To get an idea of expected cosmic background we look to previous tracker analyses while for sand muons we use study our selection's performance on a limited sample of simulated events.

Previous tracker analyses[182] predicted total cosmic ray contamination far below 0.1%, with a total out of fiducial volume background around 5%. If we conservatively[†] assume that the fraction of out of fiducial volume background from cosmic rays is the same as other analyses', contamination is well below 1% and is thus neglected.

For sand muons we do not have a full set of simulated data corresponding to $10\times$ real data (as with other Monte Carlo). Instead we pass the simulated sand events we have, corresponding to 8.9×10^{20} PoT (about 50% higher than real data

[†]Relative to other tracker analyses our out of fiducial volume background is much more likely to come from interactions near a TPC, where multiple particles can be produced and thus confuse pattern recognition.

PoT), through our selection. We find 12 ± 3.5 selected sand events which scales to a predicted 7.7 ± 2.4 in real data. Systematic uncertainty on sand muon flux is around 10%. These selected sand events and their associated error are added to our background predictions in both NEUT and GENIE.

8.3.2 Flux related systematic uncertainties

Flux related uncertainties are a broad category accounting for effects from hadron production rates to the alignment of our beams and magnetic fields. For the gas interaction analysis we use the same flux uncertainties and implementation thereof as general ND280 analyses (see for example [183]).

As shown in Figure 8.2 hadron production is the dominant contributor to our systematic uncertainty. This is constrained with data from NA61/SHINE as mentioned in Section 5.1. Other sources of error are uncertainties over the precise profile and alignment of our proton beam, the magnetic field and physical alignment of our horns and the alignment of our target. These errors are evaluated independently by varying the relevant quantity in simulation and determining its final effect on flux. All errors are ultimately combined into a covariance matrix binned in neutrino energy and type. At around 10% flux uncertainties represent the largest single source of systematic error in the gas interaction analysis.

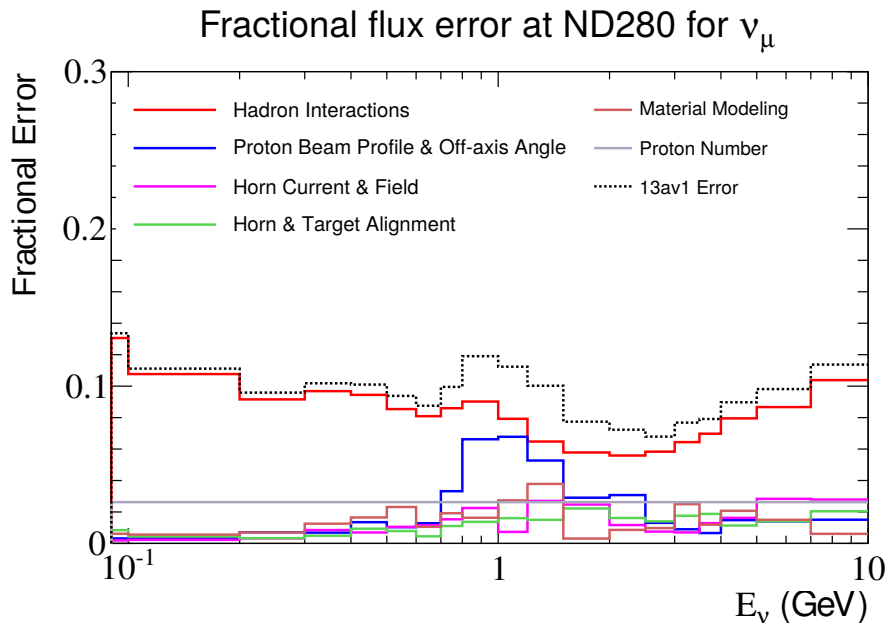


Figure 8.2: T2K fractional flux uncertainties for ν_μ [184]. 13av1 is the flux tuning used for this analysis — its total fractional uncertainty is shown by the dashed line.

8.4 Standard TPC variation uncertainties

There is a large degree of overlap between the variation based uncertainties used by our gas interaction analysis and those common to ND280 tracker analyses.

It was considered acceptable to reuse existing values for these uncertainties. This is justified by the close matching of kinematic distributions between TREx and our old reconstruction (see for example Figure 6.29) and the fact that the underlying procedures for fitting and PID were not changed for TREx .

8.4.1 Momentum scale

Uncertainties in absolute momentum scale ultimately come from uncertainties in the measurement and calibration of the ND280 B field[185]. We use the same values as pre-TREx analyses. All together our uncertainty on momentum scale is 0.57%.

The scale uncertainty is implemented as a simple variation in absolute momentum for both global and local tracks. It covers all charged tracks reconstructed in the TPC.

8.4.2 Momentum resolution

This corresponds to our uncertainty on the variation of momentum around a mean value. In practice it is the resolution of inverse momentum transverse to B field direction ($1/p_t$) for which we propagate the systematic[181].

The specific uncertainty depends on track x but generally involves a variation of 10%.

Momentum resolution uncertainty is implemented as a variation in the difference $1/p_t - 1/p_t^{\text{true}}$ between reconstructed and true inverse transverse momentum. For each toy experiment each track has its momentum translated to inverse transverse momentum, smeared by 10% of the difference with true inverse transverse momentum, then translated back.

8.4.3 B field distortion

The final momentum variation applied comes from our imperfect knowledge of distortions in the ND280 magnetic field[186] which ultimately stems from our uncertainty over magnetic yoke properties. Unlike the momentum scale error, which accounts for uncertainty on the mean field, this handles our uncertainty over deviations from this mean field and as such its computation and implementation is more involved.

The main correction to B field distortions comes from the map produced by direct measurements in the ND280 basket. Additional empirical corrections are produced from measurements from the TPC laser system. The difference between nominal track momentum and that produced by evaluating these corrections (‘refit momentum’) provides the B field distortion systematic.

As with momentum scale this uncertainty is implemented as a variation in absolute momentum of a global track and its constituents. Here the variation comes from the difference between nominal track momentum and refit momentum. This is different from track to track, even for tracks which otherwise have similar kinematics.

8.4.4 PID

Our PID uncertainties come from our uncertainty in the $\frac{dE}{dx}$ measurements used to evaluate particle identity. We reuse the values from an earlier study[187] where this is evaluated by evaluating the data–Monte Carlo difference in $\frac{dE}{dx}$ between well understood control samples of electrons, protons and minimum ionizing particles (muons and pions) as a function of momentum.

This uncertainty is implemented as a variation in the $\frac{dE}{dx}$ value for each TPC segment of each charged track reconstructed in our TPCs. The exact variation is decided based on the true particle’s identity and momentum. Electrons, protons and minimum ionizing particles use separate variations but the method is the same in each case. $\frac{dE}{dx}$ is varied based on both its mean value and the difference between mean and expected values. Pulls are recomputed after this variation (see Section 6.8).

8.5 Standard TPC efficiency uncertainties

As with variations there is a large degree of overlap between efficiency uncertainties required for the gas interaction analyses and those required for general ND280 tracker analyses. In most of these cases the overall systematic uncertainty is small, and similar performance between TREx and previous reconstruction justifies our recycling of values from pre-TREx analyses. Only track reconstruction efficiency is recomputed, due to its sensitivity to changes in pattern recognition.

8.5.1 Track efficiency


This covers our uncertainty over the chance for a charged particle which passes through the TPC’s instrumented region to be successfully reconstructed as a track. Track efficiency encapsulates both pattern recognition and likelihood fitting efficien-

cies. The methodology for computing and propagating this systematic is unchanged since pre-TREx analyses[188]. Specific values have, however, been recomputed for this analysis.


The efficiencies are tabulated in Table 8.2 and shown in Figure 8.3. Both data and Monte Carlo show high levels of efficiency which are equivalent within statistical error.

TPC	Data efficiency	Monte Carlo efficiency
TPC1	$(99.6 \pm 1.0) \%$	$(98.9 \pm 1.6) \%$
TPC2	$(99.4 \pm 0.6) \%$	$(99.3 \pm 0.7) \%$
TPC3	$(99.1 \pm 2.0) \%$	$(99.0 \pm 1.0) \%$

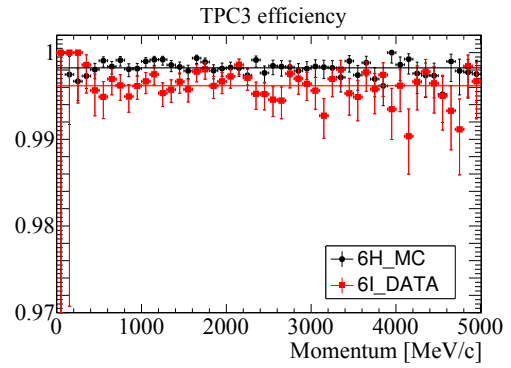
Table 8.2: TPC track finding efficiencies for data and Monte Carlo.



(a)



(b)



(c)

Figure 8.3: TPC efficiency as a function of momentum for TPC1 (Figure 8.3a), TPC2 (Figure 8.3b) and TPC3 (Figure 8.3c). The average for each TPC is shown by the horizontal black line (Monte Carlo) and red line (data).

An efficiency weight is applied based on the event’s main true track (a true muon passing through the TPC). If this true track corresponds to a reconstructed track the weight is applied as an efficiency; otherwise it is applied as an inefficiency (see subsection 8.2.3). These weights are around 1%.

8.5.2 Cluster efficiency

Cluster efficiency uncertainty comes from our uncertainty over the probability of missing one or more horizontal or vertical clusters in an otherwise reconstructed track.

The values used for this systematic uncertainty come from work done prior to the introduction of TREx[189]. In this study efficiency was compared between data and Monte Carlo for well understood samples of through-going muons in the horizontal and vertical directions.

$\cos \theta$	Extra Monte Carlo efficiency
0–0.5735	$(0.11 \pm 0.02) \%$
0.5735–1	$(0.07 \pm 0.01) \%$

Table 8.3: TPC cluster finding efficiency systematic as a function of angle in the y - z plane. Extra Monte Carlo efficiency corresponds to $(\epsilon_{MC} - \epsilon_{data}) / \epsilon_{MC}$ where ϵ_{MC} is Monte Carlo efficiency and ϵ_{data} data.

Results from this study are shown in Table 8.3. They represent our uncertainty over the chance of a single cluster being successfully reconstructed. Depending on the length of our track the loss or gain of a single cluster may not impact the analysis. For our selection it is assumed to be relevant only to the TPC quality cut, which requires 18 or more clusters. Other effects of missing a cluster, such as shifting track position or kinematics, are accounted for by other systematic uncertainties. This means that the uncertainty applies only to tracks with near 18 clusters — for the gas interaction analysis it is entirely negligible.

Only muon candidate tracks with 17 or 18 clusters are considered. For each of these the weight on a single cluster being accepted or rejected is given. This is placed to the power of the number of clusters to compute the weight from the chance of any one being lost and multiplied by the extra Monte Carlo inefficiency to account for the chance of any extra cluster being gained.

8.5.3 Charge ID efficiency

Charge ID uncertainties account for uncertainty on the chance of a muon being misidentified as its antiparticle due to incorrectly reconstructed charge. This stems from both local mis-ID (a track segment being assigned the wrong charge) and global mis-ID (a global track being assigned different charge to its best local segment).

The uncertainty is correlated strongly with the fitting error on track momentum. It also depends on the number of TPC segments in a track and their relative charge sign. Tracks with one, two and three segments require different errors as do tracks where charge sign disagrees between segments.

The efficiency is derived from a relatively complex parametrisation in momentum fitting error[190]. Since the momentum fit is in principle unchanged in **TREx** we reuse the values and propagation methodology used in previous tracker analyses. In spite of their relative complexity these systematic uncertainties contribute a final uncertainty in the gas interaction analysis of less than 1%.

Our charge ID uncertainty is propagated as an efficiency based on the muon candidate track. The precise amount by which our weights are varied depends on the muon candidate track's momentum fit error, the number of TPC segments and their charges relative to each other and relative to the global track's charge.

8.6 Gas interaction specific uncertainties

Many of the uncertainties inherent in a gas interaction analysis have not been computed before either because they are only relevant to vertices within the TPC or because they didn't have a significant effect in other analyses. We evaluate these from scratch.

Due to a lack of control samples for gas interactions many rely on Monte Carlo only studies. Since these cannot be relied on to precisely replicate data they provide conservative upper limits on our uncertainty.

8.6.1 Vertex uncertainties

These address our uncertainty over the probability of correctly identifying vertices and their properties within the TPC.

Since reconstruction differs substantially when it comes to identifying vertices in one-, two- and multi-track events the parameters for this uncertainty are binned in charged particle multiplicity. For two-track interactions vertices are identified through kink finding and in the case of multiple tracks they are identified through

the junction finding algorithm, both of which are discussed extensively in Chapter 6. The one-track case is ignored since our uncertainty for this is covered by the track efficiency uncertainty.

Due to lack of a viable control sample in data and Monte Carlo, a Monte Carlo only study was used. Nominal simulation was compared with a ‘noisy’ sample in which the variation in simulated pedestal noise was doubled and zero suppression halved to create extra noise hits. As illustrated in Figure 8.4 even such a large change in simulation results only in a slight broadening of tracks but the difference between the two can be taken to cover extra performance of data.

Three values were checked for the purposes of quantifying our uncertainty; muon efficiency, vertex resolution and secondary charged particle multiplicity.

The muon efficiency uncertainty is on the chance of reconstructing a true muon originating within the TPC’s instrumented region. It was computed using the difference in reconstruction efficiency between nominal and noisy Monte Carlo samples. Results from computing this uncertainty are shown in Figure 8.5.




Figure 8.4: Normal (Figure 8.4a) and sub-optimal (Figure 8.4b) version of the same gas interaction event. Also shown is a composite with new noisy hits generated by the sub-optimal sample highlighted in orange (Figure 8.4c).

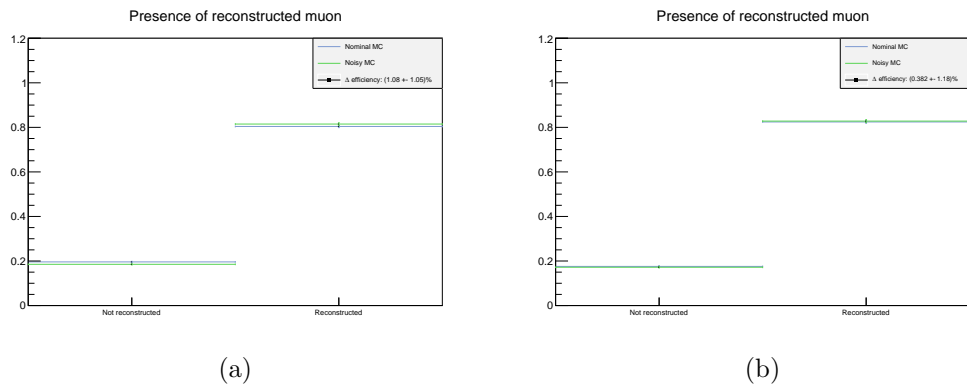


Figure 8.5: Vertex efficiency systematic uncertainties for two- (Figure 8.5a) and multi-track (Figure 8.5b) events.

The muon resolution uncertainty is our uncertainty over the resolution with which a vertex's position can be correctly identified. The difference between true and reconstructed vertex positions in each of the three dimensions was recovered and fitted with a gaussian. The difference in width of these distributions was taken between our nominal and noisy Monte Carlo samples. The results are illustrated in Figure 8.6.

The vertex multiplicity uncertainty represents our uncertainty on the chance of reconstructing the correct number of secondary charged particles emerging from a vertex. This is heavily dependent on the length of each proton track, which requires a correction at the analysis level. As a result of this it was decided that the best way to model the uncertainty would be to evaluate our uncertainty on secondary particle reconstruction efficiency as a function of true trajectory length and propagate it based on the event's corrected trajectory lengths. These uncertainties are illustrated in Figure 8.7. They are dominated by statistical uncertainties in evaluating them.

Table 8.4 summarises our vertex systematic uncertainties. For the purposes of our analysis the most significant effects are muon efficiency, for which uncertainty is of the order 1%–2%.

Muon efficiency uncertainty is propagated analogously to TPC track efficiency uncertainty. The only difference is the binning in track multiplicity. Uncertainty in resolution is accounted for simply by varying the positions of our vertices

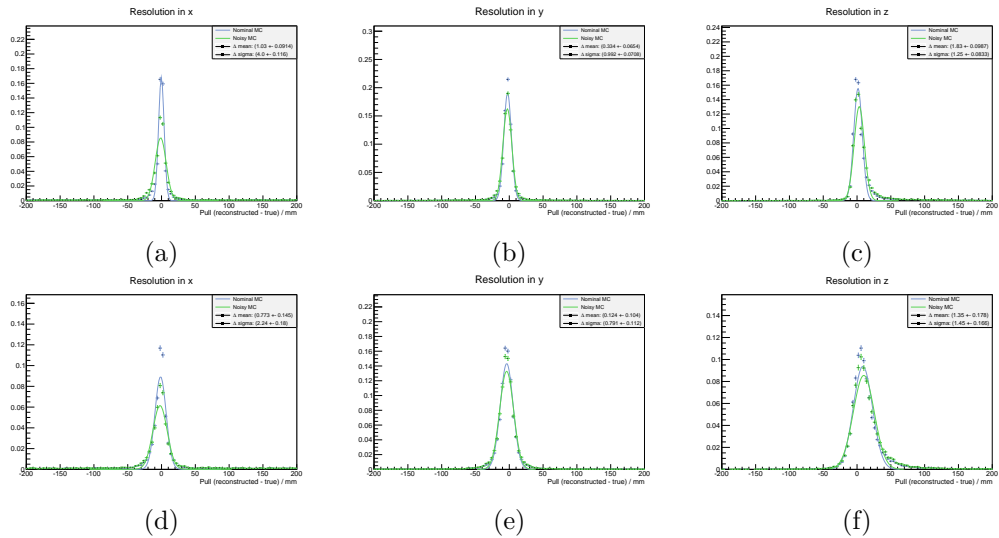


Figure 8.6: Vertex resolution systematic uncertainties for two- (Figure 8.6a, Figure 8.6b, Figure 8.6c) and multi-track (Figure 8.6d, Figure 8.6e, Figure 8.6f) events in the x , y and z directions.

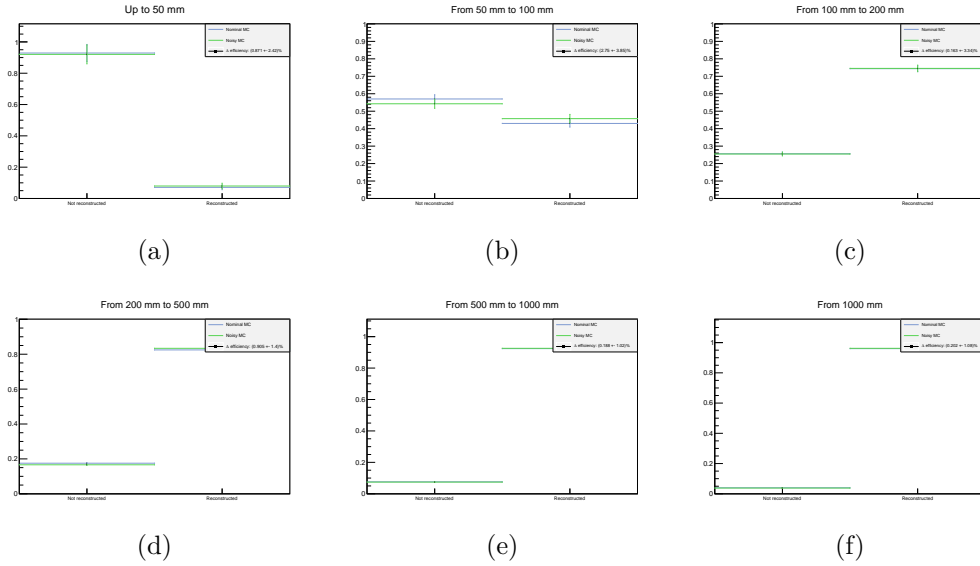


Figure 8.7: Secondary particle reconstruction efficiency uncertainties for true trajectory length up to 50 mm (Figure 8.7a), 50 mm – 100 mm (Figure 8.7b), 100 mm – 200 mm (Figure 8.7c), 200 mm – 500 mm (Figure 8.7d), 500 mm – 1000 mm (Figure 8.7e) and 1000 mm upwards (Figure 8.7f).

Systematic		Difference
Muon efficiency	Two-track	$(1.08 \pm 1.05)\%$
	Multi-track	$(0.38 \pm 1.18)\%$
Muon resolution	Two-track (x)	(4.00 ± 0.12) mm
	Two-track (y)	(0.99 ± 0.07) mm
	Two-track (z)	(1.25 ± 0.06) mm
	Multi-track (x)	(2.24 ± 0.18) mm
	Multi-track (y)	(0.79 ± 0.11) mm
	Multi-track (z)	(1.45 ± 0.17) mm
Secondary particle efficiency	0 mm – 50 mm tracks	$(0.87 \pm 2.42)\%$
	50 mm – 100 mm	$(2.75 \pm 3.85)\%$
	100 mm – 200 mm	$(0.16 \pm 3.34)\%$
	200 mm – 500 mm	$(0.91 \pm 1.40)\%$
	500 mm – 1000 mm	$(0.19 \pm 1.02)\%$
	1000 mm or more	$(0.20 \pm 1.08)\%$

Table 8.4: Summary of vertex related uncertainties.

and track starts in our toy experiments based on the uncertainties calculated.

Our secondary particle multiplicity uncertainty doesn't affect the main selection but is relevant when plotting proton multiplicity. It is propagated by randomly deleting secondary particles in each toy experiment. This is done by calculating a varied efficiency weight, $\epsilon_{\text{data}}^{\text{var}}$, for each toy. A random number δ is thrown between 0 and 1. Protons where $\delta < \epsilon_{\text{data}}^{\text{var}}$ are then removed.

8.6.2 t_0 determination

The efficiency uncertainty for t_0 determination represents uncertainty over the probability of a t_0 being found from a source (P0D , FGD or tracker ECal) through which the track in question passes. For more detail on the t_0 algorithm see Section 6.6. The uncertainty is computed by comparing efficiency for each source in data and Monte Carlo.

'Clean' events are sought. These are events where a unique single track pattern in a given TPC covers its entire length in x , y or z depending on the detector being checked. Hits are also required in nearest layers to the TPC of the two detectors either side. For z tracks these are either the P0D and FGD1, FGD1 and FGD2 or FGD2 and the downstream ECal. For x and y tracks they are the two side ECals and the top and bottom ECals respectively. Categorisation of these candidate events is illustrated in Figure 8.8.

The comparison between expected t_0 determination and actual t_0 determination shown in Figure 8.9 provides an efficiency for each detector. These are compared between data and Monte Carlo to extract uncertainties on t_0 efficiency.

Our computed t_0 efficiency uncertainties are shown in Table 8.5. Those for top and bottom and side ECals are notably higher than those from other sources. This is thought to be a result of the methodology used for evaluating the uncertainty overestimating candidate t_0 sources in data for these detectors. A more robust study could produce lower values but on our current time scale these conservative values are considered acceptable.




Figure 8.8: t_0 candidates in x (Figure 8.8a), y (Figure 8.8b) and z (Figure 8.8c). The red pattern shown in Figure 8.8c is not a candidate since it is not a unique single track pattern.

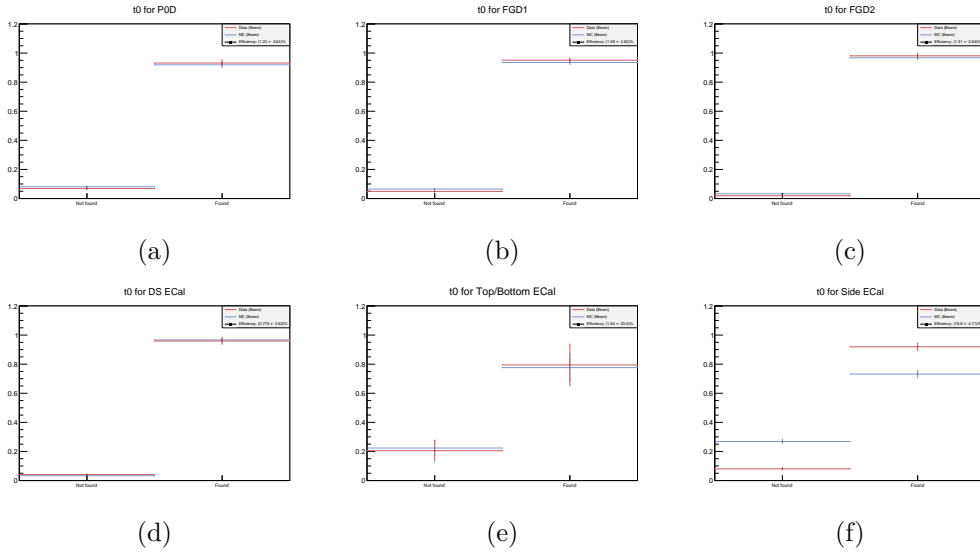


Figure 8.9: t_0 for tracker detectors, i.e. the POD (Figure 8.9a), FGD1 (Figure 8.9b), FGD2 (Figure 8.9c) and downstream (Figure 8.9d), top and bottom (Figure 8.9e) and side (Figure 8.9f) ECal, using beam data and Monte Carlo.

t_0 source	Efficiency difference
POD	$(1.23 \pm 3.64) \%$
FGD1	$(1.59 \pm 2.62) \%$
FGD2	$(1.31 \pm 2.64) \%$
Downstream ECal	$(0.78 \pm 3.62) \%$
Top and bottom ECals	$(1.8 \pm 23.0) \%$
Side ECals	$(18.8 \pm 14.7) \%$

Table 8.5: t_0 efficiency difference between data and Monte Carlo for different sources based on comparing data and Monte Carlo for samples of through-going tracks.

An efficiency uncertainty is applied to any event where the main track may have been rejected due to failing the t_0 cut. This includes any pattern where the main track's starting x position may have been moved out of the fiducial volume but excludes patterns where all possible vertex x positions are within the fiducial volume.

A candidate t_0 source is selected by choosing the t_0 source used by the track. If no valid source is found then the first detector through which the true track passes is used. Once a source is established an efficiency weight is applied to the event based on the systematic uncertainty from this source.

8.6.3 Hairy track systematic uncertainties

One of the major issues identified by the proto-analysis was the existence of ‘hairy tracks’, or low momentum proton tracks with large amounts of noise surrounding them. These are discussed in detail in Section 5.4 and subsection 6.3.3. Although reconstruction procedures were implemented to resolve them the demonstrable differences in event appearance between data and Monte Carlo necessitate a large systematic.

No viable control samples are available for these event. Some proton samples were used while improving and testing reconstruction hairy tracks but these are low in statistics and do not allow us to check tracks starting inside the TPC, which are expected to be highly sensitive to the presence of hairy tracks. For these reasons a Monte Carlo only study was used.

As seen in Figure 8.10 our current reconstruction handles hairy topologies fairly well We can in most of the worst cases reconstruct a hairy particle’s trajectory with good accuracy.

To compute our uncertainty the approach taken was to compare reconstruction efficiency and resolution between a nominal Monte Carlo sample and an artificially hairy sample created by filling every electronics unit which the lowest energy proton track passed through with low charge hits.

As seen in Figure 8.11 this sample produces far worse reconstruction performance than any hairy event witnessed in data (compare with Figure 8.10). Since our goal here was to produce a conservative, ‘worst case scenario’ uncertainty covering any unexpected drop in performance of TReX in the presence of hairy gas interactions this was considered appropriate.

As with vertex efficiency, this provides the uncertainty regarding the chance




Figure 8.10: Two hairy events; a ‘chainsaw’ event (Figure 8.10a) and a ‘low charge shadow’ event (Figure 8.10b). In each case green and red blocks represent used and unused hits respectively and blue lines represent reconstructed tracks (the single blue line in Figure 8.10a is partially obscured by unused hits).

of reconstructing a muon. In this case there are two prerequisites:

- The existence of a muon starting in the TPC.
- The existence of a hairy proton from the same vertex.

Results are shown in Figure 8.12. They demonstrate a surprisingly good muon reconstruction efficiency for our artificially hairy sample, with only 7% of muons being lost on the presence of a highly hairy proton track.

The hairy track vertex resolution systematic is analogous to the vertex resolution systematic. Our results, shown in Figure 8.13, show expected smearing of vertex position in the hairy sample. This effect is of around 10 mm or roughly the size of a single micromegas pad on average.

Hairy muon systematic uncertainties are summarised in Table 8.6. The most significant effect for our analysis is the muon efficiency uncertainty resulting from hairy tracks.

Muon efficiency and resolution for hairy tracks were propagated analogously to their analogues for general vertex systematic uncertainties. The difference is that they are only implemented for vertices featuring at least one proton below 500 MeV. This cut-off was established because we know to a reasonable level of certainty that reconstruction well handles protons above this energy in data.

Further restrictions on events for which hairy systematic uncertainties are considered may be applied in future iterations of the analyses, since we know from experience that most protons below 500 MeV are not hairy. For the time being a conservative uncertainty is used which, though our biggest source of detector related uncertainty, is still much smaller than flux related uncertainties.




Figure 8.11: Normal (Figure 8.11a) and artificially hairy (Figure 8.11b) version of the same gas interaction event.

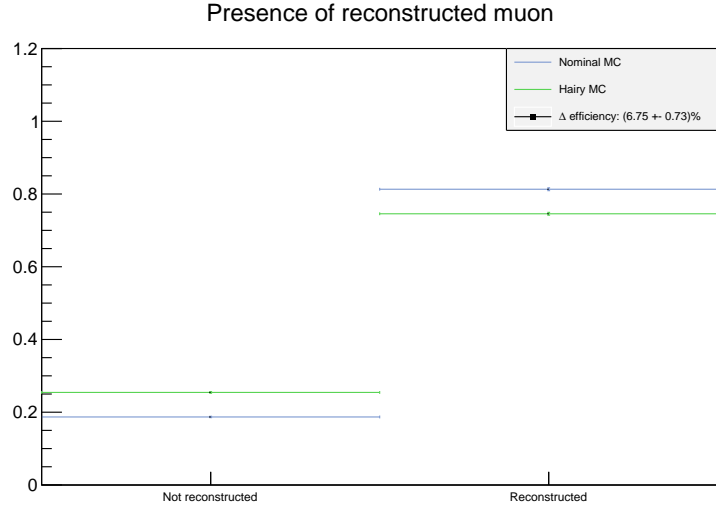


Figure 8.12: Hairy track muon efficiency for nominal and hairy Monte Carlo, with the resulting systematic shown.




Figure 8.13: Vertex resolution systematic uncertainties for hairy track events (Figure 8.13a, Figure 8.13b, Figure 8.13c) in the x , y and z directions.

Systematic		Difference
Muon efficiency	All tracks	$(6.75 \pm 0.73)\%$
Muon resolution	x	(0.64 ± 5.25) mm
	y	(0.79 ± 9.21) mm
	z	(2.55 ± 11.40) mm

Table 8.6: Summary of hairy track related uncertainties.

8.7 Out of fiducial volume uncertainties

Background in the gas interaction analysis comes almost exclusively from tracks entering the TPC from other detectors. The uncertainty in this background is split into two sources:

- A *rate uncertainty* on the total number of tracks entering from outside the TPC
- A *reconstruction uncertainty* on the probability for a track entering from outside the TPC to be reconstructed as a vertex within the TPC

Rate uncertainties come from uncertainties in the number of neutrino interactions in the detector material surrounding our TPCs. Reconstruction uncertainties can come from uncertainty on either the chance for a side-entering track to be shifted into the fiducial volume by bad t_0 or the chance for a through-going track to be broken.

8.7.1 Rate uncertainties

All of the rate uncertainties we use have been evaluated for other ND280 tracker analyses. They reflect our incomplete knowledge of the mass and neutrino interaction cross sections of target material surrounding the TPCs. We reuse the numbers acquired for previous tracker analyses[179] for these uncertainties.

8.7.2 Bad t_0 uncertainties

Tracks entering the TPCs through the sides can be shifted into the fiducial volume by a bad t_0 value. For these uncertainties the side ECal t_0 uncertainty discussed in subsection 8.6.2 is used. This accounts for uncertainty on the chance for a particle entering the TPC via a side ECal to have its t_0 correctly assigned. Since such tracks will often also be able to establish t_0 from other detectors this gives a somewhat conservative uncertainty.

This applies to tracks where the reconstructed vertex is on the pattern edge nearest to a TPC and where a true track passes into the TPC through the same side. It is propagated as an efficiency.

8.7.3 Broken track uncertainties

Broken track uncertainties come from our uncertainties in the chance for a through-going TPC track to be reconstructed with a vertex in the fiducial volume. To

compute these, a study was performed using the same sample of through-going muons as the track efficiency systematic (subsection 8.5.1).

Through-going tracks were checked in both data and Monte Carlo for the presence of reconstructed junctions. These could stem from, for example, delta rays in the TPC. The number of tracks with at least one reconstructed junction was compared between data and Monte Carlo. As can be seen in Figure 8.14 there is a high amount of agreement between data and Monte Carlo on the number of reconstructed junctions. The overall broken track uncertainty is 0.4%.

These uncertainties are applied to all out of fiducial volume tracks for which the bad t_0 uncertainty (subsection 8.7.2) does not apply. For each of these tracks a constant 0.4% efficiency uncertainty systematic is applied.

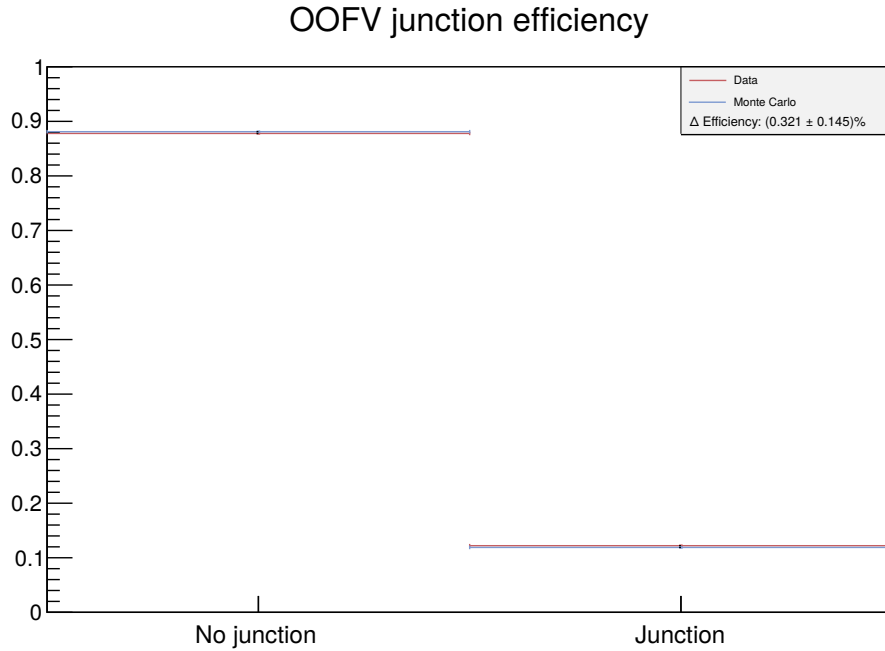


Figure 8.14: Broken track uncertainties comparing data and Monte Carlo.

Chapter 9

Argon Cross Section

This analysis was performed blind, with both selection and systematic uncertainties decided before data was observed. The results are presented here.

9.1 Single bin analysis

The main goal of this first iteration of a gas interaction analysis is to determine the charged current cross section for ν_μ interactions on argon, integrated over the T2K flux.

9.1.1 Methodology

Our cross section is given by

$$\sigma_{\nu_\mu \text{CC}}^{40\text{Ar}} = \frac{N}{\epsilon^{\text{exp}} \phi T}, \quad (9.1)$$

where N is our number of signal events in data, ϵ^{exp} the expected selection efficiency, ϕ the integrated T2K ν_μ flux at the ND280 and T the number of target nucleons.

We cannot directly measure N since our real data contains some unknown level of background contamination. In order to derive it we have opted for a purity correction method using

$$N = N^{\text{meas}} \frac{N_{\text{signal}}^{\text{exp}}}{N^{\text{exp}}}, \quad (9.2)$$

where N^{meas} is our total measured number of surviving events and $N_{\text{signal}}^{\text{exp}}$ and N^{exp} our expected signal events and total expected measured events respectively, as derived from our Monte Carlo.

The expected efficiency ϵ^{exp} corresponds to the fraction of actual signal interactions (that is, charged current ν_μ interactions on argon) picked up by our sample. This is also evaluated through our Monte Carlo.

The integrated flux ϕ is calculated by integrating the total flux at the ND280 over all of the spills we used. Flux values for each run are provided by the T2K beam group. For the total of 5.73×10^{21} protons on target used by our analysis the integrated value is

$$\phi = 1.04 \times 10^{18} \text{ m}^{-2}. \quad (9.3)$$

Finally T is the number of target nucleons in our fiducial volume. The definition for ‘fiducial volume’ is identical to that used to evaluate expected efficiency and purity. It corresponds to a total area of 6.53 m^3 . Given an ^{40}Ar density in our TPC of $(1.58 \pm 0.04) \text{ kg m}^{-3}$ we have a total of

$$T = (1.55 \pm 0.24) \times 10^{26}, \quad (9.4)$$

target nucleons.

9.1.2 Predictions from Monte Carlo

GENIE simulations corrected for sand muons predict 318 ± 44 events surviving all selection cuts. Taking systematic uncertainties and corrections into account, for the selection of charged current ν_μ interactions on argon it gives an efficiency of $(43.4 \pm 6.2)\%$ and purity of $(45.6 \pm 6.5)\%$ taking both statistical and systematic uncertainties into account.

NEUT obtains similar results but with fewer background events. 303 ± 42 events were spotted in Monte Carlo with a predicted efficiency of $(42.4 \pm 6.1)\%$ and a purity of $(49.5 \pm 7.1)\%$.

In data we have 258 events surviving our selection. Monte Carlo predictions are compared with this number of measured events in Figure 9.1. Whilst the number of events in both NEUT and GENIE overshoots data their values differ by around 1σ (taking both statistical and systematic uncertainties into account).

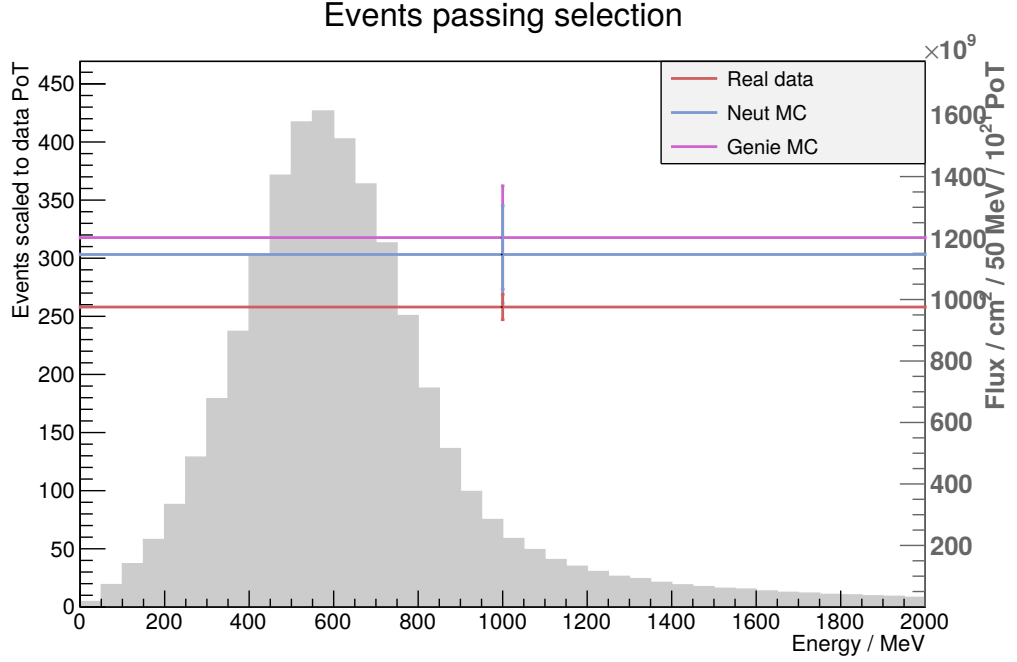


Figure 9.1: Selected events in real data, GENIE and NEUT. Errors on data are statistical only whilst those on Monte Carlo are mostly systematic. Also shown is the T2K flux.

9.1.3 Cross section calculation

Using our 258 measured data events expected purity corrections and efficiency from GENIE Monte Carlo we have a final cross section value of

$$\sigma_{\nu_{\mu}\text{CC}}^{40\text{Ar}} = (4.19 \pm 0.70(\text{stat}) \pm 1.04(\text{sys})) \times 10^{-39} \text{ cm}^2, \quad (9.5)$$

for charged current ν_{μ} interactions on argon integrated over the T2K flux.

NEUT's predicted efficiency and purity provide a similar result of

$$\sigma_{\nu_{\mu}\text{CC}}^{40\text{Ar}} = (4.61 \pm 0.78(\text{stat}) \pm 0.94(\text{sys})) \times 10^{-39} \text{ cm}^2. \quad (9.6)$$

We can naïvely extrapolate the ArgoNeuT cross section (Equation 3.1) to T2K energies by multiplying their $\sigma/E_{\nu_{\mu}} = (0.66 \pm 0.09) \times 10^{-38} \frac{\text{cm}^2}{\text{GeV}}$ measurement [141] by T2K's 0.6 GeV and weighting the result by the relative charged current cross sections at T2K and ArgoNeuT energies shown in Figure 2.3a. This gives us a loose prediction for a charged current inclusive ν_{μ} cross section of $(4.7 \pm 0.6) \times 10^{-39} \text{ cm}^2$, consistent with our results.

9.2 Distributions in other variables

As well as a single bin cross section measurement we examined distributions in other variables between data and Monte Carlo. Systematic uncertainties are less thoroughly evaluated for these than for the single bin analysis so the data should be interpreted with care. Nonetheless these distributions are useful for two reasons:

- Differences between generators can be resolved by determining which best matches data. In this case the results presented should be taken as strictly preliminary and to be refined in future generations of gas interaction analysis.
- Areas where there are notable differences between generators highlight targets for future development. We have already identified improved reconstruction of low energy protons as a target for next generation **TREx** software for this reason.

9.2.1 Muon kinematics

The kinematic properties of our muon candidate were evaluated. We expect these measurements to be somewhat reliable given our reliable simulation of muons and propagation of systematic uncertainties relating to muon momentum.

Muon momentum is shown in Figure 9.2. When both statistical or systematic uncertainty are taken into account most of the bins agree within 1σ .

The angles of our muons relative to the forward beam direction are also shown (Figure 9.3). Again most of our bins agree within 1σ between real data and simulation. There does appear to be a slight but consistent deficit in muons in the highly forwards direction. This could hint at a higher ratio of out of fiducial volume background relative to true signal events in data compared with Monte Carlo[†]. Nonetheless these differences are not hugely significant when taking statistical and systematic uncertainties into account.

9.2.2 Secondary particle kinematics

Little development has so far been done with the specific goal of accurately reconstructing secondary particle kinematics in real data, nor on the systematic uncertainties of such measurements in simulation. This is a long term goal for the gas interaction analysis. The raw distributions presented here give an indication of what we are aiming to precisely quantify in the future.

[†]Signal events are expected to produce predominantly forwards going muons whilst background features many muons entering from the sides of the TPC.

Events passing selection by momentum

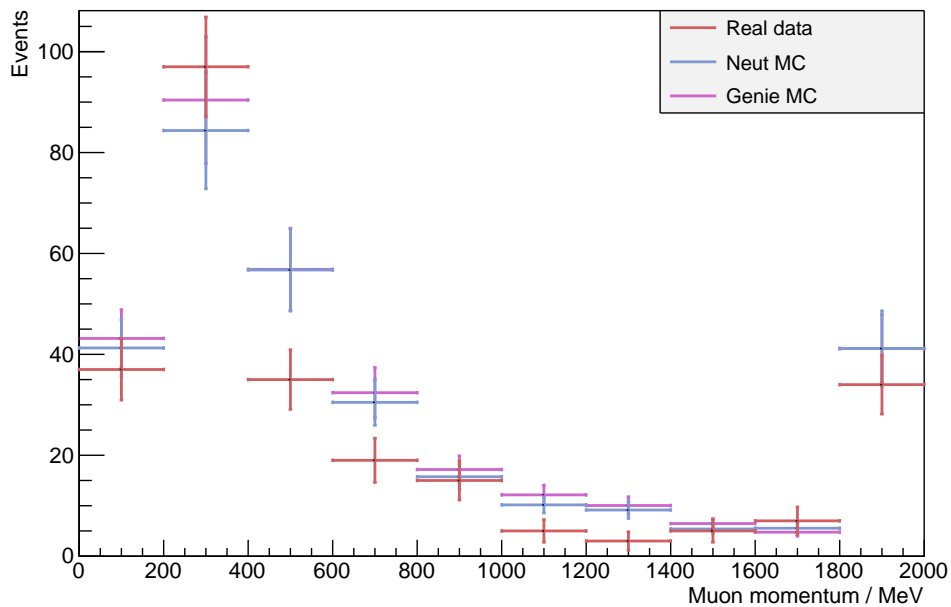


Figure 9.2: Selected muon momentum in real data and Monte Carlo.

Events passing selection by muon angle relative to beam axis




Figure 9.3: Selected muon angle relative to the beam axis in real data and Monte Carlo.

These quantities are particularly interesting because our gas TPC allows for a comparison between generators over variables which have not yet been precisely measured in experiment. In particular there is no experimentally verified model for MEC final states, and precise measurements of secondary particles and their kinematics present a chance to rectify this. GENIE generally favours events with a larger number of protons than NEUT, each carrying a lower proportion of the event's total momentum. GENIE also predicts a much larger number of protons with energy between 50 MeV and the Fermi momentum at 250 MeV whilst NEUT predicts a sharp drop off in energies below around 250 MeV. Presently no data exists for tuning the generators in this low proton energy region.

We restrict ourselves to measuring particles of 100 MeV momentum or more, corresponding to ranges starting at around 30 cm. This is due primarily to our currently inability to trust PID algorithms at low energies. It also cuts out short tracks for which systematic uncertainties are presently still uncertain.

We checked the overall number of paths emerging from each candidate vertex (Figure 9.5). In doing this we found notably fewer high multiplicity events in real data. A possible non-physics cause for this is additional noise obscuring the true number of tracks emerging from a vertex. Despite our handling of hairy events it can still be very difficult to both accurately reconstruct the number of tracks emerging from a messy vertex and provide the successful fit needed to evaluate the track's momentum. This is something which will require further study while developing TREx for future analyses.

As shown in Figure 9.5 we also checked the number of selected paths with proton PID (defined in both real data and Monte Carlo as a proton likelihood $\mathcal{L}_p > 0.9$; see Section 7.2.1). Results here are similar to the case of total path multiplicity, with substantially fewer particles identified in data than in Monte Carlo. This supports the hypothesis that unsimulated effects of messy topologies are causing us to lose candidate secondary particles. If a vertex is difficult enough that we cannot reliably reconstruct a track as having over 100 MeV momentum it also makes sense that we cannot reliably assign it a good PID value.

Finally we checked the highest momentum of our candidate protons. This is illustrated in Figure 9.6. The normalisation is much smaller for real data than Monte Carlo simply because so many fewer events were identified as containing even one proton.

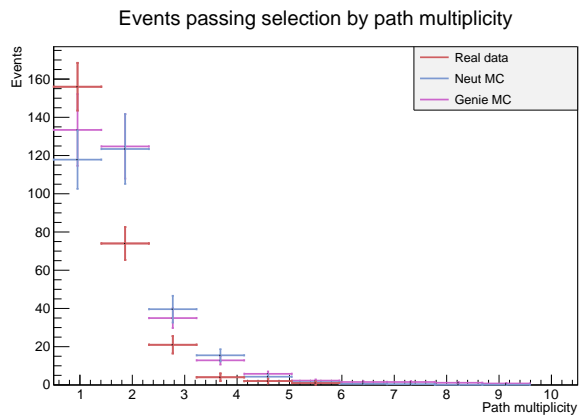


Figure 9.4: Selected vertex multiplicity, counting paths of 100 MeV or more, in real data and Monte Carlo.

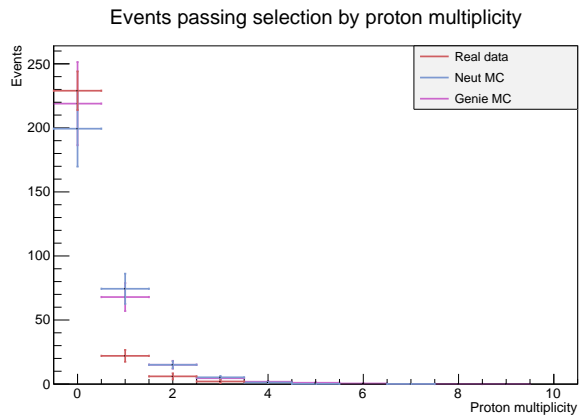


Figure 9.5: Selected vertex multiplicity, counting paths of 100 MeV or more which also have proton PID, in real data and Monte Carlo.

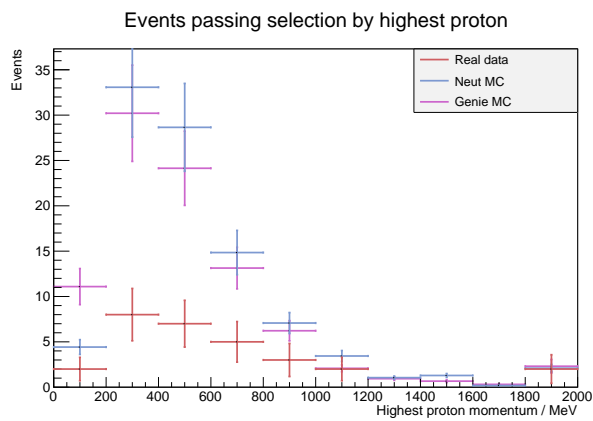


Figure 9.6: Selected highest proton momentum, from among paths of 100 MeV or more which also have proton PID, in real data and Monte Carlo.

9.3 Hand scan

A hand scan was performed on data events. This was primarily to identify interesting signal events but also functioned as a cross check for our purity prediction as well as providing indicators of the biggest sources of real background.

9.3.1 Methodology

Events were categorised as ‘signal’, ‘background’ or ‘undetermined’ with additional notes made on background topology and visible multiplicity. The hand scan was primarily done using hit level information although reconstructed tracks were also used while determining why some background events were accepted.

Events categorised as ‘signal’ featured correctly reconstructed tracks of muon PID clearly starting within the TPC gas volume. No attempt was made to distinguish between events where the muon started and stopped within the TPC. ‘Background’ events constituted any interactions clearly starting outside the TPC. In addition many events were seen where it was not clear whether the signal or background category was appropriate, generally because it was unclear whether the interaction occurred in a TPC or in the solid material of the cathode or detector frame. These were marked as ‘undetermined’.

Some examples of undetermined events are shown in Figure 9.7. Mostly these are cases where tracks appear to start within the TPC at first glance but are close enough to the edge of our fiducial volume that they have a substantial chance of originating outside.

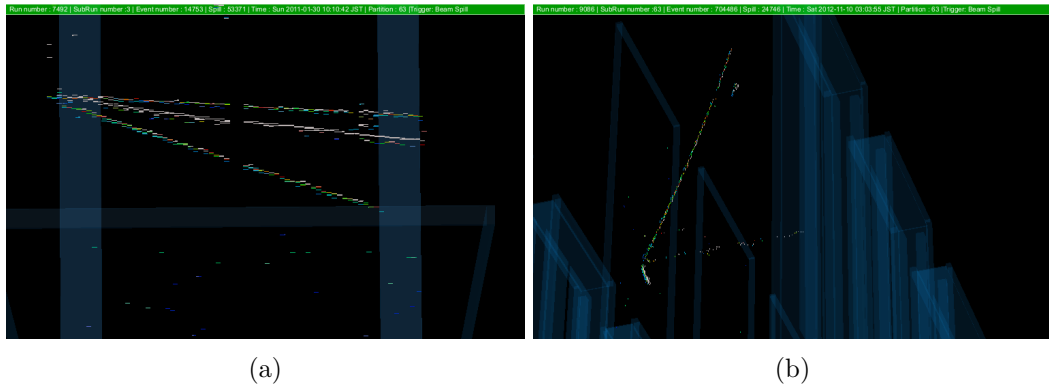


Figure 9.7: Event displays for undetermined events. Figure 9.7a shows an event at the far upstream end of a TPC whilst Figure 9.7b shows an event directly next to the central cathode.

9.3.2 Results

Category	Proportion of events	Number of events
Signal	36.4%	94
Background	57.4%	148
Undetermined	6.2%	16
Total		258

Table 9.1: Signal and background events found by hand scan.

The results of our hand scan are shown in Table 9.1. If we assume a systematic uncertainty provided by the number of ‘undetermined’ events then we have a real data purity of

$$\frac{N_{\text{signal}}^{\text{meas}}}{N_{\text{meas}}} = (39 \pm 5)\%. \quad (9.7)$$

Taking all of our computed errors into account the purity is within 0.8σ of the GENIE predicted purity and 1.2σ of the NEUT predicted purity, slightly below both estimates.

For signal events we also quantified the visible multiplicity of charged secondary particles. Results of this check for multiplicity are shown in Figure 9.8. As expected, the majority of events appear to have two charged tracks — one muon and one secondary track — with a notable tail extending to higher multiplicities.

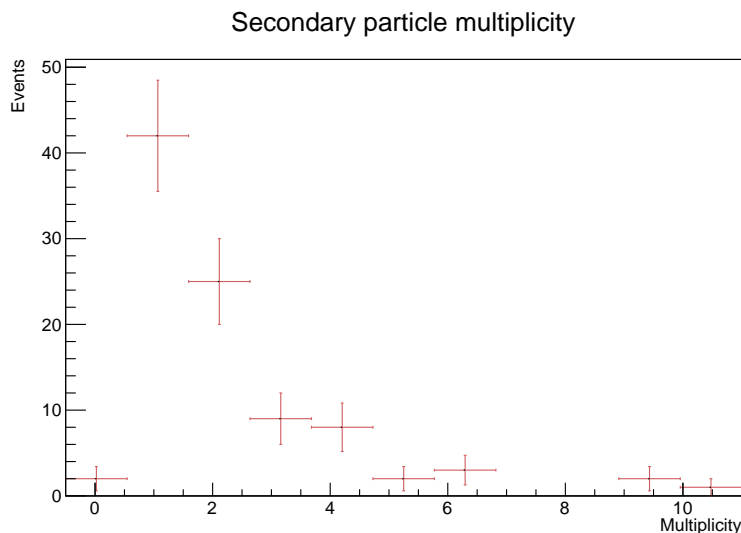


Figure 9.8: Multiplicity of charged secondary particles from our hand scan. Error bars are statistical only.

The difficulties inherent in counting particle multiplicity are illustrated by Figure 9.9. Very short tracks (some only around 30 mm) can be seen at clean vertices but are completely lost where multiplicity or noise around the vertex is high. For this hand scan we counted any vertex which seemed obvious to the scanner, an admittedly subjective criterion.

Two track events can generally be counted reliably but this becomes less true as multiplicity increases. As such the values shown in Figure 9.8 should be interpreted with care.

9.3.3 Background analysis

As part of our hand scan we aimed to categorise our background events. Our findings are presented here.

Some background still remains due to t_0 issues. In these cases a track is shifted so that it starts in the fiducial volume due to a bad x position from wrong t_0 . However these events appear to be a much smaller source of background than in the proto-analysis, where they made it necessary to disregard single track events. This is in consistent with our background predictions from Monte Carlo.

Two ‘signal’ events were seen where the lepton appeared to be travelling in a backwards direction (Figure 9.10). These were almost certainly cases of muon decay within the TPC fiducial volume. Since we make no attempt to exclude such interactions for the purposes of the hand scan they are counted towards our signal total.

The largest single source of background is tracks broken by delta-rays (see Figure 9.11). Most of these cases occur with a delta on the edges of micromegas gaps which makes matching the two track halves more difficult. Highly messy deltas can also cross and re-cross a track, leading to a difficult to reconstruct event which again complicates matching.

One interesting background source is collinear tracks (Figure 9.12). These are cases where tracks formed upstream of a TPC move apart from each other inside the fiducial volume. Since pattern recognition searches for the points at which paths diverge to form junctions this creates a fake vertex inside the TPC.

A less common but still interesting background topology is overlapping tracks (Figure 9.13). In these cases two tracks come close enough to each other to form a fake vertex. These are rare in isolation but can occur more frequently in the presence of messy out of fiducial volume events.

A big source of our misreconstruction background is high multiplicity events (probably deep inelastic scattering) occurring outside the TPCs (Figure 9.21). In

cases like this misreconstruction due to collinear or overlapping tracks becomes very likely. It is hard to determine the precise cause of the reconstruction failure in events such as this but fake vertices are most likely to emerge where large numbers of tracks are in close proximity with each other.

The last major source of misreconstruction background is missing hits near one end of a through-going track (Figure 9.15). This appears in tracks where hits are spread out, making connections between them at the pattern recognition stage difficult. It seems to be mostly associated either with tracks travelling in the x direction or with pairs of tracks sharing the same micromegas pads but at different x positions. Delta-rays in tracks with a large x component can exacerbate the problem.

There are also edge cases, two of which are shown in Figure 9.16. Interestingly background from hairy tracks as illustrated in Figure 9.16b was not prevalent in the sample, indicating that our reconstruction modifications to handle them were successful.

These hand scan results are intended to present an overview rather than a completely quantitative picture of background sources. It is often difficult to determine the precise source of a background event. A fake signal track may have bad t_0 *because* of a delta ray, or result from missing hits near the start of a through going track in a high multiplicity event.

Another take away message from these events is the wide variety of failures which can produce background events. We do not have a single dominant background topology. At our current level of purity remaining background comes from many different edge cases which are able to slip through our reconstruction and analysis. General improvements in the reconstruction can further reduce this background but it is important to take care that general reconstruction is not adversely affected while attempting to remove a very specific subset of our gas interaction background.

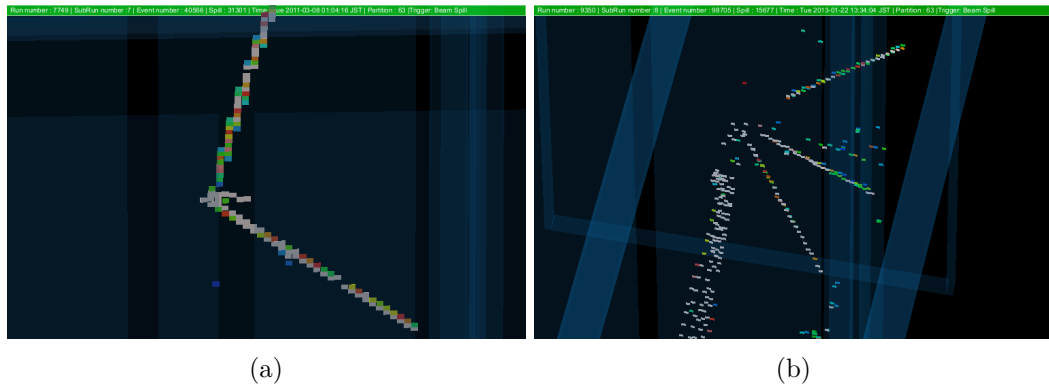


Figure 9.9: Hand scan events of difficult to determine multiplicity. Figure 9.9a was counted as having two secondary particles and Figure 9.9b was counted as four secondary particles.

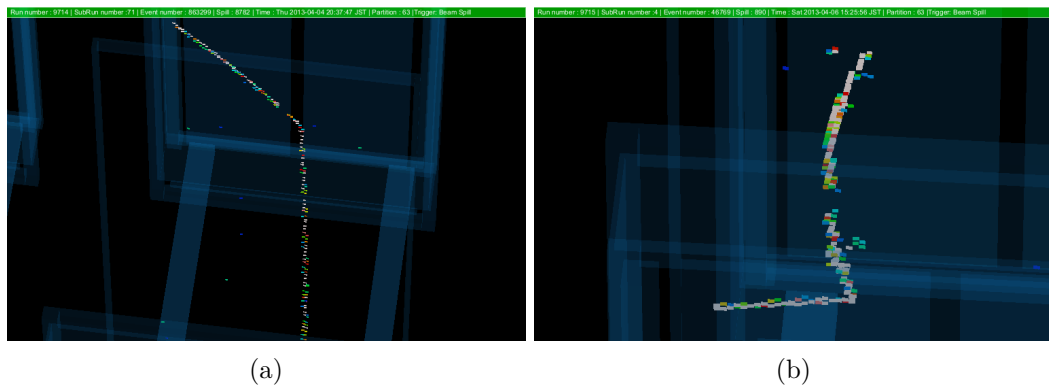


Figure 9.10: Hand scan events with backwards going products. In Figure 9.10a the secondary particle is an electron whilst in Figure 9.10b it has proton PID.

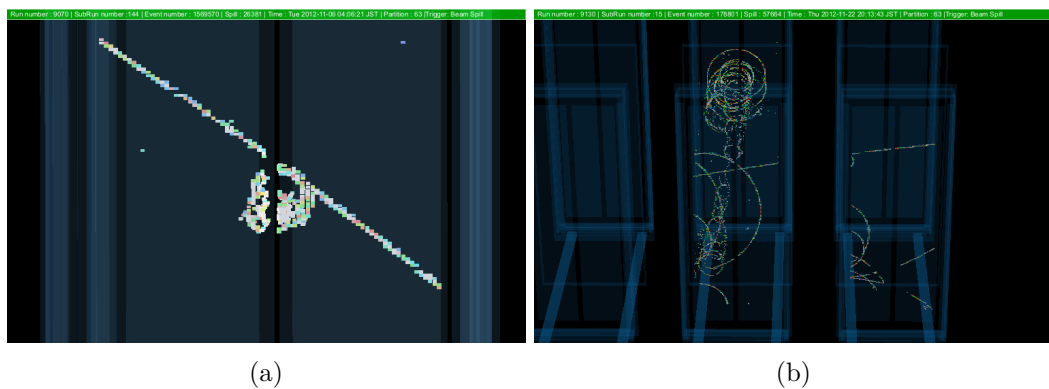


Figure 9.11: Hand scan background events with delta-rays.

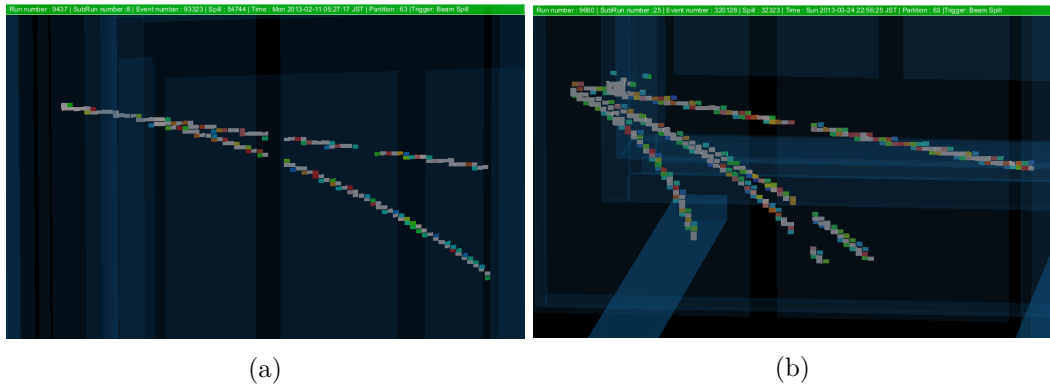


Figure 9.12: Hand scan events of collinear track background.




Figure 9.13: Hand scan background events with overlapping tracks.




Figure 9.14: Hand scan background events with messy out of fiducial volume interactions.




Figure 9.15: Hand scan background events with unassociated hits. In both Figure 9.15a and Figure 9.15b green lines represent reconstructed tracks.

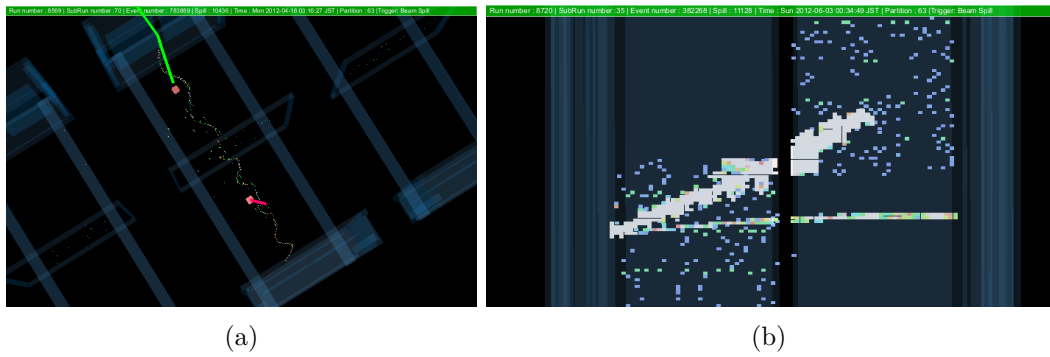


Figure 9.16: Miscellaneous hand scan background events. Figure 9.16a shows a delta ray which has somehow acquired muon PID and Figure 9.16b shows an out of fiducial volume track where start position is obscured by hairiness.

9.3.4 Overview of signal

Here we present a qualitative overview of some of our selected signal events. The real power of our gas interaction analysis is the ability to resolve features inaccessible to denser detectors so the precise details of these events are highly interesting. In the future the T2K gas interaction analysis will be able to provide a more quantitative description of these topologies.

Our largest single source of signal is two track events (Figure 9.17). In the majority of these cases tracks both the primary and secondary particle were travelling in the forwards (downstream) direction.

An interesting subset of two track events is a number of events where muon and secondary particle were travelling almost back-to-back (Figure 9.18). Cases where the particles travel at even shallower angles may have been lost at the reconstruction level due to unavoidable limitations in our kink finding algorithms.

A small number of events were observed which seemed to match up with the ‘hammer’ topology observed by ArgoNeuT[142][143] with two back to back proton tracks (Figure 9.19).

One event was seen where a secondary interaction appeared to be taking place immediately downstream of the vertex Figure 9.20. The muon was associated with the primary vertex and passed through the detector whilst none of the products of the secondary interaction made it through the ECals.

A handful of events with fairly high multiplicity were seen Figure 9.21. In these cases it is difficult to be sure of the true number of charged secondary particles as short tracks may become lost or obscured.

Many of these events featured what appeared to be hairy features Figure 9.23. These features are much more common in signal than background, supporting the hypothesis that they are a result of low momentum protons (which would generally not feature in out of fiducial volume background) depositing large amounts of charge in the TPC electronics. Finally a large number of our signal events feature what can be considered ‘high vertex activity’. Large amounts of charge was deposited at the vertex, possibly indicating the presence of additional short tracks, but these were only counted if a track was clearly resolvable.

During the course of this hand scan we identified many interesting topologies. We also encountered events where properties of the vertex were obscured by hairiness. Overcoming these difficulties to quantifying the features of our signal events constitutes an exciting challenge for future gas interaction analysers.

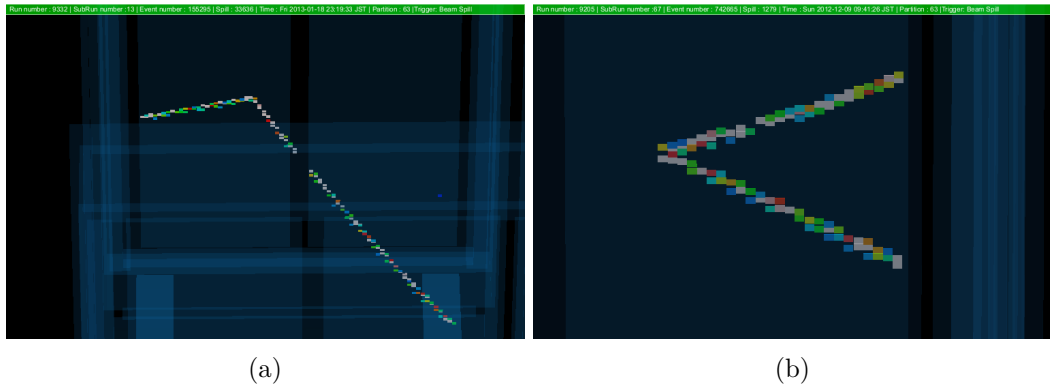


Figure 9.17: Hand scan signal events with two tracks.

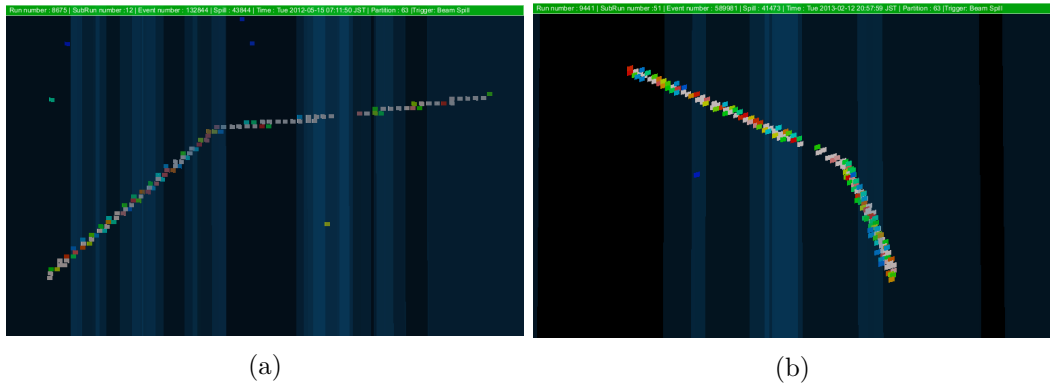


Figure 9.18: Hand scan signal events with back to back tracks.

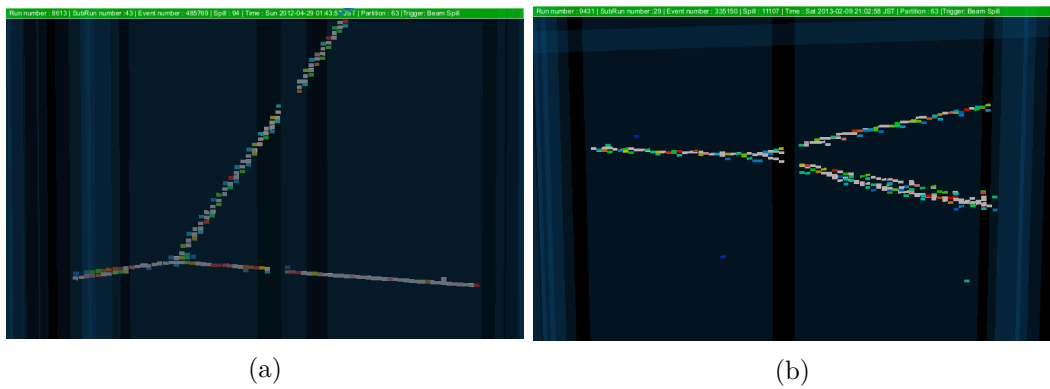


Figure 9.19: Hand scan signal events with hammer-like topology. In both cases the muon is not part of the back-to-back pair.






Figure 9.20: Hand scan signal event with a secondary interaction close to the vertex.



(a)

(b)

Figure 9.21: Hand scan signal events with high multiplicity.



(a)

(b)

Figure 9.22: Hand scan signal events with hairy appearance.




Figure 9.23: Hand scan signal events with high vertex activity. Figure 9.23a was counted as having one secondary particle and Figure 9.23b as two.

Chapter 10

Conclusions

This thesis has covered two areas relating to interactions in an argon gas TPC:

- Developing **TREx reconstruction** algorithms to identify candidate vertices and fit their paths.
- Writing a **gas interaction analysis** to extract cross section information and other useful information.

Reconstruction development has been a huge success. The TREx pattern recognition algorithms written over the course of my PhD have allowed us to identify candidate vertices and their paths in our TPC for the first time and present them in a fittable format. These algorithms are now being applied to event reconstruction in high pressure TPC designs that may form part of the near detector systems for future long baseline experiments such as DUNE and Hyper-Kamiokande.

The analysis was also a success. Our primary goal — a measurement of the cross section for ν_μ charged current interactions on argon — was attained. This represents the first time such a measurement has been reliably performed in the T2K flux. Our hand scan of signal events also revealed interesting topological features and showed us in incredible resolution activity near the vertex.

More importantly reaching this first milestone of an inclusive cross section measurements brings us closer to our final goal of precise measurements of the multiplicity and kinematics of secondary particles at neutrino interaction vertices. As discussed earlier a lack of experimental data causes uncertainty over the correct physics models for neutrino interactions, particularly MEC and FSI effects, as evidenced by the disagreements between GENIE and NEUT over such variables as proton multiplicity and momentum. A precise measurement of these variables, particularly at the low energies our TPC is capable of resolving, will provide crucial

information for such models. We briefly go over some potential improvements to the software and analysis in moving towards this goal.

Software and reconstruction

Modelling ‘hairy’ events even to a limited degree will greatly improve the predictive power of our Monte Carlo when it comes to evaluating proton multiplicity and kinematics. It will also help in the further development of reconstruction algorithms for handling them, an effort which has already seen much success.

TREx pattern recognition is highly effective at identifying tracks and vertices but potential improvements exist which could make TREx2.0 even better. Voxelisation (subsection 6.3.4) is not actually an essential part of the algorithm and could be phased out to increase our flexibility, particularly in reconstructing tracks with a large x component. Junction formation also needs some improvements to adequately handle the sort of hairy and high activity vertices seen in data; so long as we can accurately associate hits with junctions the rest of pattern recognition can proceed smoothly with minimum changes to the code. As seen in the analysis delta-rays still represent a particular reconstruction challenge so specialised procedures for ensuring their hits are not associated with genuine paths should be a priority.

More sophisticated approaches to tracking is an obvious area for improvement. A logical extension to our current method would be the implementation of a multi-pass approach. This would involve preliminary tracking results seeding a second pass of pattern recognition. This could allow extra hits to be recovered by paths, very short tracks to be detected near junctions and delta rays to be isolated so as not to interfere with second pass fitting. Another very interesting improvement would be simultaneous fitting of tracks around each vertex. This would allow both more accurate kinematics for interaction products, particularly short tracks, and a reliable vertex position established at the reconstruction level.

In our gas interaction analysis one thing we have repeatedly seen is that reconstruction *is* a fundamental part of the analysis. The next big improvement in TREx2.0 could quite conceivably take us to the point where our analysis can accurately quantify complex properties of vertices.

Analysis

Once we have a greater understanding of our background contamination a background subtraction method can be used to calculate cross sections, insulating our measurements from dependencies on our generators.

In order to measure detailed properties of vertices such as proton momentum and multiplicity some sort of unfolding may be required. This requires an accurate understanding of our selection efficiency and purity. As seen in our results there is potentially slight tension between purity in real data and simulation although this conclusion results from a somewhat unreliable hand scan. We also need a robust evaluation of systematic uncertainties on proton multiplicity and momentum which so far have not undergone a full study.

Final thoughts

We have developed a full reconstruction algorithm for a gas TPC which has been corroborated by a reliable cross section measurement. TREx now forms T2K's official TPC reconstruction and I hope for it to inspire future algorithms.

The gas interaction analysis is well positioned for further development. It is important to remember the unusual position of the T2K TPCs as relatively low density and high resolution and thus able to observe neutrino interaction vertices in uniquely fine detail. We are the only experiment able to resolve and count very low energy protons emerging from vertices and a quantitative result taking advantage of this fact will be very helpful developing the theory and simulation of neutrino interactions. Here we have taken a first big stride towards that goal.

Bibliography

- [1] Frank Close. *Neutrino*. Oxford University Press, 2010.
- [2] G. M. Lewis. *Neutrinos*. Wykeham Publications, 1970.
- [3] Wolfgang Pauli. On the earlier and more recent history of the neutrino. In Klaus Winter, editor, *Neutrino Physics*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 2nd edition, 2000.
- [4] Rudolf Peierls. The early days of neutrino physics. In Ettore Fiorini, editor, *Neutrino Physics and Astrophysics*, NATO conference series. Plenum Press, 1980.
- [5] E. Fermi. An attempt of a theory of beta radiation. 1. *Z.Phys.*, 88:161–177, 1934.
- [6] F. Wilson. Fermi’s Theory of Beta Decay. *Am.J.Phys*, 36:1150–1160, 1968.
- [7] S. Schael et al. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006.
- [8] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956.
- [9] F. Reines and C.L. Cowan. Detection of the free neutrino. *Phys.Rev.*, 92:830–831, 1953.
- [10] G. Danby, J.M. Gaillard, Konstantin A. Goulianos, L.M. Lederman, Nari B. Mistry, et al. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys.Rev.Lett.*, 9:36–44, 1962.
- [11] Martin L. Perl, G.S. Abrams, A. Boyarski, Martin Breidenbach, D. Briggs, et al. Evidence for Anomalous Lepton Production in $e^+ - e^-$ Annihilation. *Phys.Rev.Lett.*, 35:1489–1492, 1975.

- [12] K. Kodama et al. Observation of tau neutrino interactions. *Phys.Lett.*, B504:218–224, 2001.
- [13] S. M. Bilenky. Bruno pontecorvo and the neutrino. *Physics-Uspokhi*, 57(5):489, 2014.
- [14] John N. Bahcall. Electron capture and nuclear matrix elements of be^7 . *Phys. Rev.*, 128:1297–1301, Nov 1962.
- [15] Trevor Pinch. *Confronting Nature: The Sociology of Solar-Neutrino Detection*. D. Reidel Publishing Company, 1986.
- [16] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. New solar opacities, abundances, helioseismology, and neutrino fluxes. *Astrophys. J.*, 621:L85–L88, 2005.
- [17] JohnN. Bahcall. Solar neutrinos: Theory versus observation. *Space Science Reviews*, 24(2):227–251, 1979.
- [18] John N. Bahcall and Raymond Davis Jr. An account of the development of the solar neutrino problem. In Charles A. Barnes, Donald D. Clayton, and David Schramm, editors, *Essays in Nuclear Astrphysics*. Cambridge University Press, 1982.
- [19] B.T. Cleveland, Timothy Daily, Jr. Davis, Raymond, James R. Distel, Kenneth Lande, et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys.J.*, 496:505–526, 1998.
- [20] John N. Bahcall et al. Proposed Solar Neutrino Experiment Using Ga-71. *Phys. Rev. Lett.*, 40:1351–1354, 1978.
- [21] J. N. Abdurashitov et al. Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle. *J. Exp. Theor. Phys.*, 95:181–193, 2002. [Zh. Eksp. Teor. Fiz.122,211(2002)].
- [22] J. N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period. *Phys. Rev.*, C80:015807, 2009.
- [23] W. Hampel et al. GALLEX solar neutrino observations: Results for GALLEX IV. *Phys. Lett.*, B447:127–133, 1999.

- [24] M. Altmann et al. Complete results for five years of GNO solar neutrino observations. *Phys. Lett.*, B616:174–190, 2005.
- [25] K. S. Hirata et al. Real time, directional measurement of B-8 solar neutrinos in the Kamiokande-II detector. *Phys. Rev.*, D44:2241, 1991. [Erratum: *Phys. Rev.* D45,2170(1992)].
- [26] S. Fukuda et al. Solar B-8 and hep neutrino measurements from 1258 days of Super-Kamiokande data. *Phys. Rev. Lett.*, 86:5651–5655, 2001.
- [27] J. Boger et al. The Sudbury neutrino observatory. *Nucl. Instrum. Meth.*, A449:172–207, 2000.
- [28] Q. R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002.
- [29] C.V. Achar, M.G.K. Menon, V.S. Narasimham, P.V.Ramana Murthy, B.V. Sreekantan, K. Hinotani, S. Miyake, D.R. Creed, J.L. Osborne, J.B.M. Pattison, and A.W. Wolfendale. Detection of muons produced by cosmic ray neutrinos deep underground. *Physics Letters*, 18(2):196 – 199, 1965.
- [30] F. Reines, M. F. Crouch, T. L. Jenkins, W. R. Kropp, H. S. Gurr, G. R. Smith, J. P. F. Sellschop, and B. Meyer. Evidence for high-energy cosmic-ray neutrino interactions. *Phys. Rev. Lett.*, 15:429–433, Aug 1965.
- [31] M. R. Krishnaswamy, M. G. K. Menon, V. S. Narasimham, K. Hinotani, N. Ito, S. Miyake, J. L. Osborne, A. J. Parsons, and A. W. Wolfendale. The kolar gold fields neutrino experimen. i. the interactions of cosmic ray neutrinos. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 323(1555):489–509, 1971.
- [32] M. F. Crouch, P. B. Landecker, J. F. Lathrop, F. Reines, W. G. Sandie, H. W. Sobel, H. Coxell, and J. P. F. Sellschop. Cosmic-ray muon fluxes deep underground: Intensity vs depth, and the neutrino-induced component. *Phys. Rev. D*, 18:2239–2252, Oct 1978.
- [33] H. E. Bergeson, G. L. Cassiday, and M. B. Hendricks. Neutrino-induced muons deep underground. *Phys. Rev. Lett.*, 31:66–70, Jul 1973.
- [34] W. Frati, T. K. Gaisser, A. K. Mann, and Todor Stanev. Atmospheric neutrino data and neutrino oscillations. *Phys. Rev. D*, 48:1140–1149, Aug 1993.

- [35] Y. Fukuda et al. Atmospheric muon-neutrino / electron-neutrino ratio in the multiGeV energy range. *Phys. Lett.*, B335:237–245, 1994.
- [36] D. Casper, R. Becker-Szendy, C. B. Bratton, D. R. Cady, R. Claus, S. T. Dye, W. Gajewski, M. Goldhaber, T. J. Haines, P. G. Halverson, T. W. Jones, D. Kielczewska, W. R. Kropp, J. G. Learned, J. M. LoSecco, C. McGrew, S. Matsuno, J. Matthews, M. S. Mudan, L. Price, F. Reines, J. Schultz, D. Sinclair, H. W. Sobel, J. L. Stone, L. R. Sulak, R. Svoboda, G. Thornton, and J. C. van der Velde. Measurement of atmospheric neutrino composition with the imb-3 detector. *Phys. Rev. Lett.*, 66:2561–2564, May 1991.
- [37] W. W. M. Allison et al. The Atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton year exposure of Soudan-2. *Phys. Lett.*, B449:137–144, 1999.
- [38] M. Ambrosio et al. Measurement of the atmospheric neutrino induced upgoing muon flux using MACRO. *Phys. Lett.*, B434:451–457, 1998.
- [39] K. Daum et al. Determination of the atmospheric neutrino spectra with the Frejus detector. *Z. Phys.*, C66:417–428, 1995.
- [40] M. Aglietta et al. Experimental study of atmospheric neutrino flux in the NUSEX experiment. *Europhys. Lett.*, 8:611–614, 1989.
- [41] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.
- [42] M. Nakahata. Super-Kamiokande. *Nucl. Phys. Proc. Suppl.*, 87:125–134, 2000.
- [43] M. H. Ahn et al. Indications of neutrino oscillation in a 250 km long baseline experiment. *Phys. Rev. Lett.*, 90:041801, 2003.
- [44] M. H. Ahn et al. Measurement of Neutrino Oscillation by the K2K Experiment. *Phys. Rev.*, D74:072003, 2006.
- [45] Bruno Pontecorvo. Fifty years of neutrino physics: A few examples. In Ettore Fiorini, editor, *Neutrino Physics and Astrophysics*, NATO conference series. Plenum Press, 1980.
- [46] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *J. Exptl. Theoret. Phys. (U.S.S.R.)*, 34:247–249, Jan 1958.
- [47] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962.

- [48] Kai Zuber. *Neutrino Physics*. CRC Press, 2nd edition, 2012.
- [49] Julien Lesgourgues, Mangano Gianpiero, Gennaro Miele, and Sergio Pastor. *Neutrino Cosmology*. Cambridge University Press, 2013.
- [50] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.
- [51] D.V. Forero, M. Tortola, and J.W.F. Valle. Neutrino oscillations refitted. *Phys.Rev.*, D90(9):093006, 2014.
- [52] Evgeny K. Akhmedov. Neutrino physics. In *Particle physics. Proceedings, Summer School, Trieste, Italy, June 21-July 9, 1999*, pages 103–164, 1999.
- [53] Jouni I Peltoniemi and Juha T Peltoniemi. Coherence conditions for the forward scattering of neutrinos. *JHEP*, 08:008, 1999.
- [54] Tsuyoshi Nakaya and Robert K. Plunkett. Neutrino Oscillations with the MINOS, MINOS+, T2K, and NOvA Experiments. 2015.
- [55] Hiroshi Nunokawa, Stephen J. Parke, and Renata Zukanovich Funchal. Another possible way to determine the neutrino mass hierarchy. *Phys. Rev.*, D72:013009, 2005.
- [56] D. G. Michael et al. Observation of muon neutrino disappearance with the MINOS detectors and the NuMI neutrino beam. *Phys. Rev. Lett.*, 97:191801, 2006.
- [57] A. Cervera, A. Donini, M. B. Gavela, J. J. Gomez Cadenas, P. Hernandez, Olga Mena, and S. Rigolin. Golden measurements at a neutrino factory. *Nucl. Phys.*, B579:17–55, 2000. [Erratum: *Nucl. Phys.*B593,731(2001)].
- [58] Patrick Huber, Michele Maltoni, and Thomas Schwetz. Resolving parameter degeneracies in long-baseline experiments by atmospheric neutrino data. *Phys. Rev.*, D71:053006, 2005.
- [59] J. A. Formaggio and G. P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev. Mod. Phys.*, 84:1307, 2012.
- [60] C.H. Llewellyn Smith. Neutrino reactions at accelerator energies. *Physics Reports*, 3(5):261 – 379, 1972.

- [61] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. The nucleon axial mass and the MiniBooNE Quasielastic Neutrino-Nucleus Scattering problem. *Phys. Lett.*, B707:72–75, 2012.
- [62] R. A. Smith and E. J. Moniz. NEUTRINO REACTIONS ON NUCLEAR TARGETS. *Nucl. Phys.*, B43:605, 1972. [Erratum: *Nucl. Phys.*B101,547(1975)].
- [63] L. A. Ahrens et al. Measurement of Neutrino - Proton and anti-neutrino - Proton Elastic Scattering. *Phys. Rev.*, D35:785, 1987.
- [64] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. Two Particle-Hole Excitations in Charged Current Quasielastic Antineutrino–Nucleus Scattering. *Phys. Lett.*, B721:90–93, 2013.
- [65] Dieter Rein and Lalit M. Sehgal. Neutrino Excitation of Baryon Resonances and Single Pion Production. *Annals Phys.*, 133:79–153, 1981.
- [66] Dieter Rein and Lalit M. Sehgal. Coherent π^0 Production in Neutrino Reactions. *Nucl. Phys.*, B223:29, 1983.
- [67] D. Rein and L. M. Sehgal. PCAC and the Deficit of Forward Muons in π^+ Production by Neutrinos. *Phys. Lett.*, B657:207–209, 2007.
- [68] M. Glück, E. Reya, and A. Vogt. Dynamical parton distributions revisited. *Eur. Phys. J.*, C5:461–470, 1998.
- [69] A. Bodek and U. K. Yang. Modeling neutrino and electron scattering inelastic cross- sections in the few GeV region with effective LO PDFs TV Leading Order. In *2nd International Workshop on Neutrino-Nucleus Interactions in the Few GeV Region (NuInt 02) Irvine, California, December 12-15, 2002*, 2003.
- [70] A Bodek and U. K. Yang. Higher twist, $\xi(\omega)$ scaling, and effective LO PDFs for lepton scattering in the few GeV region. *J. Phys.*, G29:1899–1906, 2003.
- [71] T. Leitner and U. Mosel. Neutrino-nucleus scattering reexamined: Quasielastic scattering and pion production entanglement and implications for neutrino energy reconstruction. *Phys. Rev.*, C81:064614, 2010.
- [72] U. Mosel and O. Lalakulich. Neutrino-nucleus interactions. In *Proceedings, 14th International Workshop on Neutrino Factories, Super Beams and Beta*

Beams (NuFact 2012): Williamsburg, USA, July 23-28, 2012, pages 17–24, 2012.

- [73] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti-neutrino disappearance. *Phys. Rev. Lett.*, 90:021802, 2003.
- [74] Carlo Bemporad, Giorgio Gratta, and Petr Vogel. Reactor based neutrino oscillation experiments. *Rev. Mod. Phys.*, 74:297, 2002.
- [75] F. Boehm, J.F. Cavaignac, F.v. Feilitzsch, A.A. Hahn, H.E. Henrikson, D.H. Koang, H. Kwon, R.L. Müssbauer, B. Vignon, and J.L. Vuilleumier. Experimental study of neutrino oscillations at a fission reactor. *Physics Letters B*, 97(2):310 – 314, 1980.
- [76] G. Zacek, F. v. Feilitzsch, R. L. Mössbauer, L. Oberauer, V Zacek, V., F. Boehm, P. H. Fisher, J. L. Gimlett, A. A. Hahn, H. E. Henrikson, H. Kwon, J. L. Vuilleumier, and K. Gabathuler. Neutrino-oscillation experiments at the gösgen nuclear power reactor. *Phys. Rev. D*, 34:2621–2636, Nov 1986.
- [77] J.F. Cavaignac, A. Hoummada, D.H. Koang, B. Vignon, Y. Declais, H. de Kerret, H. Pessard, and J.M. Thenard. Indication for neutrino oscillation from a high statistics experiment at the bugey reactor. *Physics Letters B*, 148(4–5):387 – 394, 1984.
- [78] Y. Declais et al. Search for neutrino oscillations at 15-meters, 40-meters, and 95-meters from a nuclear power reactor at Bugey. *Nucl. Phys.*, B434:503–534, 1995.
- [79] Z. D. Greenwood, W. R. Kropp, M. A. Mandelkern, S. Nakamura, E. L. Pasierb-Love, L. R. Price, F. Reines, S. P. Riley, H. W. Sobel, N. Baumann, and H. S. Gurr. Results of a two-position reactor neutrino-oscillation experiment. *Phys. Rev. D*, 53:6054–6064, Jun 1996.
- [80] AI Afonin, SA Bogatov, AA Borovoi, AG Vershinskii, SL Gavrilov, Yu L Dobrynin, SN Ketov, Yu V Klimov, VI Kopeikin, LA Levina, et al. Search for neutrino oscillations in an experiment in the reactor of the rovno nuclear power plant. *Jetp Lett*, 42(5):285, 1985.
- [81] M. Apollonio et al. Search for neutrino oscillations on a long baseline at the CHOOZ nuclear power station. *Eur. Phys. J.*, C27:331–374, 2003.
- [82] F. Boehm et al. Final results from the Palo Verde neutrino oscillation experiment. *Phys. Rev.*, D64:112001, 2001.

- [83] G. Mention, M. Fechner, Th. Lasserre, Th. A. Mueller, D. Lhuillier, M. Cri-bier, and A. Letourneau. The Reactor Antineutrino Anomaly. *Phys. Rev.*, D83:073006, 2011.
- [84] J. Ashenfelter et al. The PROSPECT Physics Program. 2015.
- [85] Celine Moortgat. Technology of the SoLid detector and construction of the first submodule. In *Proceedings, 2015 European Physical Society Conference on High Energy Physics (EPS-HEP 2015)*, 2015.
- [86] Ba Ro Kim et al. Development and Mass Production of a Mixture of LAB- and DIN-based Gadolinium-loaded Liquid Scintillator for the NEOS Short-baseline Neutrino Experiment. 2015.
- [87] J. Gaffiot. The SOX experiment. *Nucl. Part. Phys. Proc.*, 265-266:129–131, 2015.
- [88] M. Pequignot. The Nucifer and Stereo reactor antineutrino experiments. *Nucl. Part. Phys. Proc.*, 265-266:126–128, 2015.
- [89] F. P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.
- [90] J. K. Ahn et al. Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 108:191802, 2012.
- [91] F. Ardellier et al. Double Chooz: A Search for the neutrino mixing angle θ_{13} . arXiv preprint, 2006.
- [92] Fengpeng An et al. Neutrino Physics with JUNO. 2015.
- [93] Zelimir Djurcic et al. JUNO Conceptual Design Report. 2015.
- [94] Soo-Bong Kim. New results from RENO and prospects with RENO-50. *Nucl. Part. Phys. Proc.*, 265-266:93–98, 2015.
- [95] P. Berge et al. Prompt neutrino results from a proton beam dump experiment. *Z. Phys.*, C56:175–180, 1992.
- [96] J. Dorenbosch et al. Prompt neutrino production in 400 gev proton-copper interactions. *Z. Phys.*, C40(4):497–522, 1988.
- [97] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the obser-vation of anti-neutrino(electron) appearance in a anti-neutrino(muon) beam. *Phys. Rev.*, D64:112007, 2001.

- [98] S. J. Freedman et al. Limits on neutrino oscillations from anti-electron-neutrino appearance. *Phys. Rev.*, D47:811–829, 1993.
- [99] B. Armbruster et al. Upper limits for neutrino oscillations muon-anti-neutrino \rightarrow electron-anti-neutrino from muon decay at rest. *Phys. Rev.*, D65:112001, 2002.
- [100] L. Borodovsky et al. Search for muon-neutrino oscillations muon-neutrino \rightarrow electron-neutrino (anti-muon-neutrino \rightarrow anti-electron-neutrino in a wide band neutrino beam. *Phys. Rev. Lett.*, 68:274–277, 1992.
- [101] L. A. Ahrens et al. Comparison of narrow-band and wide-band neutrino beams in the search for $\nu_\mu \rightarrow \nu_e$ oscillations. *Phys. Rev. D*, 36:702–706, Aug 1987.
- [102] P. Astier others. Search for $\nu_\mu \rightarrow \nu_e$ oscillations in the {NOMAD} experiment. *Physics Letters B*, 570(1–2):19 – 31, 2003.
- [103] A. Romosan et al. High statistics search for $\nu_\mu(\bar{\nu}_\mu) \rightarrow \nu_e(\bar{\nu}_e)$ oscillations in the small mixing angle regime. *Phys. Rev. Lett.*, 78:2912–2915, Apr 1997.
- [104] A. A. Aguilar-Arevalo et al. Improved search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations in the miniboone experiment. *Phys. Rev. Lett.*, 110:161801, Apr 2013.
- [105] A. A. Aguilar-Arevalo et al. A Search for electron neutrino appearance at the $\Delta m^2 \sim 1\text{eV}^2$ scale. *Phys. Rev. Lett.*, 98:231801, 2007.
- [106] A. A. Aguilar-Arevalo et al. Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations. *Phys. Rev. Lett.*, 105:181801, 2010.
- [107] Sowjanya Gollapinni. Accelerator-based Short-baseline Neutrino Oscillation Experiments. In *12th Conference on the Intersections of Particle and Nuclear Physics (CIPANP 2015) Vail, Colorado, USA, May 19-24, 2015*, 2015.
- [108] Jan Kisiel. ICARUS T600 - a large Liquid Argon Time Projection Chamber. *J. Phys. Conf. Ser.*, 650(1):012004, 2015.
- [109] Nicola McConkey. The LAr1-ND Experiment. *J. Phys. Conf. Ser.*, 650(1):012007, 2015.
- [110] M. Antonello et al. A Proposal for a Three Detector Short-Baseline Neutrino Oscillation Program in the Fermilab Booster Neutrino Beam. 2015.
- [111] L. A. Ahrens et al. New limit on the strength of mixing between ν_μ and ν_e . *Phys. Rev. D*, 31:2732–2736, Jun 1985.

- [112] F. Bergsma et al. A Search for Neutrino Oscillations. *Z. Phys.*, C40:171, 1988.
- [113] C. Angelina, A. Apostolakis, A. Baldini, M. Baldo-Ceolin, L. Bertanza, F. Bobisut, E. Calimani, U. Camerini, S. Ciampolillo, R. Fantechi, V. Flaminio, W.J. Fry, H. Huzita, P. Ioannou, S. Katsanevas, C. Kourkoumelis, J. Koutentakis, M. Loreti, R. Loveless, G. Miari, R. Pazzi, P. Pramantiotis, M. Procario, G. Puglierin, D.D. Reeder, L.K. Resvanis, B. Saitta, and M. Vassiliou. New experimental limits on $\bar{\nu}_e$ oscillations. *Physics Letters B*, 179(3):307 – 312, 1986.
- [114] F. Dydak, G.J. Feldman, C. Guyot, J.P. Merlo, H.-J. Meyer, J. Rothberg, J. Steinberger, H. Taureg, W. von Rüden, H. Wachsmuth, H. Wahl, J. Wotschack, H. Bläijmer, P. Buchholz, J. Duda, F. Eisele, K. Kleinknecht, J. Knobloch, B. Pszola, B. Renk, R. Belusevic, B. Falkenburg, T. Flottmann, J.G.H. de Groot, C. Geweniger, H. Keilwerth, K. Tittel, P. Debu, A. Para, P. Perez, B. Peyaud, J. Rander, J.P. Schuller, R. Turlay, H. Abramowicz, and J. KrÅšlikowski. A search for $\bar{\nu}_e$ oscillations in the $\bar{\nu}_e$ range 0.3–90 eV. *Physics Letters B*, 134(3–4):281 – 286, 1984.
- [115] I.E. Stockdale, A. Bodek, F. Borcherding, N. Giokaris, K. Lang, D. Garfinkle, F.S. Merritt, M. Oreglia, P. Reutens, P.S. Auchincloss, R. Blair, C. Haber, M. Ruiz, F. Sciulli, M.H. Shaevitz, W.H. Smith, R. Zhu, R. Coleman, H.E. Fisk, B. Jin, D. Levinthal, W. Marsh, P.A. Rapidis, H.B. White, and D. Yovanovitch. Search for muon neutrino and antineutrino oscillations in the mass range 15–1,000 eV/c². *Zeitschrift für Physik C Particles and Fields*, 27(1):53–56, 1985.
- [116] S. Avvakumov et al. A Search for muon-neutrino \rightarrow electron-neutrino and muon-anti-neutrino \rightarrow electron-anti-neutrino oscillations at NuTeV. *Phys. Rev. Lett.*, 89:011804, 2002.
- [117] E. Eskut et al. The CHORUS experiment to search for muon-neutrino \rightarrow tau-neutrino oscillation. *Nucl. Instrum. Meth.*, A401:7–44, 1997.
- [118] N. Ushida et al. Limits to Muon-neutrino, electron-neutrino \rightarrow tau-neutrino Oscillations and Muon-neutrino, electron-neutrino \rightarrow tau- Direct Coupling. *Phys. Rev. Lett.*, 57:2897–2900, 1986.
- [119] Sacha E. Kopp. The NuMI neutrino beam at Fermilab. In *Particle accelerator. Proceedings, Conference, PAC'05, Knoxville, USA, May 16-20, 2005*, 2005.

- [120] Alexandre B. Sousa. First MINOS+ Data and New Results from MINOS. *AIP Conf. Proc.*, 1666:110004, 2015.
- [121] R. B. Patterson. The NOvA Experiment: Status and Outlook. In *Proceedings, 25th International Conference on Neutrino Physics and Astrophysics (Neutrino 2012)*, 2012. [Nucl. Phys. Proc. Suppl.235-236,151(2013)].
- [122] N. Agafonova et al. Search for $\nu_\mu \rightarrow \nu_e$ oscillations with the OPERA experiment in the CNGS beam. *JHEP*, 07:004, 2013. [Addendum: JHEP07,085(2013)].
- [123] N. Agafonova et al. Discovery of τ Neutrino Appearance in the CNGS Neutrino Beam with the OPERA Experiment. *Phys. Rev. Lett.*, 115(12):121802, 2015.
- [124] K. Abe et al. A Long Baseline Neutrino Oscillation Experiment Using J-PARC Neutrino Beam and Hyper-Kamiokande. 2014.
- [125] C. Adams et al. The Long-Baseline Neutrino Experiment: Exploring Fundamental Symmetries of the Universe. 2013.
- [126] Andrew Renshaw. Solar Neutrino Results from Super-Kamiokande. *Phys. Procedia*, 61:345–354, 2015.
- [127] G. Bellini et al. Precision measurement of the ^7Be solar neutrino interaction rate in Borexino. *Phys. Rev. Lett.*, 107:141302, 2011.
- [128] M. Agostini et al. Spectroscopy of geoneutrinos from 2056 days of Borexino data. *Phys. Rev.*, D92(3):031101, 2015.
- [129] Ondrej Sramek, William F. McDonough, Edwin S. Kite, Vedran Lekic, Steve Dye, and Shijie Zhong. Geophysical and geochemical constraints on geoneutrino fluxes from Earth’s mantle. *Earth Planet. Sci. Lett.*, 361:356–366, 2013.
- [130] G. Fiorentini, G. L. Fogli, E. Lisi, F. Mantovani, and A. M. Rotunno. Mantle geoneutrinos in KamLAND and Borexino. *Phys. Rev.*, D86:033004, 2012.
- [131] Mark C. Chen. The SNO+ Experiment. In *Proceedings, 34th International Conference on High Energy Physics (ICHEP 2008)*, 2008.
- [132] Francesca Di Lodovico. Neutrino physics perspectives with Hyper-Kamiokande. *Nucl. Part. Phys. Proc.*, 265-266:275–280, 2015.
- [133] D. Indumathi. India-based neutrino observatory (INO): Physics reach and status report. *AIP Conf. Proc.*, 1666:100003, 2015.

- [134] Ulrich F. Katz. The ORCA Option for KM3NeT. In *Proceedings of the 15th International Workshop on Neutrino Telescopes (Neutel 2013)*, 2014.
- [135] M. G. Aartsen et al. PINGU Sensitivity to the Neutrino Mass Hierarchy. In *Cosmic Frontier Workshop: Snowmass 2013 Menlo Park, USA, March 6-8, 2013*, 2013.
- [136] L. Aliaga et al. Design, Calibration, and Performance of the MINERvA Detector. *Nucl. Instrum. Meth.*, A743:130–159, 2014.
- [137] Yasuhiro Nakajima. Status of FNAL SciBooNE experiment. *J. Phys. Conf. Ser.*, 120:052043, 2008.
- [138] A. A. Aguilar-Arevalo et al. First measurement of the muon antineutrino double-differential charged-current quasielastic cross section. *Phys. Rev.*, D88(3):032001, 2013.
- [139] Teppei Katori. Cross section analyses in MiniBooNE and SciBooNE experiments. *AIP Conf. Proc.*, 1663:020001, 2015.
- [140] C. Anderson et al. First Measurements of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon. *Phys. Rev. Lett.*, 108:161802, 2012.
- [141] R. Acciarri et al. Measurements of Inclusive Muon Neutrino and Antineutrino Charged Current Differential Cross Sections on Argon in the NuMI Antineutrino Beam. *Phys. Rev.*, D89(11):112003, 2014.
- [142] R. Acciarri et al. Detection of back-to-back proton pairs in charged-current neutrino interactions with the ArgoNeuT detector in the NuMI low energy beam line. *Phys. Rev.*, D90(1):012008, 2014.
- [143] A. M. Szec. Recent Results from ArgoNeuT and Status of MicroBooNE. *Nucl. Part. Phys. Proc.*, 265-266:208–211, 2015.
- [144] R. Acciarri et al. First Measurement of Neutrino and Antineutrino Coherent Charged Pion Production on Argon. *Phys. Rev. Lett.*, 113(26):261801, 2014. [Erratum: *Phys. Rev. Lett.* 114, no. 3, 039901 (2015)].
- [145] K. Abe et al. The T2K Experiment. *Nucl. Instrum. Meth.*, A659:106–135, 2011.
- [146] K. Abe et al. Indication of Electron Neutrino Appearance from an Accelerator-produced Off-axis Muon Neutrino Beam. *Phys. Rev. Lett.*, 107:041801, 2011.

- [147] K. Abe et al. Observation of Electron Neutrino Appearance in a Muon Neutrino Beam. *Phys. Rev. Lett.*, 112:061802, 2014.
- [148] Taku Ishida. Operational Status and Power Upgrade Prospects of the Neutrino Experimental Facility at J-PARC. *JPS Conf. Proc.*, 8:023012, 2015.
- [149] K. A. Olive et al. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014.
- [150] T. Sekiguchi et al. Development and operational experience of magnetic horn system for T2K experiment. *Nucl. Instrum. Meth.*, A789:57–80, 2015.
- [151] Y. Fukuda et al. The Super-Kamiokande detector. *Nucl. Instrum. Meth.*, A501:418–462, 2003.
- [152] K. Abe et al. Calibration of the Super-Kamiokande Detector. *Nucl. Instrum. Meth.*, A737:253–272, 2014.
- [153] M. Otani et al. Design and construction of INGRID neutrino beam monitor for T2K neutrino experiment. *Nucl. Instrum. Meth.*, A623:368–370, 2010.
- [154] Takatomi Yano. The Side Muon Range Detector for the T2K experiment. *Nucl. Phys. Proc. Suppl.*, 229-232:454, 2012.
- [155] D. Allan et al. The Electromagnetic Calorimeter for the T2K Near Detector ND280. *JINST*, 8:P10019, 2013.
- [156] S. Assylbekov et al. The T2K ND280 Off-Axis Pi-Zero Detector. *Nucl. Instrum. Meth.*, A686:48–63, 2012.
- [157] P. A. Amaudruz et al. The T2K Fine-Grained Detectors. *Nucl. Instrum. Meth.*, A696:1–31, 2012.
- [158] N. Abgrall et al. Time Projection Chambers for the T2K Near Detectors. *Nucl. Instrum. Meth.*, A637:25–46, 2011.
- [159] B. Jamieson. Experience with the Time Projection Chambers for the T2K Near Detectors. *Phys. Procedia*, 37:535–542, 2012.
- [160] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth.*, A389:81–86, 1997.
- [161] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.

- [162] Rene Brun, Federico Carminati, and Simone Giani. GEANT Detector Description and Simulation Tool. 1994.
- [163] K. Abe et al. T2K neutrino flux prediction. *Phys. Rev.*, D87(1):012001, 2013. [Addendum: *Phys. Rev.*D87,no.1,019902(2013)].
- [164] Alfredo Ferrari, Paola R. Sala, Alberto Fasso, and Johannes Ranft. FLUKA: A multi-particle transport code (Program version 2005). 2005.
- [165] N Abgrall et al. Measurements of Cross Sections and Charged Pion Spectra in Proton-Carbon Interactions at 31 GeV/c. *Phys. Rev.*, C84:034604, 2011.
- [166] Y. Hayato. NEUT. *Nucl. Phys. Proc. Suppl.*, 112:171–176, 2002. [,171(2002)].
- [167] Omar Benhar and Adelchi Fabrocini. Two nucleon spectral function in infinite nuclear matter. *Phys. Rev.*, C62:034304, 2000.
- [168] R. Gran, J. Nieves, F. Sanchez, and M. J. Vicente Vacas. Neutrino-nucleus quasi-elastic and 2p2h interactions up to 10 GeV. *Phys. Rev.*, D88(11):113007, 2013.
- [169] C. Andreopoulos et al. The GENIE Neutrino Monte Carlo Generator. *Nucl. Instrum. Meth.*, A614:87–104, 2010.
- [170] Costas Andreopoulos, Christopher Barry, Steve Dytman, Hugh Gallagher, Tomasz Golan, Robert Hatcher, Gabriel Perdue, and Julia Yarba. The GENIE Neutrino Monte Carlo Generator: Physics and User Manual. 2015.
- [171] A. Bodek and J. L. Ritchie. Further Studies of Fermi Motion Effects in Lepton Scattering from Nuclear Targets. *Phys. Rev.*, D24:1400, 1981.
- [172] Phillip Hamilton. *A study of neutrino interactions in argon gas*. PhD thesis, Imperial College London, 2015.
- [173] Thomas Lindner, Daniel Brook-Roberge, Anthony Hillairet, Blair Jamieson, and Federico Sahchez. ND280 Reconstruction. T2K Technical Note, 2011.
- [174] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- [175] Gary Barker, Steven Boyd, Stephen Dennis, Andrew Furmanski, David Hadley, Callum Lister, Phillip Lichfield, Daniel Scully, and Leigh Whitehead.

- Second Generation Particle ID with the ND280 Tracker ECALs. T2K Technical Note, 2012.
- [176] P.E. Hart, N.J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *Systems Science and Cybernetics, IEEE Transactions on*, 4(2):100–107, July 1968.
- [177] E. W. Dijkstra. A note on two problems in connexion with graphs. *NUMERISCHE MATHEMATIK*, 1(1):269–271, 1959.
- [178] F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975.
- [179] Hillairet, A and others. ν_μ CC event selections in the ND280 tracker using Run 2+3+4 data. T2K Technical Note, 2015.
- [180] J.F. Ziegler. Srim/trim version 2003.00. Computer software.
- [181] Cervera, Anselmo and Escudero, L. Study of momentum resolution and scale using tracks that cross multiple TPCs. T2K Technical Note, 2014.
- [182] Di Lodovico, F and others. CCQE-like and CC-non-QE-like ν_μ event selections in the ND280 tracker using Run 1+2+3 data. T2K Technical Note, 2013.
- [183] K. Abe et al. Measurement of the ν_μ charged current quasielastic cross section on carbon with the T2K on-axis neutrino beam. *Phys. Rev.*, D91(11):112002, 2015.
- [184] From T2K beam group.
- [185] Eike, Frank and others. B-field calibration and systematic errors. T2K Technical Note, 2011.
- [186] Bojecho, C and others. Measurement and Correction of Magnetic Field Distortions in the Time Projection Chambers. T2K Technical Note, 2013.
- [187] Bordoni, S and others. TPC-PID status and performances with Production 6. T2K Technical Note, 2014.
- [188] Petrov, Yevgeniy and Hillairet, Anthony. ND280 TPC Track Reconstruction Efficiency. T2K Technical Note, 2016.
- [189] Fiorentini, Arturo. ND280 TPC cluster Efficiency. T2K Technical Note, 2015.

- [190] Sanchez, Federico and others. ND280 global charge identification systematic error. T2K Technical Note, 2016.