Measuring Neutrino Oscillations with NOvA and T2K

Thesis by
Varun Ranga Raj

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended May 23, 2025

Varun Ranga Raj
ORCID: 0000-0003-0790-4814

ACKNOWLEDGEMENTS

While a thesis only bears one author, it is impossible to imagine I could have written this document without the help of so many others.

I first got into neutrino physics as an undergraduate at UC Berkeley. Special thanks go to my advisor Kam-Biu Luk, who gave us undergrads an up-close look at life as a working physicist. Additional thanks to the other mentors I had there: Herb Steiner, Dan Dwyer, Yasuhiro Nakajima, and Sam Kohn.

As part of my time on NOvA, I've had the good fortune to work with many of our lovely collaborators. Between the NOvA Test Beam and time spent as a NOvA DAQ expert, I learned from people like Alex Sousa, Mike Wallbank, Will Flanagan, Anne Norrick, Maria Martinez-Casales, Teresa Lackey, Andrew Sutton, and Tom Wieber.

The NOvA-T2K joint working group is small but mighty, and I'm appreciative of the opportunity to work with these fine folks. On the NOvA side, thank you to Erika Cataño-Mur, Liudmila Kolupaeva, and Artur Sztuc. On the T2K side, thank you to Kendall Mahn, Patrick Dunne, Ed Atkin, Evan Goodman, and Marvin Pfaff.

The people I've worked with in the Caltech neutrino group have been altogether excellent. Thanks to Dan Pershey, Joe Lozier, Kirk Bays, Scott Schwartz, and Yuxuan (Billy) Wang. Special thanks to Leon Mualem for sharing a glimpse of his NOvA DAQ wizardry and for being an amazing role model.

Also at Caltech, Zoya Vallari and Kathryn Sutton were amazing mentors and were instrumental to the NOvA-T2K joint fit. In particular, I worked closely with Zoya throughout my time in grad school, and she was always able to lend an ear or get me excited about our work again. Finally, I've enjoyed the camaraderie with Alex Diaz. I look forward to continuing my friendship with these three.

I am immensely grateful to my advisor, Ryan Patterson. His guidance, interest, patience, confidence, and support through this process have made this work possible. I am sincerely thankful for the opportunity to learn from him.

Thank you to my parents, Ranga and Hema, as well as my grandparents and extended family. Finally, thank you to all my friends from childhood all the way to grad school for keeping me sane, especially those going through academic journeys of their own: Vikas Aragam, Joey Barreto, Tom Naragon, and Albert Wandui.

ABSTRACT

The discoveries of the twentieth century proved that neutrinos have mass and can change flavor. For the past few decades, a major focus of research has been the measurement of the physical parameters which govern this flavor oscillation. These measurements remain inconclusive on a few key questions, including the ordering of the neutrino masses and whether neutrinos violate CP symmetry. NOvA and T2K are two long-baseline accelerator neutrino experiments working in this space. By placing detectors in a beam of muon (anti-)neutrinos, these experiments interrogate neutrino oscillations by measuring $\overset{(-)}{\nu}_\mu$ disappearance and $\overset{(-)}{\nu}_e$ appearance. The complementarity of NOvA's and T2K's oscillation measurements motivated the experiments to pursue a joint oscillation analysis. After bracketing the potential impacts of correlations between the two experiments' systematic uncertainties and constructing a joint likelihood function, we share in this thesis the first results from the NOvA-T2K joint oscillation analysis. We report the world's most precise measurement of Δm^2_{32} to date: $+2.429^{+0.039}_{-0.035}(-2.477 \pm 0.035) \times 10^{-3}$ eV^2 assuming the normal (inverted) mass ordering, showing a slight preference for the inverted mass ordering. The maximally CP-violating value of $\delta_{CP} = +\frac{\pi}{2}$ is excluded by 3σ credible intervals, and if we assume neutrinos are in the inverted mass ordering, we see evidence of CP violation at 3σ.

Additionally, we present an effort to encapsulate neutrino cross-section models in a parametrization-agnostic way. We have created a suite of systematic parameters that are capable of mimicking the action of NOvA's cross-section model. This method could be used in a future joint data analysis between long-baseline neutrino experiments. We also introduce Voronoi histograms, an ancillary technique developed as part of this program. Voronoi histograms are a new way to efficiently bin high-dimensional data. This method preserves bin density in regions of interest, while tightly controlling the total number of bins used. The performance gains from using Voronoi binnings over standard rectangular binnings scale dramatically with the dimensionality of the data.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Chapter 1

THE LANDSCAPE OF NEUTRINO OSCILLATIONS

1.1 History

The story of neutrinos begins in the early twentieth century with the measurement of missing energy in beta decay experiments. In 1930, Wolfgang Pauli proposed an invisible particle which carried off this energy. This particle was dubbed the "neutrino" by Enrico Fermi in his 1934 theory of beta decay [1, 2].

The existence of the neutrino was finally proven by Reines, Cowan, et al. in 1953 [3] who detected electron antineutrinos coming from nuclear reactors. Within the decade, while in exile in the Soviet Union [4], Bruno Pontecorvo first proposed a theory of neutrino oscillations (initially in the context of neutrino-antineutrino mixing) [5].

In 1962 Leon Lederman, Jack Steinberger, et al. made the discovery of a second type of neutrino, the muon neutrino [6]. That same year, Maki, Nakagawa, and Sakata proposed the mixing-angle formalism of oscillations [7] that would allow a neutrino to oscillate from one flavor to another. In 1965, muon neutrinos were found in cosmic ray showers in the Kolar Gold Fields of South India and in South Africa [8, 9].

A curious puzzle had emerged by the 1970s: solar neutrino experiments were showing much lower event rates than expected based on the standard solar model [10]. This issue, known as the solar neutrino problem, would not see resolution until the end of the century.

By 1989, measurements of the Z width gave evidence that there were in fact three weak-interacting neutrino generations [11]. Indeed, in 2000 the DONUT collaboration discovered the final fermion in the standard model: the tau neutrino [12].

The resolution of solar neutrino problem came in the late 1990s in the form of the SNO and Super-Kamiokande experiments proving that the missing solar electron neutrinos had in fact oscillated into other flavors [13, 14].

1.2 Theory of Neutrino Oscillation

The theory of massive neutrinos begins with the notion that the neutrino mass eigenstates $|\nu_1\rangle$, $|\nu_2\rangle$, and $|\nu_3\rangle$ are not the same as the three flavor eigenstates $|\nu_e\rangle$, $|\nu_\mu\rangle$, and $|\nu_\tau\rangle$ [1, 15, 16].

Instead, there exists a 3×3 unitary change-of-basis matrix U_{PMNS} between the mass and flavor eigenbases:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.1}$$

This matrix is known as the PMNS matrix, named after Pontecorvo, Maki, Nakagawa, and Sakata [5, 17, 7].

A neutrino in a pure mass eigenstate is stationary under the action of the free-space Hamiltonian. However, a flavor eigenstate is a superposition of mass eigenstates. These mass eigenstates propagate independently, meaning that the overall time evolution of the neutrino state is more complicated. Let the neutrino begin in a flavor state α:

$$\begin{aligned} |\nu_\alpha(x,t)\rangle &= \sum_i U_{\alpha i} |\nu_i(x,t)\rangle \\ &= \sum_i U_{\alpha i} e^{i(p_i x - E_i t)} |\nu_i\rangle \\ &= \sum_{i,\beta} U_{\alpha i} U^*_{\beta i} e^{i(p_i x - E_i t)} |\nu_\beta\rangle \end{aligned}$$

Over time, a neutrino produced as a flavor eigenstate may pick up components belonging to other flavor states. This is the core of neutrino oscillation.

Oscillations in a Vacuum

Consider a neutrino produced in flavor α traveling in a vacuum. Because neutrinos are so light, we can assume they are very relativistic, so $t \approx x := L$, where L is the distance traveled by the neutrino. We can also use this fact to rewrite the energy-momentum relation for the individual mass eigenstates $E_i^2 - p_i^2 = m_i^2$:

$$p_i - E_i = -\frac{m_i^2}{E_i + p_i} \approx -\frac{m_i^2}{2E}$$

where, again assuming relativistic velocities, we denote the neutrino's energy as $E \approx p_i \approx E_i$.

Let us construct the inner product of a neutrino in the pure flavor state β and the time-evolved state of a neutrino emitted in the α state:

$$\begin{aligned}\langle \nu_\beta | \nu_\alpha(x,t)\rangle &= \sum_i U_{\alpha i} U^*_{\beta i} e^{i(p_i x - E_i t)} \\ &= \sum_i U_{\alpha i} U^*_{\beta i} e^{i(p_i - E_i)L} \\ &= \sum_i U_{\alpha i} U^*_{\beta i} \exp\left(-i \frac{m_i^2 L}{2E}\right)\end{aligned}$$

From this, we can calculate the transition probability:

$$\begin{aligned}P(\nu_\alpha \to \nu_\beta) &= \left|\langle \nu_\beta | \nu_\alpha(x,t)\rangle\right|^2 \\ &= \sum_{i,j} U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i} \exp\left(-i \frac{\Delta m^2_{ij}}{2} \frac{L}{E}\right) \end{aligned} \tag{1.2}$$

where $\Delta m^2_{ij} := m_i^2 - m_j^2$. As we shall see, certain Δm^2_{ij} are relevant in different oscillation contexts, so two are given special names: $\Delta m^2_{\text{sol}} := \Delta m^2_{21}$ from solar neutrino oscillations, and $\Delta m^2_{\text{atm}} := \Delta m^2_{32} \approx \Delta m^2_{31}$ from atmospheric oscillations.

We can simplify equation 1.2 by making the further substitution $\Delta_{ij} = \frac{\Delta m^2_{ij} L}{4E}$ and splitting it into real and imaginary parts:

$$\begin{aligned}P\left(\nu_\alpha \to \nu_\beta\right) &= \sum_{i,j} U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i} e^{-i2\Delta_{ij}} \\ &= \sum_i \left|U_{\alpha i} U^*_{\beta j}\right|^2 + \sum_{i \neq j} U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i} \cos 2\Delta_{ij} \\ &\qquad\qquad - \sum_{i \neq j} i U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i} \sin 2\Delta_{ij} \\ &= \delta_{\alpha\beta} - 4 \sum_{j>i} \mathfrak{Re}\left[U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i}\right] \sin^2 \Delta_{ij} \\ &\quad + 2 \sum_{j>i} \mathfrak{Im}\left[U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i}\right] \sin 2\Delta_{ij}\end{aligned} \tag{1.3}$$

Note that in the case of antineutrino oscillations, the oscillation probability is similar, with a change of sign to the $\mathfrak{Im}\left[U_{\alpha i} U^*_{\alpha j} U_{\beta j} U^*_{\beta i}\right]$ term.

The PMNS matrix is, in general, a complex unitary matrix. This would give it nine degrees of freedom: three real parameters, and six complex phases. However, we can absorb some of these complex phases into the original six flavor and mass eigenstates, leaving us with only a single complex degree of freedom that affects oscillations.

An immediate consequence of Eq. 1.3 is that if U_{PMNS} has an imaginary component, then the $\mathfrak{Im}\left[U_{\alpha i}U^*_{\alpha j}U_{\beta j}U^*_{\beta i}\right]$ term is nonzero and therefore $P(\nu_\alpha \to \nu_\beta) \neq P(\overline{\nu}_\alpha \to \overline{\nu}_\beta)$. This would violate CP symmetry, so the complex free parameter to the PMNS matrix is dubbed the CP-violating phase δ_{CP}.

The three real parameters are commonly written as three rotation angles, also known as the mixing angles: θ_{12}, θ_{13}, and θ_{23}. These are traditionally known as the solar, reactor, and atmospheric angles respectively.

In this parametrization, the PMNS matrix can be rewritten as follows:

$$\begin{aligned}
U_{PMNS} &= U_{\text{solar}} \times U_{\text{reactor}} \times U_{\text{atmospheric}} \\
&= \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \\
&= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix}
\end{aligned} \tag{1.4}$$

where $s_{ij} = \sin\theta_{ij}$ and $c_{ij} = \cos\theta_{ij}$. This parametrization for δ_{CP} was chosen especially because the mixing angle θ_{13} was known to be quite small. As a result, the overall PMNS matrix will have a small imaginary component, no matter the value of δ_{CP}.

This means that, to leading order, we can further simplify equation 1.3:

$$P\left(\overset{(-)}{\nu}_\alpha \to \overset{(-)}{\nu}_\beta\right) \approx \delta_{\alpha\beta} - 4\sum_{j>i} \mathfrak{Re}\left[U_{\alpha i}U^*_{\alpha j}U_{\beta j}U^*_{\beta i}\right] \sin^2 \frac{\Delta m^2_{ij} L}{4E} \tag{1.5}$$

From equation 1.5 we can readily see that transition probabilities show sinusoidal oscillations in L/E. The frequency of these oscillations is proportional to Δm^2_{ij}, and the amplitude of the oscillations is a product of PMNS matrix elements (and therefore a function of the mixing angles θ_{ij}). Because the oscillations go as $\sin^2\Delta_{ij}$, under our assumption that U_{PMNS} has small imaginary components, oscillation signals are the same no matter the sign of Δm^2_{ij}. As such, disambiguating the sign of a Δm^2_{ij} requires interactions between neutrinos and matter to adjust the oscillation probability, as will be discussed in the next section.

We can rewrite Δ_{ij}, the argument of the sine, by reintroducing suppressed units:

$$\Delta_{ij} = \frac{\Delta m^2_{ij} L}{4E} = 1.27 \times \frac{\Delta m^2_{ij}}{[\text{eV}^2]} \times \frac{L/E}{[\text{m/MeV}]} \tag{1.6}$$

This provides a good numerical rule-of-thumb for evaluating oscillations: if the L/E of a travelling neutrino (expressed in units of m/MeV) is the same order of magnitude as $\left(\Delta m^2_{ij}\right)^{-1}$ (expressed in units of eV^{-2}), then it will be dominated by oscillations due to Δm^2_{ij}. Put another way: the L/E range accessible in a given neutrino oscillation experiment determines which values of Δm^2_{ij} the experiment is most sensitive to. Experimentally, we see that $\Delta m^2_{21} \approx 7 \times 10^{-5}\ \mathrm{eV}^2$, and $\left|\Delta m^2_{32}\right| \approx 2 \times 10^{-3}$, so we should look for L/E ratios of $O\left(10^4 - 10^5\right)$ or $O\left(10^3\right)$ m/MeV respectively [15].

Substituting in PMNS matrix elements into 1.5 allows us to write vacuum oscillation probabilities as functions of mixing angles. As an example, here are the leading-order flavor transition and survival probabilities for muon neutrinos (Assuming $L/E \gg 1/\Delta m^2_{\mathrm{sol}}$, so the oscillation is dominated by $\Delta m^2_{31} \approx \Delta m^2_{32} = \Delta m^2_{\mathrm{atm}}$):

$$
\begin{aligned}
P(\nu_\mu \to \nu_e) &\approx \sin^2 2\theta_{13} \sin^2 \theta_{23} \sin^2 \left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right) \\
P(\nu_\mu \to \nu_\tau) &\approx \sin^2 2\theta_{23} \cos^4 \theta_{13} \sin^2 \left(\frac{\Delta m^2_{\mathrm{atm}} L}{4E}\right) \\
P(\nu_\mu \to \nu_\mu) &= 1 - P(\nu_\mu \to \nu_e) - P(\nu_\mu \to \nu_\tau)
\end{aligned}
$$

As a final note on vacuum oscillations, equation 1.3 can be used to show:

$$
\begin{aligned}
P(\nu_\mu \to \nu_e) - P(\overline{\nu}_\mu \to \overline{\nu}_e) &= 4 \sum_{j>i} \mathfrak{Im}\left[U_{\mu i} U^*_{\mu j} U_{ej} U^*_{ei}\right] \sin 2\Delta_{ij} \\
&= 16 J \sin \Delta_{21} \sin \Delta_{32} \sin \Delta_{31} \qquad (1.7)
\end{aligned}
$$

where $J = \mathfrak{Im}\left[U_{e1} U^*_{e2} U_{\mu 2} U^*_{\mu 1}\right]$ is the Jarlskog invariant. First defined analogously for quark mixing, J provides a parametrization-agnostic measure of CP violation [18]. In terms of the standard PMNS parametrization:

$$
J = \cos(\theta_{12}) \cos^2(\theta_{13}) \cos(\theta_{23}) \sin(\theta_{12}) \sin(\theta_{13}) \sin(\theta_{23}) \sin(\delta_{CP})
$$

If J is nonzero then CP symmetry is violated.

Oscillations in Matter

Neutrinos are subject to weak interactions as they pass through matter. This has practical importance to all neutrino experiments, whether the neutrinos are produced deep within the Sun and need to travel through its outer layers, or need to pass through Earth's crust before reaching a detector.

Weak interactions via the neutral current (NC) between neutrinos and atomic matter are the same regardless of flavor. As a result, it can be shown that NC interactions contribute only an unobservable phase change to neutrino interactions, so will be neglected for the rest of this section [15].

Charged current (CC) interactions, on the other hand, are only relevant for electron neutrinos and antineutrinos (see Figure 1.1), as there are essentially no muons or taus embedded in atomic matter.

(a) ν_e CC interaction (b) $\overline{\nu}_e$ CC interaction

Figure 1.1: Feynman diagrams depicting ν_e (left) and $\overline{\nu}_e$ (right) interactions via the charged current.

The effect of these charged current interactions can be thought of as a potential V_e that only affects electron neutrinos:

$$V_e(x) = \pm\sqrt{2}G_F n_e(x)$$

where G_F is the Fermi constant, $n_e(x)$ is the local electron density, and the sign is positive for neutrinos and negative for antineutrinos.

We can account for this potential by perturbing the kinetic Hamiltonian. In the mass basis, we can write the original as:

$$H_m = \frac{1}{2E}\mathrm{diag}\left(m_1^2, m_2^2, m_3^2\right)$$

while in the flavor basis, our modification to the Hamiltonian can be written as:

$$\Delta H_\alpha = \mathrm{diag}\left(V_e, 0, 0\right)$$

Thus, using U_{PMNS} to convert between the flavor and mass bases, the full modified Hamiltonian in the mass basis is:

$$H'_m = H_m + U^\dagger \Delta H_\alpha U$$

For simplicity, let us take the case of two neutrino flavors ν_e and ν_μ, and two vacuum mass eigenstates ν_1 and ν_2. Here, the PMNS matrix would have only one mixing angle and can be expressed as:

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}$$

The modified Hamiltonian will show two new mass eigenstates:

$$m^2_{1m,2m} = EV_e + \frac{m_1^2 + m_2^2}{2} \pm \frac{\Delta m^2}{2}\sqrt{\left(\cos(2\theta) + \frac{2EV_e}{\Delta m^2}\right)^2 + \sin^2(2\theta)}$$

where $\Delta m^2 = m_1^2 - m_2^2$. We can use these to construct a modified mass-squared difference:

$$\Delta m^2_m = \Delta m^2 \sqrt{\left(\cos(2\theta) + \frac{2EV_e}{\Delta m^2}\right)^2 + \sin^2(2\theta)}$$

We can also construct a modified PMNS matrix U' which connects the new mass eigenstates ν_{1m} and ν_{2m} to the flavor states. We can then define a modified mixing angle θ_m:

$$\tan 2\theta_m = \frac{\sin 2\theta}{\cos 2\theta - \frac{2EV_e}{\Delta m^2}}$$

As before, we can consider Δm^2_m to be the frequency of oscillations in matter, and the amplitude of oscillation is some function of θ_m. However, unlike before, Δm^2_m and θ_m are functions of V_e, which as discussed earlier is positive for neutrinos and negative for antineutrinos.

Importantly, because θ_m is also a function of Δm^2, oscillation signals are sensitive to the sign of Δm^2. This was not the case in a vacuum, so matter effects are useful for placing the relative order of the neutrino masses.

This means that matter effects would manifest as a difference in oscillation behavior between neutrinos and antineutrinos. This is similar to the signal we would expect from CP violation, and indeed in certain regimes of matter effects and values of δ_{CP}, these effects can compete with each other. For more on this tradeoff in the context of NOvA and T2K, see Section 3.3.

The MSW effect

Solar neutrinos provide an interesting case for studying matter effects. Electron neutrinos produced in nuclear reactions in the core of the Sun need to pass through the outer layers in order to escape into space. It turns out that, while in a vacuum $|\nu_e\rangle$ has a large component of $|\nu_1\rangle$, in solar matter it is mostly $|\nu_{2m}\rangle$. The Sun's density changes gradually enough that as the neutrino escapes, it is adiabatically pumped into the $|\nu_{2m}\rangle$ state [15]. As the neutrino leaves the solar medium, $|\nu_{2m}\rangle \to |\nu_2\rangle$. As a result, essentially all solar neutrinos are in the pure $|\nu_2\rangle$ mass eigenstate. This effect is due to Mikheev, Smirnov, and Wolfenstein, and is known as the MSW effect in their honor [19, 20].

Note that the MSW effect allows for the sign of Δm^2_{21} to be disambiguated. As such, only the sign of Δm^2_{32} (and therefore Δm^2_{31} as well) is still unknown.

Open Questions in Three-Flavor Oscillations

There are two open questions surrounding three-flavor oscillations that will be discussed in this thesis: CP violation and the mass ordering.

The extent to which neutrinos violate CP is currently unknown. A nonzero value of $\sin\delta_{CP}$ may have implications for theories of leptogenesis and may contribute to the matter-antimatter asymmetry observed in the universe [21, 22].

While the MSW effect disambiguates the sign of Δm^2_{sol}, we do not yet know the sign of Δm^2_{atm}. Figure 1.2 shows graphically the two options, known as the Normal Ordering (NO) and the Inverted Ordering (IO). These are also known in literature as the Normal Hierarchy (NH) and Inverted Hierarchy (IH) respectively.

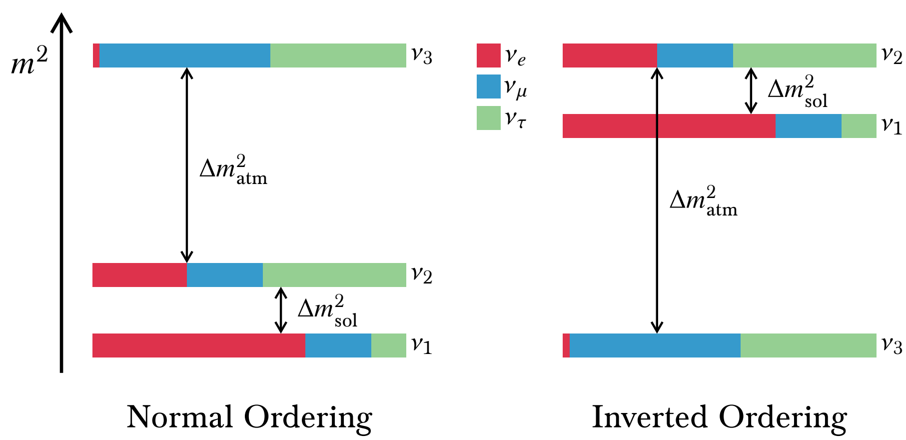

Figure 1.2: Diagram depicting the two possible neutrino mass orderings. Colors depict the approximate flavor composition of each mass eigenstate. In the Normal Ordering (NO), ν_3 is the heaviest eigenstate (left); in the Inverted Ordering (IO), ν_3 is the lightest eigenstate (right).

Neutrinos being in the IO would be especially interesting, as this would provide constraints on the sum of neutrino masses for cosmological experiments. Determining the mass ordering would also provide context necessary for interpreting neutrinoless double beta decay measurements [23–25].

1.3 Experimental Status of Neutrino Oscillation

Solar Neutrinos

An enormous number of electron neutrinos are produced in the Sun. The spectral flux of solar neutrinos is shown in Figure 1.3.

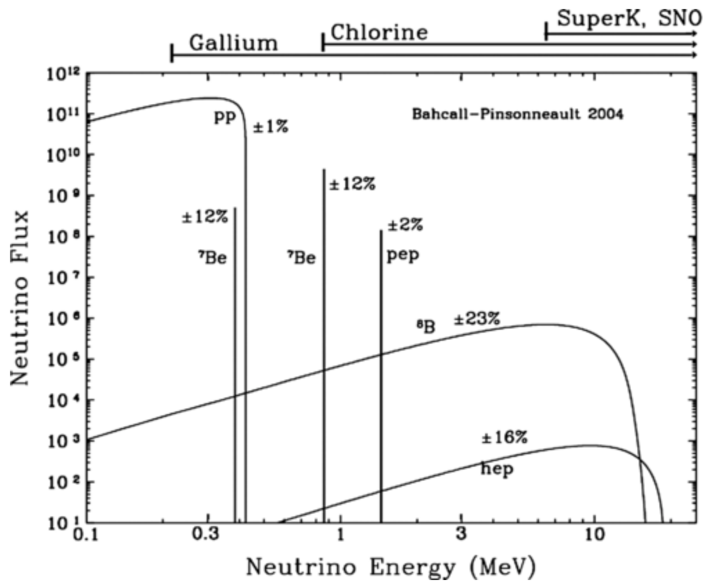

Figure 1.3: Flux of solar neutrinos from various solar processes at Earth's orbital distance. Continuous fluxes are in units of $\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{MeV}^{-1}$, while the line fluxes are in units of $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [26]. Marked above the plot are the sensitivity ranges of various solar neutrino detector technologies. Liquid Scintillator detectors like BOREXINO and SNO+ have sensitivities have the most sensitivity to ^{8}B.

Because of the MSW effect, solar neutrinos are the best ways to interrogate θ_{12} as well as Δm^2_{21}. However, almost all solar neutrinos are well under the ~30 MeV minimum energy required to create a muon in a charged-current interaction. Thus, to observe solar neutrino oscillations we need to observe one of three phenomena: ν_e charged-current (ν_e CC) neutrino-nucleus interactions, neutral-current (NC) neutrino-nucleus interactions (which can happen with neutrinos of any flavor), or neutrino-electron elastic scattering (ES) events (which can happen with neutrinos of any flavor, but favor electron neutrinos).

The Super-Kamiokande detector (SK), an enormous water Cherenkov detector, has excellent angular resolution [14]. This let SK make a measurement of neutrino-electron elastic scattering, by observing an excess of 5-20 MeV neutrinos coming from the direction of the Sun.

The Sudbury Neutrino Observatory (SNO) used heavy water as a detection medium. ν_e CC events were detectable via the Cherenkov ring associated with the final state electron. Additionally, if an incident neutrino had enough energy, NC interactions could cause deuterons to break apart. The resultant free neutron could capture on another atom, releasing a detectable electromagnetic shower.

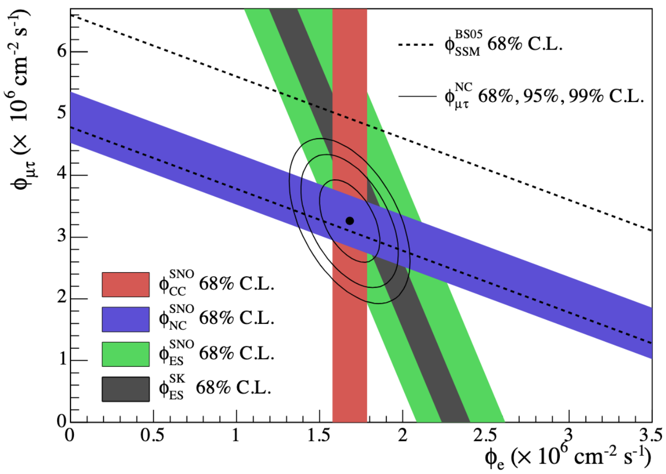

Figure 1.4: Evidence of solar oscillations. The three detection methods for solar neutrino methods yield different conclusions about the solar neutrino flux. ν_e CC events only provide information about the ν_e flux at Earth. ES events provide some information about the ν_μ and ν_τ fluxes at Earth, but are dominated by ν_e events. Finally, NC events are identical regardless of neutrino flavor. The joint fit shows that the overall neutrino flux is consistent with the standard solar model (dashed lines) [13].

Reactor Neutrinos

Nuclear reactors provide a very pure, very intense source of ~1 MeV electron antineutrinos. Since the L/E of short-baseline reactor experiments are around $10^2 - 10^3$ m/MeV, reactor neutrinos are dominated by Δm^2_{atm} oscillations. From Equation 1.2, it can be shown that to leading order at this L/E regime, the size of the $\overline{\nu}_e$ disappearance signal is dependent on θ_{13}.

The Daya Bay Reactor Neutrino Experiment was the premier experiment measuring short-baseline reactor oscillations. The eight functionally identical detectors were located in three experimental halls. Each detector was filled with a liquid scintillator, doped with gadolinium. Two were close to the Daya Bay and Ling Ao nuclear reactors, and the third hall was ~1 km away from the six reactor cores. This allowed Daya Bay to directly observe the L/E oscillation dip (see Figure 1.5).

The primary detection mechanism was inverse beta decay (IBD):

$$\overline{\nu}_e + p \rightarrow n + e^+$$

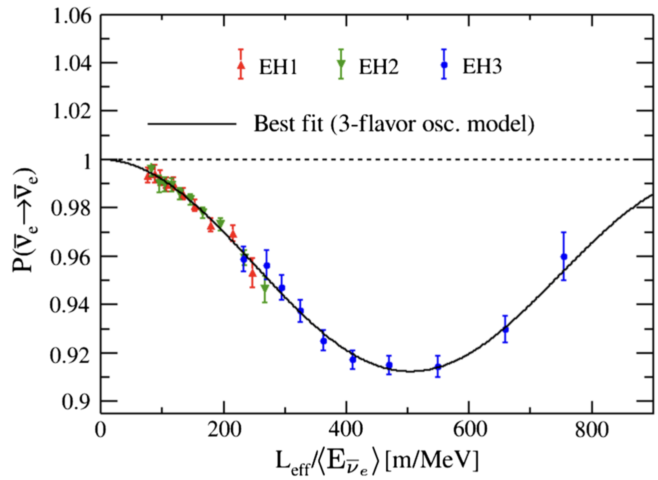

Figure 1.5: Electron Antineutrino survival probability as a function of $L_{\text{eff}}/\langle E_{\bar{\nu}_e} \rangle$, where L_{eff} is the effective baseline, and $\langle E_{\bar{\nu}_e} \rangle$ is the average antineutrino energy at that baseline. Data from the far experimental hall, EH3, shows clear indications of an oscillation dip [27].

A prompt signal from annihilation of the positron would be followed by a delayed signal from the capture of the neutron on the gadolinium dopant. This provided a robust signal of an IBD event. Daya Bay was able to measure θ_{13} to world-leading precision [27].

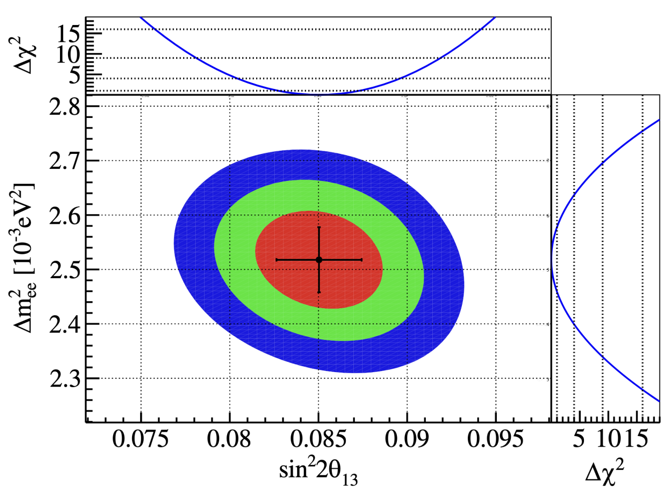

Figure 1.6: Daya Bay's most recent oscillation results [27]. Instead of plotting Δm^2_{32} or Δm^2_{31}, Daya Bay reports mass-squared differences in Δm^2_{ee}, a closely related empirical quantity based on a two-flavor approximation. Specifically, $\Delta m^2_{ee} \simeq \cos^2\theta_{12}\left|\Delta m^2_{31}\right| + \sin^2\theta_{12}\left|\Delta m^2_{32}\right| \approx \left|\Delta m^2_{32}\right|$ [28].

As a final note, long baseline (~100 km) reactor neutrino experiments like KamLAND have an L/E values around $10^4 - 10^5$ m/MeV [29]. This gave them sensitivity to solar oscillations, as that E/L is close in magnitude to Δm^2_{21}. This allowed KamLAND to make measurements of θ_{12}.

Atmospheric Neutrinos

High energy cosmic rays can interact with matter in Earth's atmosphere, inducing extremely energetic hadronic showers. Decaying mesons from these showers will produce predominantly muon neutrinos, but ν_e production is also common at lower energies. Atmospheric ν_τ production also occurs beyond the TeV scale [30].

Figure 1.7 shows the energy and zenithal flux of atmospheric neutrinos. Most atmospheric neutrinos are produced ~20 km above Earth's surface, but since neutrinos will pass through the matter in Earth easily, the oscillation baseline can vary from 10^4 to 10^7 m. Atmospheric neutrinos are also produced a wide range of energies (10^2 to 10^5 MeV), meaning atmospheric neutrinos can probe oscillations with $\left|\Delta m^2\right| \approx 10^{-1}$ to 10^{-4}eV^2. This puts Δm^2_{32} in range to dominate oscillations.

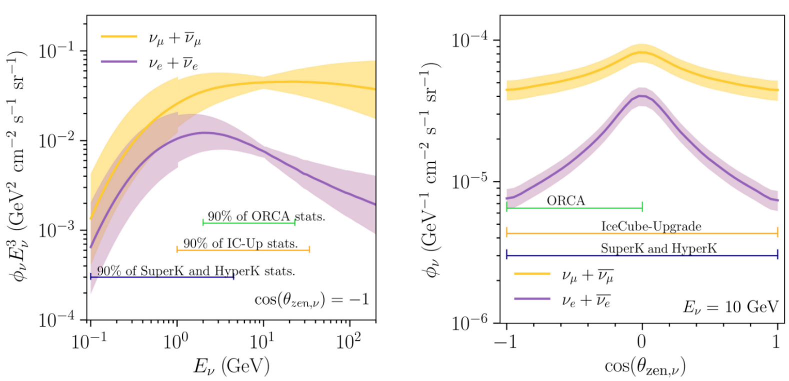

Figure 1.7: Energy (left) and zenithal angle (right) fluxes of atmospheric neutrinos at Earth's surface. Marked are the sensitivities of current and future experiments. All neutrino flavors show higher fluxes at angles close to horizontal ($\cos\theta_{\text{zen}} = 0$), as these neutrinos come from mesons which are able to pass through more of Earth's atmosphere before hitting Earth's surface and losing energy [30]. There is also more solid angle perpendicular to the zenith, leading to a geometric bias towards $\cos\theta_{\text{zen}} = 0$.

Because upward-going neutrinos must have traveled through Earth, these neutrinos are subject to matter effects (Section 1.2). This gives atmospheric neutrinos a handle to disambiguate the neutrino mass ordering.

Super-Kamiokande's excellent directional reconstruction, coupled with its ability to distinguish $\overset{(-)}{\nu}_e$ vs $\overset{(-)}{\nu}_\mu$, allowed them to make measurements of Δm^2_{32} and θ_{23} [31].

IceCube is an enormous Cherenkov detector that uses the Antarctic ice sheet at the South Pole. It is able to detect neutrinos from an extremely broad range of energies. In particular, the DeepCore region of the detector has a higher density of photodetectors, enabling the detection of neutrino events down to a few GeV, giving it sensitivity to atmospheric neutrino oscillations (Figure 1.8) [32].

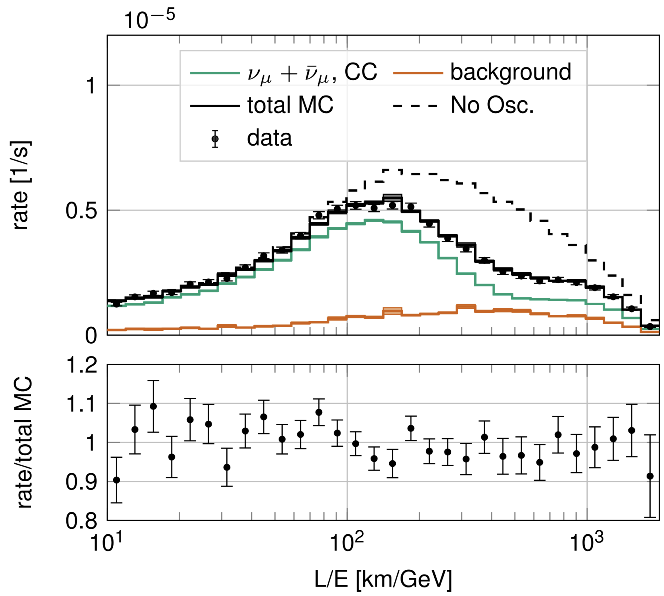

Figure 1.8: The IceCube Experiment is able to clearly measure the ν_μ oscillation dip region due to the DeepCore region of the detector [32].

Long Baseline Accelerator Neutrinos

The measurement of long-baseline (LBL) oscillations of accelerator neutrinos is the primary subject of this thesis.

While atmospheric neutrinos are typically $\overset{(-)}{\nu}_\mu$, the initial production site, flavor, and energy are unknown. As such, a high-energy $\overset{(-)}{\nu}_\mu$ source with a known baseline to a detector would provide a valuable asset for measuring neutrino oscillations. This can be accomplished using a particle accelerator. Protons with energies of ~10 to ~100 GeV can impinge on a target to produce a shower of high energy mesons, especially charged pions and kaons. Almost all charged pions are subject to the decay $\pi^\pm \to \mu^\pm + \overset{(-)}{\nu}_\mu$. A large majority of kaons also decay via analogous processes to produce high-energy $\overset{(-)}{\nu}_\mu$s. By using a strong focusing magnet and collimators, accelerators can produce a beam of either π^+ or π^-, whose decays will produce ν_μ or $\overline{\nu}_\mu$ respectively. The net result is a 100 MeV- to GeV-scale beam of mostly-collimated muon neutrinos or antineutrinos.

These neutrinos can be measured before significant oscillations at a near detector at a short baseline (~1 km) to obtain an initial neutrino spectrum before oscillations. Then, at a long baseline from the beam source, a large far detector can measure the oscillated flux.

With neutrino energies of $10^2 - 10^3$ MeV and baselines of $10^5 - 10^6$ m, LBL neutrino oscillation experiments will have greatest sensitivity to Δm^2_{32} oscillations, and as such are designed to provide a handle to measure θ_{23}. In order to travel the long baselines, beams must be angled slightly down through Earth to reach a far detector. Thus, neutrinos will experience strong matter effects, allowing for possible dissambiguation of the mass ordering. Since the parent meson charge can be chosen via the magnetic focusing horns, LBL experiments also provide a handle to measure CP violation as a difference in oscillation behavior between ν_μ and $\overline{\nu}_\mu$.

The first experiment to follow this paradigm was the KEK-to-Kamioka (K2K) experiment, which leveraged the KEK Proton Synchrotron in Ibaraki Prefecture, Japan, to produce a 1-1.5 GeV beam of muon neutrinos aimed at the Super-Kamiokande detector. SK therefore functioned as K2K's far detector, located at a baseline of 250 km from the beam source. The near detector was a much smaller 1 kt water Cherenkov detector near the beam source, constraining the beam flux and other systematic uncertainties. Operating from 1999 to 2005, K2K was able to measure neutrino oscillations.

The Main Injector Neutrino Oscillation Search (MINOS) turned on in 2006 and measured oscillations using the NuMI (Neutrinos at the Main Injector) beam at Fermilab in Illinois, USA. MINOS consisted of a ~1-kt near detector and a 5.4-kt far detector, located ~1 and 735 km from the beam source respectively. Both detectors were built of steel and plastic scintillator in a toroidal magnetic field [33]. After 2012, with the construction of NOvA and upgrades to the NuMI beam, MINOS continued taking data as MINOS+ until 2016 [34].

The current generation of LBL neutrino oscillation experiments are the successors to MINOS and K2K: NOvA and T2K. These experiments and and a new joint data analysis between them will be discussed at length in the following chapters. For more about the NuMI beam and the NOvA experiment, see Chapter 2. For more about the T2K experiment, see Section 3.1.

1.4 Neutrino Cross Section Measurements

For precise measurements of neutrino oscillations, it is critical to reduce sources of systematic error. One important area of exploration is in neutrino cross-section measurements.

MINERvA (Main Injector Neutrino ExpeRiment to study ν-A interactions) was a successful experiment to study neutrino interactions on atomic nuclei. The experiment had a unique design with several passive interaction targets surrounded by a tracking calorimeter (see Figure 1.9). [35]

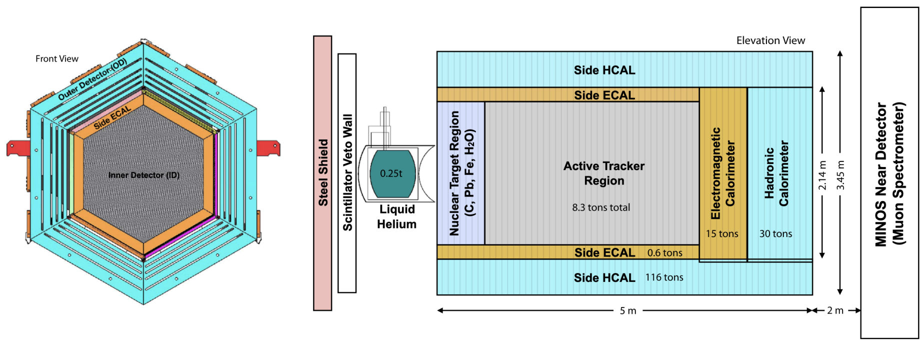

Figure 1.9: Schematic of the MINERvA Experiment. The "Nuclear Target Region" consists of planes of nuclear targets interspersed with more tracking scintillator planes. [35]

Notably, MINERvA used the NuMI beam, like NOvA and MINOS. In fact, MINERvA used the front planes of the MINOS near detector to range out muons. One important consequence is that MINERvA data is relevant to NOvA for tuning cross section models. Indeed, MINERvA's tune of the neutrino event generator GENIE proved important for NOvA and the NOvA-T2K joint fit (see Section 4.10) [36].

MINERvA has been able to measure ν_μ cross-sections on a variety of targets at several energy ranges [37–39]. In addition, it has also produced detailed measurements of the NuMI flux, probed pion and kaon production from ν_μ interactions, and measured the axial form factor of $\overline{\nu}_\mu$-proton scattering [40–43].

1.5 Future Oscillation Projects

Progress has been made on solving the outstanding questions of neutrino oscillation physics. As will be presented in Chapter 4, NOvA and T2K have made precise measurements of $\left|\Delta m^2_{32}\right|$ and have constrained regions of CP-violating space. Additional data may help strengthen these measurements, but in all cases, we need to look to the next generation of oscillation experiments for definitive information.

DUNE

The Deep Underground Neutrino Experiment (DUNE) will be Fermilab's next flagship long baseline neutrino oscillation experiment. DUNE will be primarily using Time Projection Chambers (TPCs) to detect and track charged particles, in contrast to scintillator-based calorimetry like the preceding MINOS and NOvA.

DUNE will consist of a new neutrino beam called the Long Baseline Neutrino Facility (LBNF), a near detector complex, and a far detector complex [44, 45].

LBNF will be fed from an upgraded Fermilab accelerator complex, and will produce a 1.2 MW $\overset{(-)}{\nu}_\mu$ beam, with potential upgrades bringing the beam power up to 2.4 MW. The $\overset{(-)}{\nu}_\mu$ energy spectrum will be peaked at 2.5 GeV [44].

The near detector complex is depicted in Figure 1.10. It will eventually consist of three main detectors: a liquid-argon TPC (ND-LAr), a gaseous-argon TPC (ND-GAr), and the System for on-Axis Neutrino Detection (SAND) [45].

Uniquely, the two argon TPCs will be movable, enabling beam neutrino measurements at a variety of off-axis angles. This program, known as the DUNE Precision Reaction-Independent Spectrum Measurment (DUNE-PRISM), will enable DUNE to greatly constrain neutrino interaction models by selectively intercepting neutrinos at different energies [45]. This enables the DUNE near detectors to construct, through superposition, arbitrary ND spectra. For more about how off-axis beam angles impact neutrino spectra in the context of the NuMI beam, see Section 2.2.

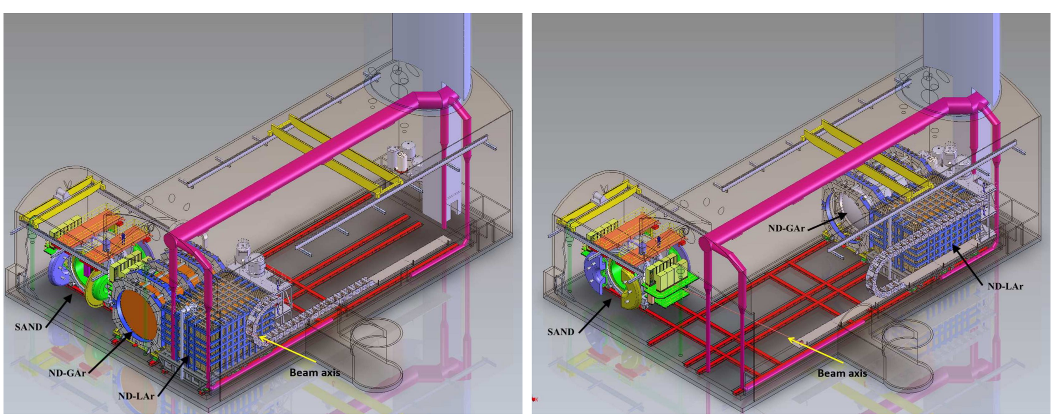

Figure 1.10: Layout of the DUNE ND Detectors. The DUNE-PRISM program will allow the two TPC detectors (ND-LAr and ND-GAr) to move from an on-axis (left) and to various off-axis angles (right). The SAND detector will remain on axis for beam monitoring. The ND suite will be used to constrain flux and interaction uncertainties, strengthening oscillation measurements from the FD [45]. Note that ND-GAr will be built at a later date. A muon range stack will take its place during DUNE Phase-I [46].

The far detectors will be enormous liquid-argon TPCs located at the Sanford Underground Research Facility (SURF) in Lead, South Dakota. SURF is in the former Homestake Gold Mine (which was the site of the experiment mentioned in Section 1.1 that discovered the solar neutrino problem) and is at a 1300 km baseline with respect to LBNF. The initial plans for the DUNE FD call for 20 kt worth of fiducial volume, upgradeable to 40 kt [44]. Due to the excellent energy reconstruction possible in a TPC, DUNE's far detectors will be on-axis relative to the LBNF beam, unlike current-generation long-baseline oscillation experiments. Each module will consist of a crystotat containing 17.5 kt of liquid Argon, and a fiducial volume of 10 kt.

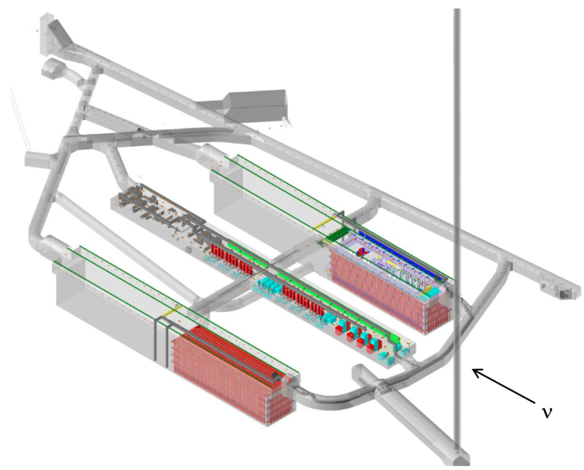

Figure 1.11: Schematic of the DUNE Far Detector cavern. The two planned FD modules will have a fiducial volume of 10 kt each. There will be enough cavern space for up to two additional modules to be installed at a later date [44].

The 1300 km baseline means that neutrinos produced by LBNF will experience strong matter effects before detection in SURF. These matter effects will lead to strong signals in favor of the correct mass ordering in fairly short order. Figure 1.12 shows how the matter effects strongly break the degeneracy of the mass ordering. DUNE will be able to determine the mass ordering to 5σ for 100% of δ_{CP} values within 2-3 years of turning on [44].

DUNE will also be able to measure CP violation ($\delta_{CP} \neq 0$ or π) to 5σ for 75% of δ_{CP} values during its run. If CP violation is maximal ($\delta_{CP} = \pm\frac{\pi}{2}$), DUNE will be able to measure CP violation within seven years of staged operation [44].

Prototypes of ND-LAr are being operated and tested in front of the NuMI beam currently. The FD caverns have been excavated, and DUNE is currently expected to turn on around 2030.

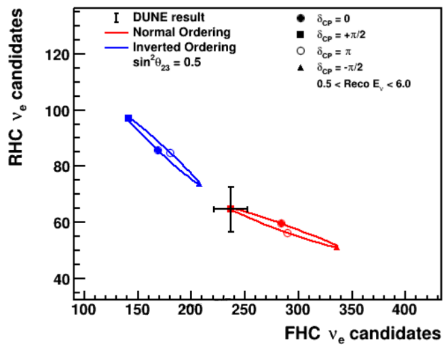

Figure 1.12: Hypothetical bi-event plot for the DUNE after four years of data taking. The "DUNE result" hypothetical data point shows that even accounting for statistical uncertainty there is no degeneracy between the normal and inverted mass orderings at DUNE. Note that this plot integrates over a broad energy range, so actual sensitivity will be even higher [47]. See Section 3.3 for further discussion of bi-event plots in the context of NOvA and T2K.

Hyper-Kamiokande

After nearly three decades of Super Kamiokande's operation, its successor Hyper-Kamiokande is now under construction. Hyper-Kamiokande (HK), like SK, is a water Cherenkov detector, but will be much larger. HK will have a total target volume of 258 kton of water, 187 kton of which will be the fiducial volume [48]. This gives HK a fiducial volume nearly 9 times larger than that of SK.

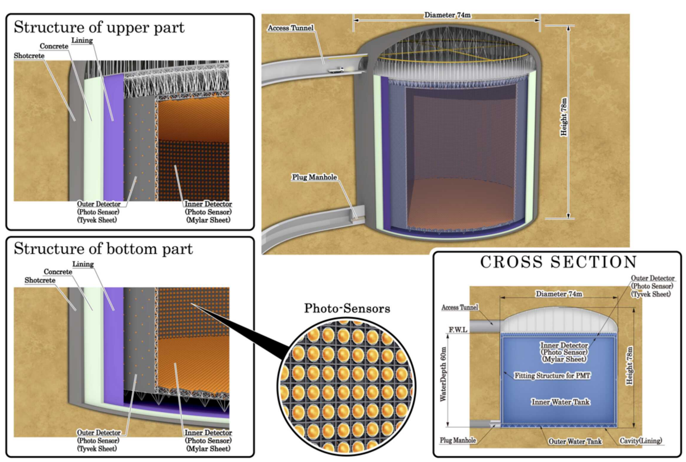

Figure 1.13: Schematic Diagram of the Hyper-Kamiokande Experimental site.

HK will be located 8 km south of SK in the Tochibora mine. This means that the HK detector will be able to intercept the J-PARC neutrino beam, at approximately the same 2.5° off-axis angle as SK relative to the J-PARC neutrino beam. Like SK, HK will be able to conduct analyses on both atmospheric as well as accelerator neutrinos.

For more about the J-PARC beam and how SK detects neutrinos, see Section 3.1.

Like T2K, HK will be using the ND280 complex at J-PARC to monitor the neutrino beam flux and composition. In addition to the ND280 complex, there are plans to build a movable ~800 tonne water Cherenkov detector roughly 1 km from the J-PARC beam source. Such a water Cherenkov detector could also be loaded with gadolinium for analysis of neutron production, among other ways it can be used to help reduce detector systematic errors at HK [48].

During its nominal run, HK will be able to measure δ_{CP} to 5σ for 60% of true δ_{CP} values. The lower beam energy than DUNE causes weaker matter effects. This can lead to degeneracies at certain values of δ_{CP} depending on the mass hierarchy. As such, combining beam and atmospheric measurements can break this degeneracy and lead to robust δ_{CP} measurements [49].

In addition to the detector currently under construction at Tochibora, there are discussions for a second identical water Cherenkov detector to be built at the second oscillation maximum. The current preferred site for such a second detector would be at Mount Bisul in Korea, located at a baseline of 1,088 km and an off-axis angle of 1.3° relative to the J-PARC neutrino beam. This would greatly increase physics sensitivities for both neutrino oscillations and other astrophysical measurements [48].

The main cavern for HK in Tochibora is currently under construction, with the top dome of the cavern excavated by early 2024. HK is currently expected to start taking data by 2027.

JUNO

The Jiangmen Underground Neutrino Observatory (JUNO) is an upcoming reactor neutrino experiment in southern China. The detector will consist of a 20 kt liquid scintillator volume [50] located at a baseline of ~50 km relative to several nuclear reactor cores. Like the Daya Bay Experiment (Section 1.3), JUNO will detect reactor antineutrinos via inverse beta decay. Unlike Daya Bay, however, the detector is much larger, the baseline is much longer, and the liquid scintillator will not be loaded with gadolinium. These differences give JUNO a different mechanism to measure neutrino oscillations.

JUNO's sensitivity to the neutrino mass ordering comes from the interference between Δm^2_{32} and Δm^2_{31} oscillations, as depicted in Figure 1.15a.

As a result, JUNO may be able to measure the neutrino mass ordering to 3σ signif-

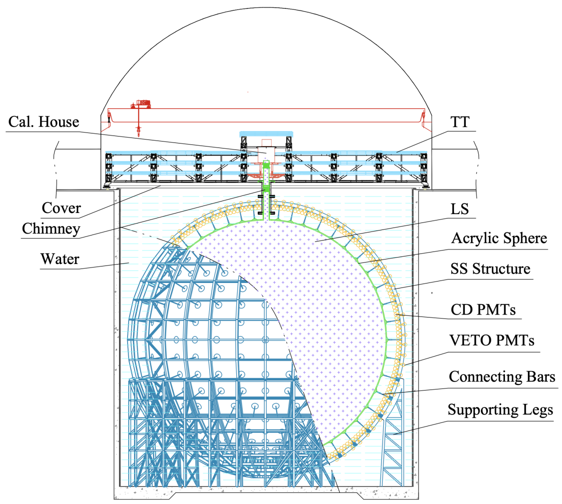

Figure 1.14: Schematic of the JUNO Detector. The liquid scintillator detector will be surrounded by a water Cherenkov detector which provides shielding as well as vetos [50].

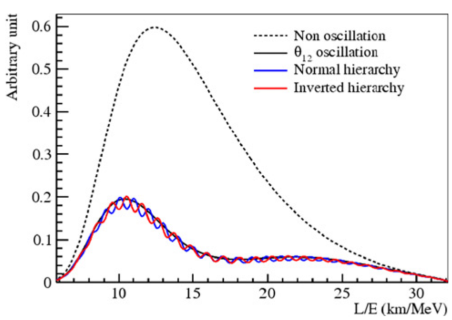

(a) JUNO MO Signal

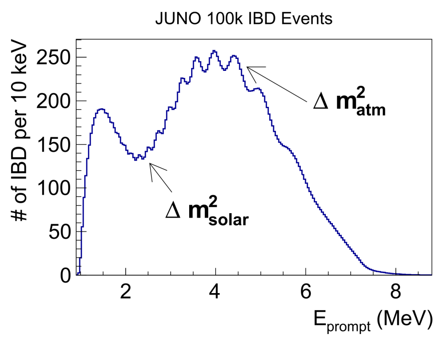

(b) JUNO Spectrum

Figure 1.15: JUNO's ability to distinguish the MO is due to the interplay between Δm^2_{32} and Δm^2_{31} (left). Due to JUNO's medium baseline, it will be able to see oscillations associated with both Δm^2_{21} and Δm^2_{32} (right) [51].

icance within six years of operation. JUNO's future results will be complementary with the LBL experiments, meaning that combined measurements will have increased sensitivities [51].

As of December 2024, JUNO is filling its detectors, with first data collection to begin in August 2025 [52].

Chapter 2

THE NOvA EXPERIMENT

2.1 The NuMI Beam

The Neutrinos at the Main Injector (NuMI) beam has served several experiments, notably MINOS, MINERvA, and NOvA.

Figure 2.1 depicts the Fermilab accelerator complex. Ionized hydrogen is sent through the linear accelerator, accelerating the protons from 750 keV to 400 MeV. After exiting the linac, protons accelerate to 8 GeV in the Booster ring and are transferred in bunches to the Recycler. Protons are then sent to the Main Injector, with which the Recycler shares a tunnel. The Main Injector accelerates protons up to a final energy of 120 GeV. Notably, new bunches of protons recharge the Recycler while the Main Injector accelerates its proton bunches. This reduces NuMI's cycle time to 1.3 s [53].

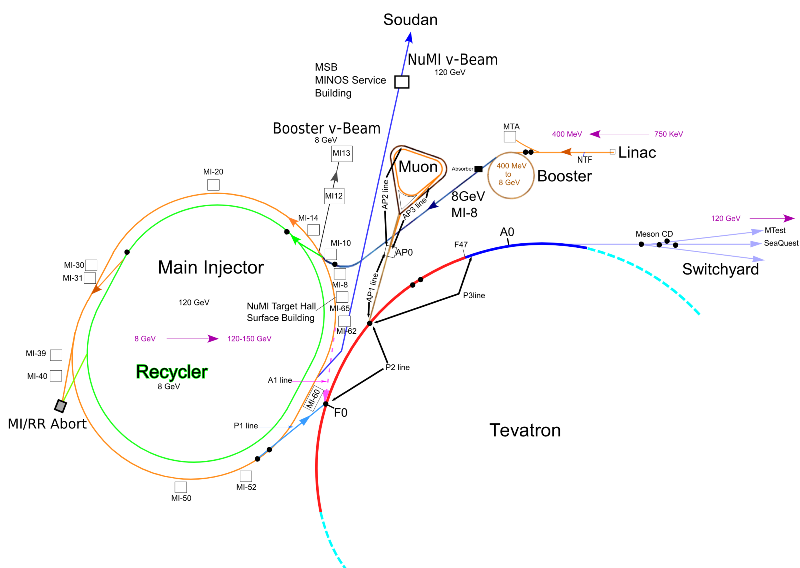

Figure 2.1: The Fermilab Accelerator Complex. Protons reach a peak energy of 120 GeV before reaching the NuMI target hall [53].

At the NuMI Target hall, protons impinge on a target composed of 48 graphite "fins" (Figure 2.2). The resultant cascade of daughter particles passes through a pair of magnetic focusing horns, which deflect and focus charged pions into a beam. These pions will decay into $\mu^{\pm}$ and $\overset{(-)}{\nu}_{\mu}$ in the 675 m long decay pipe. Any hadrons that pass through the horns and decay pipe are absorbed in the 5m thick absorber of aluminum, concrete, and steel, while the remaining muons are absorbed in the 200 m of rock between the NuMI target hall and the MINOS/NOvA Near Detector caverns.

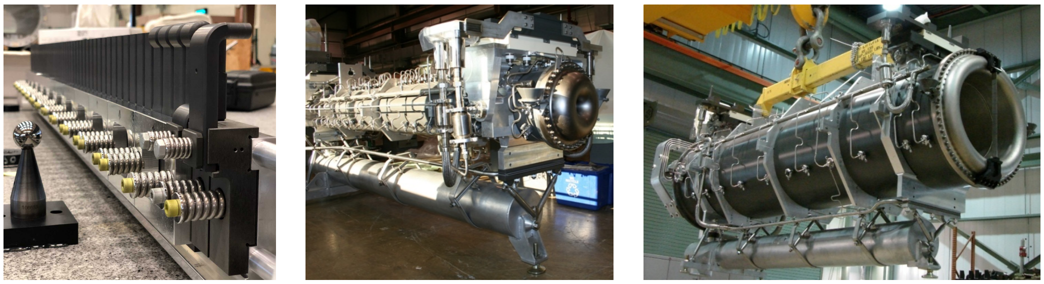

Figure 2.2: The NuMI target (left), and the two magnetic focusing horns (center, left), capable of producing a 1 MW beam [54]

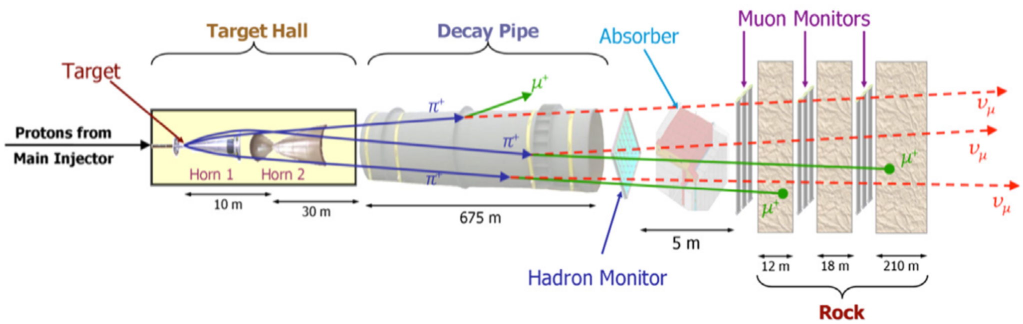

Figure 2.3: Schematic of the NuMI beam in the Forward Horn Current configuration, producing a predominantly ν_μ beam. Because the horn current can be reversed, the NuMI beam is also capable of producing a $\overline{\nu}_\mu$-dominated beam [53].

The current going to the horns is reversible. In the Forward Horn Current (FHC) configuration, the horns will focus π^+, which produce ν_μ after decaying, while in the Reverse Horn Current (RHC) configuration, the horns focus π^-, which decay into $\overline{\nu}_\mu$.

The final neutrino beam will not be pure ν_μ in FHC-mode or pure $\overline{\nu}_\mu$ in RHC-mode, since there will be significant contributions from decaying kaons, muons, and wrong-sign pions. These decays will contaminate the neutrino beam with $\overset{(-)}{\nu}_e$'s and wrong-sign $\overset{(-)}{\nu}_\mu$'s. The intrinsic $\overset{(-)}{\nu}_e$ component is an important background for $\overset{(-)}{\nu}_\mu \rightarrow \overset{(-)}{\nu}_e$ oscillation measurements. Observations of NuMI flux from the NOvA Near Detector can help constrain this background. Finally, the wrong-sign $\overline{\nu}_\mu$ contamination of FHC beam and ν_μ contamination of RHC beam is an important consideration for NOvA, especially because NOvA is not magnetized and therefore does not have the ability to detect muon charge.

After first proton delivery in 2004, NuMI operated at a beam power of 400 kW [55], beamline components were upgraded to 700 kW in 2012 for the beginning of NOvA's data taking. As NOvA came online, the power of the NuMI beam was slowly ramped up until it finally reached the designed power of 700 kW in January 2017 [56]. In 2019, NuMI went through another long shutdown to replace the target and magnetic focusing horns in order to achieve a beam power of 1 MW. NuMI briefly reached this power in June 2024 [57]. The total integrated exposures of the ND and FD up to May 31, 2024 are shown in Figure 2.4.

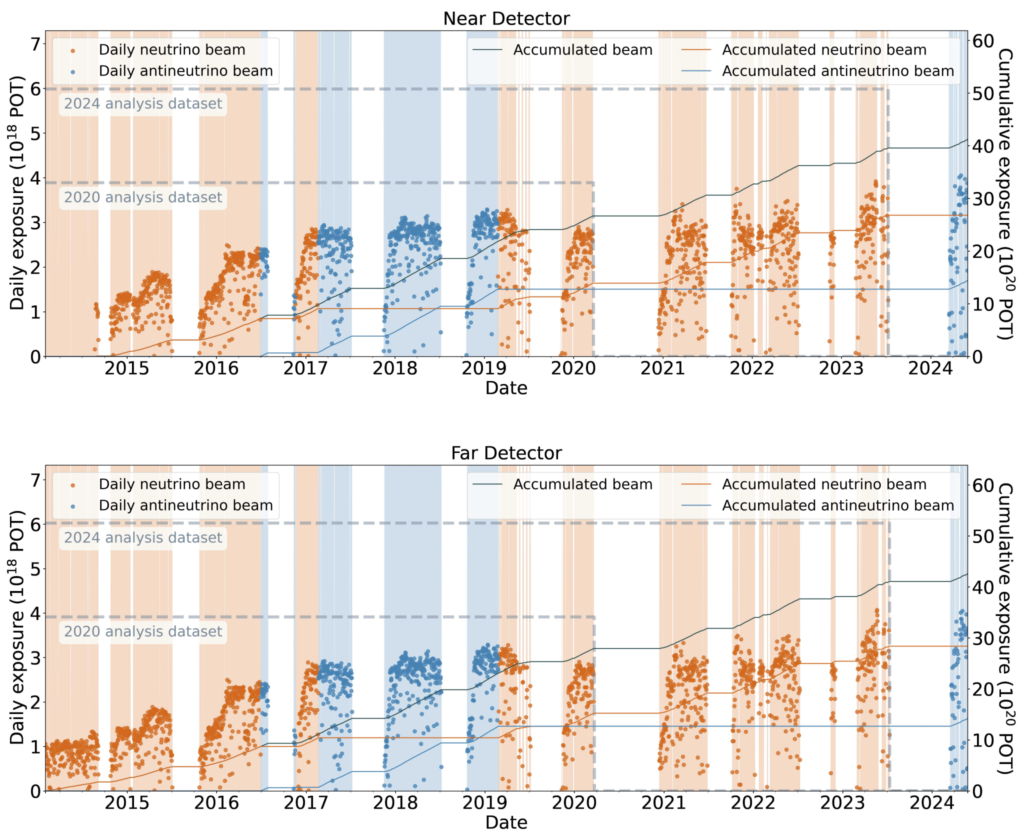

Figure 2.4: Integrated and daily exposures for the NOvA Near (top) and Far (bottom) detectors up to May 31, 2024. The gray boxes indicate the exposures used for the 2020 and 2024 NOvA 3-flavor analyses [58, 59].

Due to electrical transformer issues limiting power draw at Fermilab, NuMI will not be operating in FY25.

2.2 Off-Axis Neutrino Beams

Beam neutrinos primarily come from pion decay, which has one dominant decay mode: $\pi^{\pm} \to \mu^{\pm} + \overset{(-)}{\nu}_{\mu}$. In the center of momentum frame, the decay products are monoenergetic, with the neutrino having energy $E_{\nu}^{\mathrm{CM}} \approx 29.8$ MeV.

In the lab frame, say the pion is moving with some large speed β in the z-direction. If we take the neutrino's angle relative to the beam direction to be α in the CM frame, we can compute the Lorentz-boosted four-momentum of the neutrino in the lab frame:

$$p_\nu^\mu = \begin{pmatrix} \gamma E_\nu^{\mathrm{CM}}(1+\beta\cos\alpha) \\ E_\nu^{\mathrm{CM}}\sin\alpha \\ 0 \\ \gamma E_\nu^{\mathrm{CM}}(\beta+\cos\alpha) \end{pmatrix} \tag{2.1}$$

If the neutrino's angle relative to the axis is θ in the lab frame, we can write the following expression:

$$\cos\theta = \frac{p_\nu^z}{|\vec{p}_\nu|} = \frac{p_\nu^z}{E_\nu} = \frac{\beta+\cos\alpha}{1+\beta\cos\alpha}$$

Solving for $\cos\alpha$, we get:

$$\cos\alpha = \frac{\cos\theta-\beta}{1-\beta\cos\theta} \tag{2.2}$$

Substituting this expression into Eq (2.1), we see that the neutrino's energy in the lab frame is:

$$E_\nu = \frac{E_\nu^{\mathrm{CM}} m_\pi}{1-\beta\cos\theta}\cdot\frac{1}{E_\pi} \tag{2.3}$$

where we use the fact that $E_\pi = \gamma m_\pi$ is the energy of the parent pion. Note that when off-axis ($\theta > 0$), at high values of E_π we see that $E_\nu \propto \frac{1}{E_\pi}$.

Alternatively, we can rewrite Eq. (2.3) to an equivalent formulation:

$$E_\nu = \frac{1-\beta^2}{1-\beta\cos\theta}\cdot\frac{E_\nu^{\mathrm{CM}}}{m_\pi}E_\pi \tag{2.4}$$

From this formulation, we can see that at small values of E_π, or on-axis ($\theta = 0$), $E_\nu \propto E_\pi$.

Thus, we see that when off-axis, E_ν begins rising linearly with E_π before leveling off and falling as $\frac{1}{E_\pi}$. Figure 2.5 shows this behavior graphically, highlighting the off-axis angles of NOvA and T2K. Note that the vast majority of pions produced at NuMI are around 10-20 GeV.

The upshot is that off-axis, the energy spectrum of neutrinos is more sharply peaked. This means that the L/E of neutrinos from the beam source will be more constant, meaning that proportionally more neutrinos will be in the oscillation region. Also, an overly broad beam can lead to backgrounds from NC events in the oscillation regions. Because not all the energy exchanged in an NC interaction will be visible, high-energy improperly reconstructed NC events may enter the oscillation region. Thus,

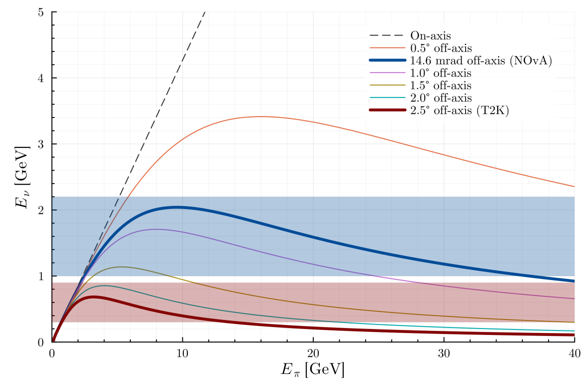

Figure 2.5: Neutrino Energy as a function of parent pion energy at several off-axis angles relative to the beam center. The bolded curves correspond to the off-axis angles of NOvA (blue) and T2K (red). The shaded bands indicate the oscillation dip regions corresponding to each experiment's baselines.

by keeping the beam spectrum more sharply peaked, there is less ability for these NC backgrounds to affect the oscillation measurement.

The NuMI Off-Axis ν_e Appearance experiment (NOvA) was designed to take advantage of this effect. The far detector site was chosen to have an off-axis angle of 14.6 mrad and is as far as possible from the NuMI beam source without leaving the United States. The far detector site is situated along the Ash River, just south of Voyageurs National Park and the Canadian border (Figure 2.6).

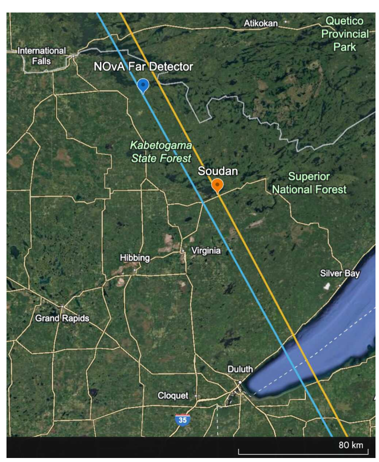

Figure 2.6: Map of northern Minnesota. The orange line depicts the center of the NuMI beampath, which passes through the MINOS far detector hall in the Soudan mine. The NOvA Far Detector site is located 11.8 km away from the beam center along the blue line 14.6 mrad separated from the NuMI beam. The NOvA Far Detector is sited by the Ash River, just south of Voyageurs National Park and the Canadian border. Maps data: Google, ©2021, Landsat/Copernicus, NOAA.

Using beam simulations to get the actual energy distribution of daughter neutrinos, the unoscillated neutrino spectrum NOvA sees from NuMI is sharply peaked at around 2 GeV (Figure 2.7).

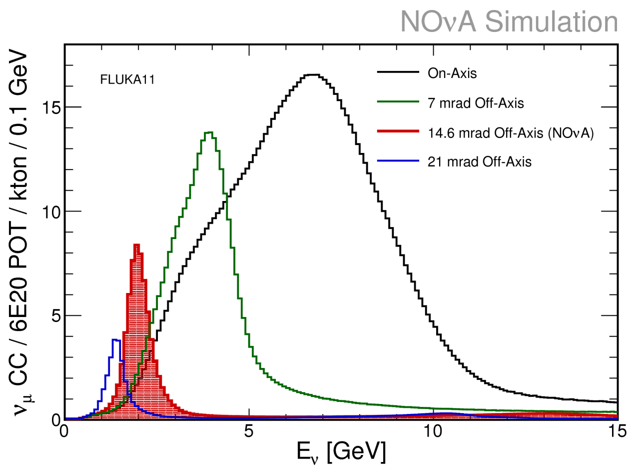

Figure 2.7: Simulated neutrino spectra at the NOvA far detector on-axis and at some off-axis angles relative to the NuMI beamline, assuming no oscillations. Note that larger off-axis angles result in a narrower spectrum at the expense of total flux [60].

2.3 NOvA Detectors and Data Acquisition

The NOvA detectors are functionally identical, segmented tracking calorimeters. The Far Detector (FD) is 14 kton and consists of 896 planes of liquid scintillator-containing cells. The Near Detector (ND) is 290 ton, and consists of 192 planes. In addition, the ND includes a muon catcher with an additional 22 planes interspersed with 4 inch thick iron plates. These extra planes only cover $\frac{2}{3}$ the height of the detector and help contain high energy muons for better energy estimation [61]. The ND is located at a baseline of ~1 km relative to the NuMI source, and the FD is at a baseline of 810 km.

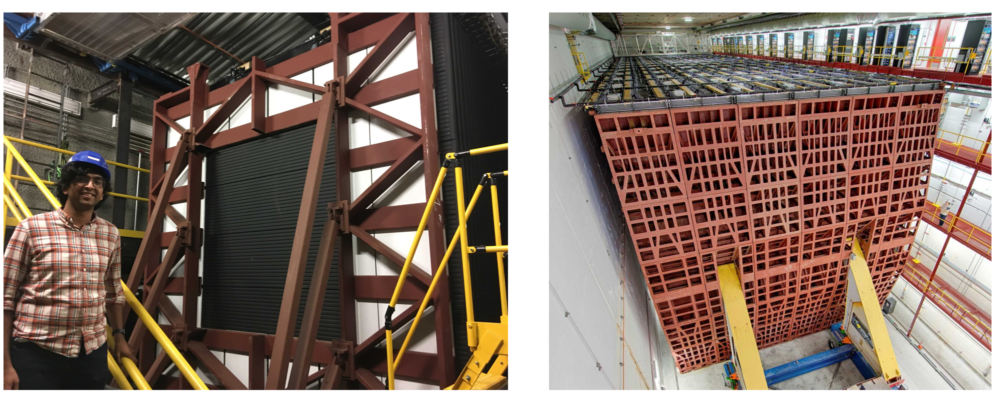

Figure 2.8: The NOvA Near Detector (left, author for scale), and Far Detector (right).

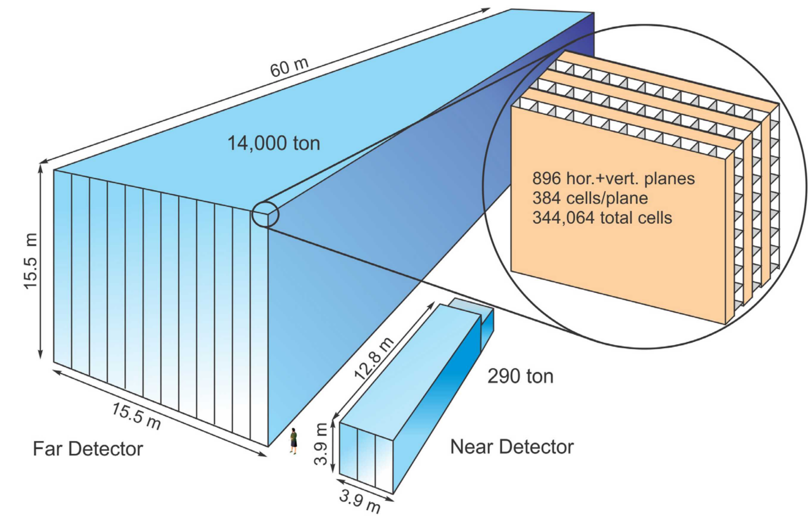

Figure 2.9: Schematic depicting the NOvA near and far detectors. Both detectors are made of alternating planes of horizontal and vertically oriented PVC extrusions [61].

The basic data-gathering unit of a NOvA detector is a cell. The detector is constructed of alternating planes of extruded PVC cells that run either the width or height of the detector.

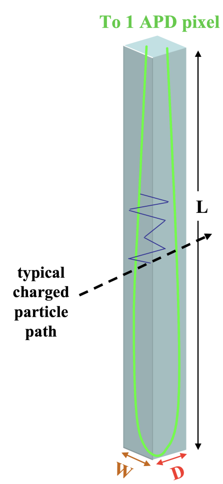

Figure 2.10: Depiction of the NOvA cell. Each cell is 3.8 cm wide and 5.9 cm deep along the beam direction. A loop of wavelength-shifting fiber twice the length of the cell is inserted to collect any scintillation light and send it to the APDs [62, 63].

Each cell contains a loop of fiber and is filled with a mineral oil-based liquid scintillator. This fiber is twice the length of the cell, and is responsible for collecting light and transporting it to be read out. Charged particles passing through the cell will excite pseudocumene, the primary scintillant. Pseudocumene primarily emits in the ultraviolet, but our readout electronics require longer wavelenghts of light. As a result, the scintillation light is shifted by a chain of wavelength shifters: the initial pseudocumene scintillation excites the wavelength shifter PPO, which emits and excites the wavelength shifter bis-MSB, which emits and excites wavelength shifters embedded in the fiber. The final green photons travel the length of the fiber to the detectors' readout electronics [63].

When a charged particle deposits enough energy to pass a threshold, we term this a "cell hit". Cells run either the full height or width of the detector, so each cell hit only contains two spatial dimensions: vertical cells determine xz, and horizontal cells determine yz (where z is the beam direction, x is the horizontal axis, and y is the vertical axis). Because the planes of cells alternate vertical and horizontal, each time a charged particle leaves a vertical cell it will enter a horizontal cell and vice versa. Thus, by combining adjacent events, NOvA can reconstruct full three-dimensional tracks (see Figure 2.11).

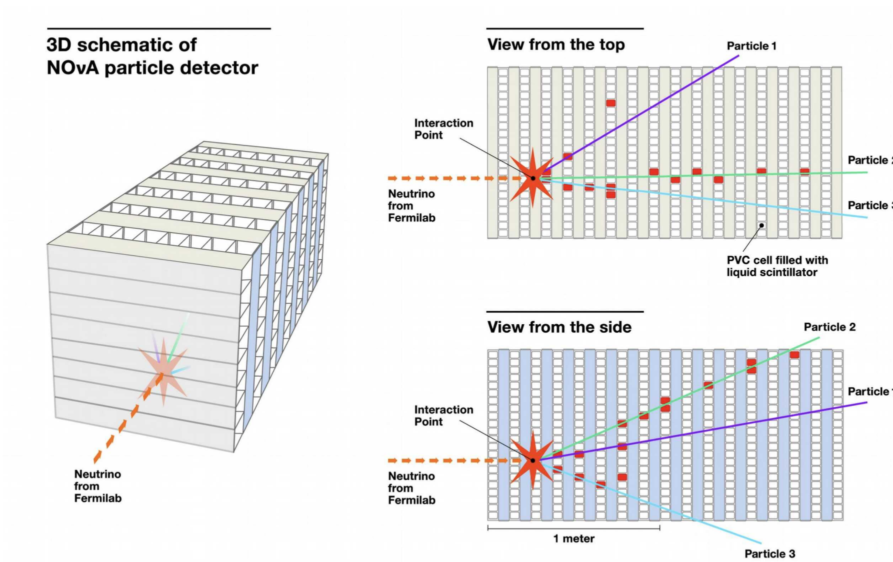

Figure 2.11: Illustration of particles passing through the NOvA detector. Because each plane alternately contains horizontal and vertical cells, each hit only provides two spatial coordinates, along with timing information. However, this is sufficient to reconstruct full three-dimensional tracks. Image courtesy of Fermilab.

Light from 32 cells is gathered in wavelength shifting fibers and sent to an avalanche photodiode (APD). These solid state components are capable of a quantum efficiency of 85% in the 500-550 nm range output by the wavelength shifting fibers. The output from the APDs above threshold is digitized via a Front End Board (FEB).

Because the APD noise reduces with temperature, each APD is outfitted with a thermoelectric cooler (TEC) capable of bringing their teperature down to −15°C. To prevent

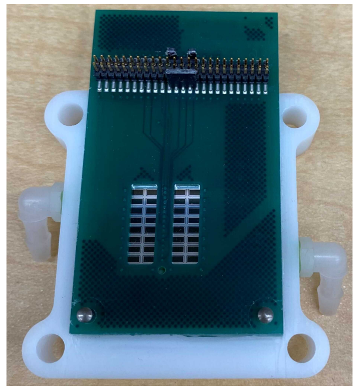

Figure 2.12: A NOvA APD. Each of the 32 dark rectangles is a single pixel, designed to be optically mated to both ends of a single loop of wavelength-shifting fiber coming from a cell. Hoses delivering dry gas are connected to the plastic connectors on either side of the APD, protecting it from condensation when cooled.

condensation on the APD, it is constantly bathed in dry gas. The TECs themselves are more efficient when cooled, so they are actively cooled with 4° C chilled water. The TECs are controlled by a daughterboard attached to the FEB.

Each FEB recieves power at three different voltages from Power Distribution Units (PDUs) arranged around the detector. The high-voltage supply is used by the APDs, a lower-voltage supply powers the TECs, and the lowest-voltage supply powers the rest of the FEB electronics [61].

Data from 64 FEBs is sent to a Data Concentrator Module (DCM). The DCM is responsible for collating incoming data from FEBs and sends packets to a set of buffer nodes nearby.

For accurate results from the NOvA detectors, all electronics must have excellent timing synchronization. To do this, NOvA developed custom electronics called Timing Distribution Units (TDU) to synchronize all parts of the detector to GPS time. The NOvA timing system builds a tree with TDUs as root nodes, FEBs as leaves, and DCMs in the middle. Each level of the tree carefully synchronizes its timing information with the level below, before repeating with the next level until the entire detector is synchronized. This system is accurate enough for all electronics in the Far Detector to agree on the current time within 15.6 ns [64].

NOvA streams all data from the DCMs into a circular buffer that can hold a maximum of 20 s worth of data. As a result, no event triggering occurs on the FEBs or DCMs, and instead occurs on the live data stored in the circular buffer nodes. Each NuMI spill is treated as a trigger, and all data from the spill window is saved.

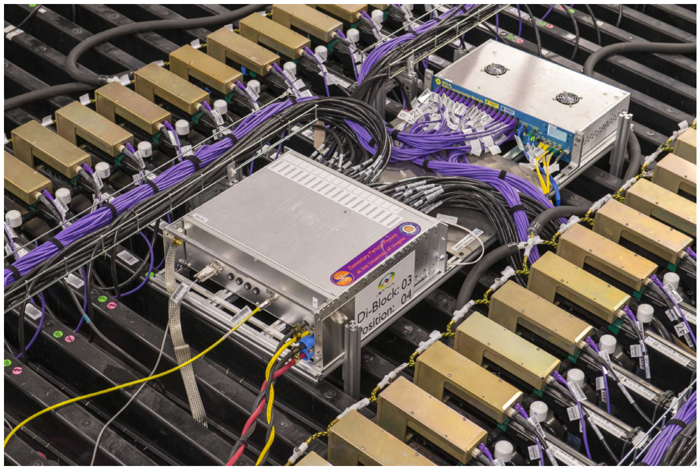

Figure 2.13: Example PDU (foreground) and DCM (background) on the Far Detector. The gold boxes surrounding are FEBs, and are responsible for digitizing APD output from 32 cells [61].

Given the 1.3 s NuMI rep rate and the 10 μs spill time, if NOvA were only to store data associated with the NuMI trigger, it would throw away well over 99% of NOvA's live data readout [65]. Instead, NOvA makes use of its buffered readout to use Data-Driven Triggers (DDTs) to construct various events, beyond the NuMI trigger. This enables NOvA to use live detector readout to improve detector calibration, conduct Beyond-the-Standard Model (BSM) searches, and use external triggers like those from SNEWS and LIGO [66, 67].

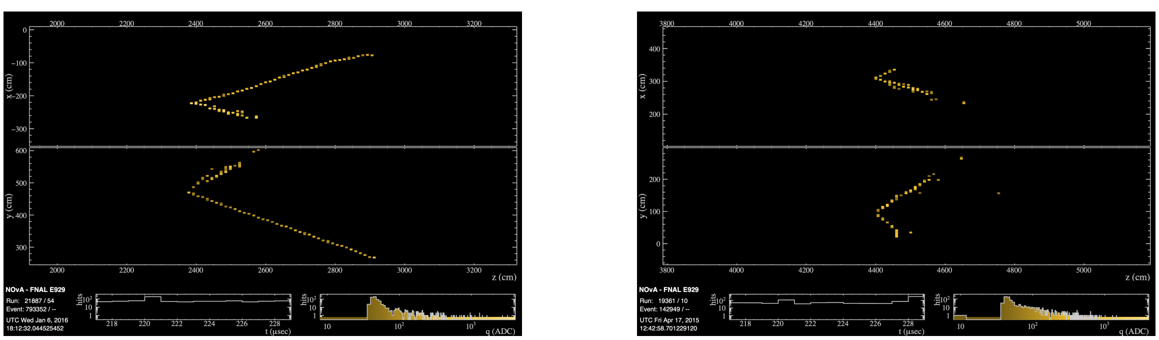

(a) ν_μ CC event

(b) ν_e CC event

Figure 2.14: Event Displays from the NOvA 2020 neutrino event samples. ν_μ events display a long track corresponding to the muon, whereas ν_e events show an electromagnetic shower corresponding to an electron [68, 69].

2.4 Event Selection and Reconstruction

The oscillation analysis requires selecting and reconstructing beam neutrinos. The data associated with a NuMI trigger gets clustered into "slices", where each slice corresponds to regions with a high density of cell hits in time and space, taking care that cell hits that happen at the same time and z coordinates but come from different physics events should not be merged into the same slice [70].

First, slices pass basic data quality cuts. Any slice caused by a beam neutrino needs to be temporally associated with the 10 μs NuMI spill window. Additionally, slices need to be physically contained within the fiducial volume, as if an interaction occurs too close to the edge of the detector, visible energy can escape the detector, preventing accurate energy estimation.

NOvA uses a convolutional neural network (CNN_{evt}) to score events on whether they look most like $\overset{(-)}{\nu}_\mu$ CC, $\overset{(-)}{\nu}_e$ CC, NC, or cosmogenic background events. Figure 2.15 shows a t-SNE (t-distributed stochastic neighbor embedding) plot visualizing the clusters developing among feature vectors at the second-to-last layer of the neural network.

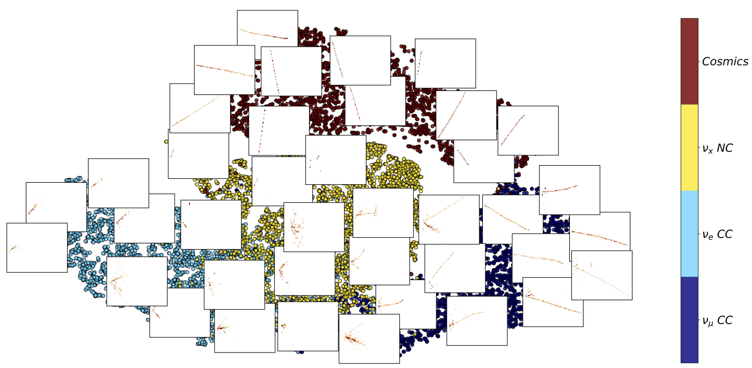

Figure 2.15: t-SNE visualization of CNN_{evt} feature vectors at the penultimate layer. Each point is colored by true interaction type. Event displays from the training sample are overlaid to show the differences in topology between events [71]. Note that t-SNE is purely for visualizing high-dimensional data. Distances between points do not have physical meaning.

At the Far Detector, cosmic rates are about 130 kHz since the detector is on the surface. This means that there should be an average of ~1.3 cosmic ray particles per 10 μs NuMI spill window. As such, NOvA uses energy depositions outside of the NuMI spill window to train both a Boosted Decision Tree (BDT) as well as another convolutional neural network to identify cosmics. These models are able to bring cosmogenic backgrounds down to under 5% of the selected sample.

Neutrino energy estimation is handled separately for ν_μ and ν_e CC events. In both cases, the general principle is to measure the energy of the charged lepton along with the energy of the hadronic system. The sum of these quantities provides an estimate for the incident neutrino energy.

2.5 Simulation and Calibration

Monte Carlo Simulation (MC) begins with modeling of the NuMI beamline using the GEANT4 package to produce an estimated flux of neutrinos from the NuMI beam [72]. NOvA then uses PPFX (Package for Predicting the FluX), which leverages data from MINERvA to reweight the naïve GEANT4 flux prediction to better match observed NuMI data [40].

NOvA feeds these flux predictions into a modified version of the GENIE event generator to create simulated events in the detectors [73]. GENIE handles neutrino scattering, and produces a set of final-state daughter particles of the interaction. These final-state particles are then sent through GEANT4 to propagate through the detector geometry, depositing energy. Cosmic rays are simulated via the CRY library, and like GENIE, the daughter particles of cosmic ray interactions are propagated through the detector via GEANT4 [74].

Next, NOvA uses a custom simulation stack to turn energy depositions into simulated APD signals. A ray tracer model estimates light collection in the wavelength-shifting fibers. Finally, a parametrized model of the APD and readout electronics turns simulated photoelectrons into the digitized readout for the given event [75].

If simulation is the process of turning raw energy depositions in the detector into electronic output, calibration is the process of turning electronic outputs into estimated energy depositions. NOvA calibrates each detector in two stages: relative calibration to ensure detector performance is uniform across the detector, and absolute calibration to set the detector's energy scale [76].

Relative calibration itself is a multi-step process, accounting for a number of effects that can cause output photoelectrons to be inconsitent spatially across the detector. These include effects due to attenuation between the signal source and the readout electronics, and "shadowing" effects where (especially in the Far Detector) energy deposition is higher near the top of the detector compared to the lower parts which are shielded by the overlying detector matter.

The absolute calibration is determined via minimum-ionizing particles as they stop in the detector [77]. NOvA has observed a 0.3% drop in observed light per year, so this absolute calibration happens multiple times over shorter time periods.

2.6 Extrapolation

NOvA is able to leverage the identical technology underpinning each detector. Any inadequacies of the simulation for the Near Detector would also affect the Far Detector. For example, if NOvA simulation consistently underpredicts the expected neutrino counts at a given kinematic bin in the ND, we can expect the simulation would also undercount that bin in the FD by a similar degree. This decision led NOvA (and before it, MINOS) to use a technique called extrapolation, where ND data is not fit to, but is instead used to build a data-driven FD prediction using simulation (Figure 2.16).

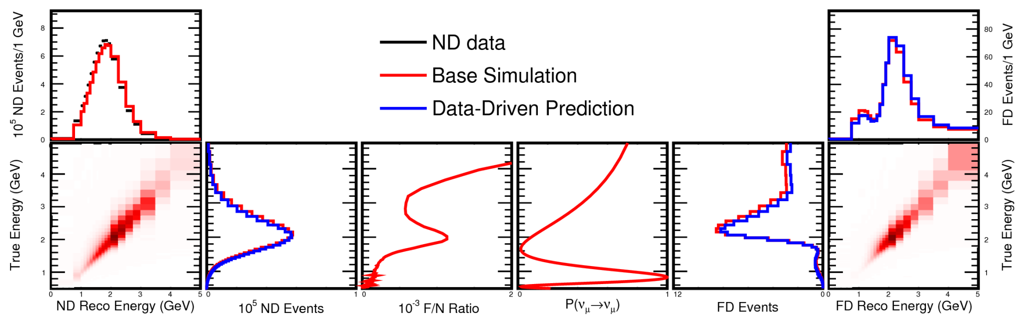

Figure 2.16: NOvA's extrapolation process to use Near Detector Data to form a data-driven prediction for the Far Detector. Because the Near and Far Detectors are functionally identical, we expect very strong correlations in systematic errors between the two detectors. Extrapolation takes advantage of this fact by using Data-MC discrepancies to build a more accurate prediction of Far Detector spectra [78].

To get from ND Data to an FD prediction, we construct a series of weights. The first weights convert ND Data into estimated true spectra. These spectra are then converted into true FD spectra, accounting for beam geometry, attenuation, and oscillation. Finally, the FD true spectra are used to create a predicted reconstructed spectrum. This data-driven prediction is then fit to FD data in the final oscillation analysis. For more on the extrapolation procedure, see [77].

The benefit of extrapolation is that it greatly reduces the systematic error on NOvA's oscillation measurements. Figure 2.17 shows the effect of extrapolation on the systematics budget. Because the extrapolation process expects some homogeneity between the detectors, those systematic errors which do depend on the detector are slightly increased. However, this is more than offset by the reduction in uncertainty for those which extrapolation greatly helps.

Rather than exclusively binning in E_ν, NOvA also separates the extrapolation into bins of transverse lepton momentum p_T. Because the FD is so much bigger than the ND, it

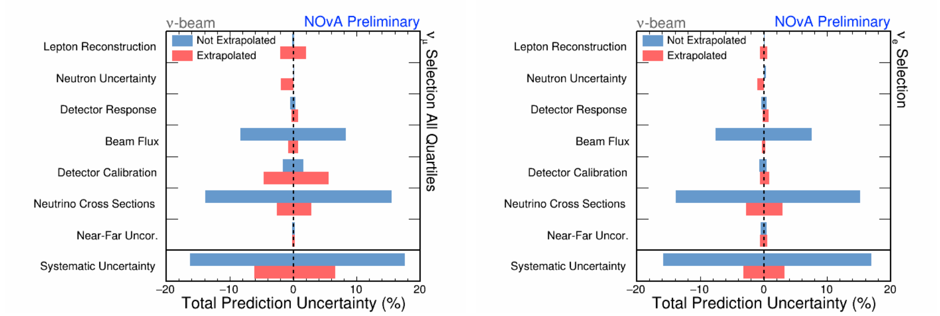

Figure 2.17: Systematic Errors with and without extrapolation for ν_μ (left) and ν_e (right) reconstruction. Large systematic errors which are correlated between the near and far detector are traded for smaller systematic errors where the detectors differ, as in calibration and in lepton reconstruction [79].

is better able to contain leptons with larger transverse momenta (Figure 2.18). Thus, by separating events by p_T, extrapolated predictions are less sensitive to modeling of p_T, at the expense of a modest increase in lepton reconstruction uncertainty.

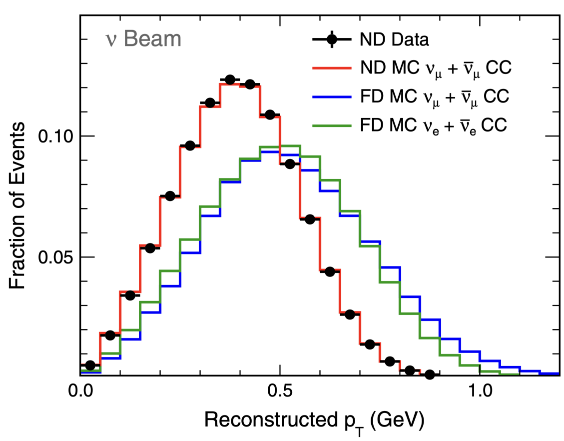

Figure 2.18: The larger size of the FD means that leptons can have larger transverse momenta while still being contained.

2.7 Binning

As the strength of oscillations depends on the neutrino energy E_ν, NOvA's fits are done over binnings that are primarily in reconstructed neutrino energy. NOvA considers several factors when optimizing the binning, including the energy resolution and memory demands on the final fits [80].

ν_μ Analysis Binning

The ν_μ energy binning is fine between 1 and 2 GeV, where we expect the oscillation signal, and is coarser elsewhere. The E_ν bin widths are chosen to make the contents of each bin roughly equal in an unoscillated spectrum [81].

NOvA's ν_μ energy resolution depends on how much of the incoming neutrino's energy is sent into the daughter muon. Since the muon energy estimation is much better than the hadronic energy estimation, we divide each bin of the FD E_ν spectrum into four quantiles based on the hadronic energy fraction $E_{\text{frac}} = \frac{E_{\text{had}}}{E_\nu} = \frac{E_\nu - E\mu}{E_\nu}$. Figure 2.19 shows the 2-dimensional ν_μ spectrum, with quantiles demarcated. This also helps, since cosmic and NC backgrounds are reconstructed to have high hadronic energy fractions, which increases the signal purity of the lower E_{frac} quantiles.

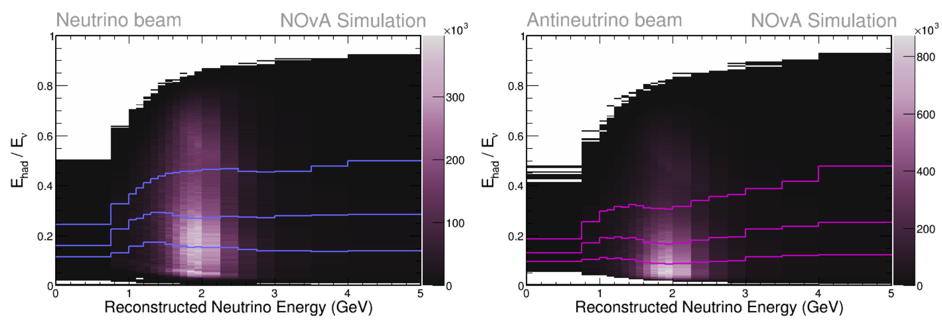

Figure 2.19: A 2D ND Spectrum depicting Hadronic Energy Fraction E_{frac} vs Neutrino energy for FHC (left) and RHC (right). The E_{frac} quantiles are drawn over the spectrum. [82].

Figure 2.20a shows the spectrum at the ND using these bins, split into the four quantiles.

ν_e Analysis Binning

Due to the low statistics of ν_e events, we cannot bin as finely as in ν_μ events. NOvA splits the ν_e sample into three categories to optimize the signal-to-background ratio. Events which pass all prior cuts and have a high ν_e score from CNN_{evt} are binned separately from those with lower scores. These two samples correspond to the "core" ν_e samples. However, our cosmic rejection and containment cuts may discard real ν_e events. To improve efficiency, we construct a "peripheral" bin from events which fail either cut but have very high CNN_{evt} scores and/or high BDT scores for cosmic rejection.

The "core" samples are binned in 6 E_ν equal-width bins each. The "peripheral" sample is treated as a separate, energy-integrated bin, since peripheral events have no ND equivalent. As a result, there is no ND spectrum that can be extrapolated to construct a FD peripheral spectrum.

Spectra at the ND from the 2020 NOvA analysis using this binning are depicted in Figure 2.20. Note that the peripheral sample is only in the FD.

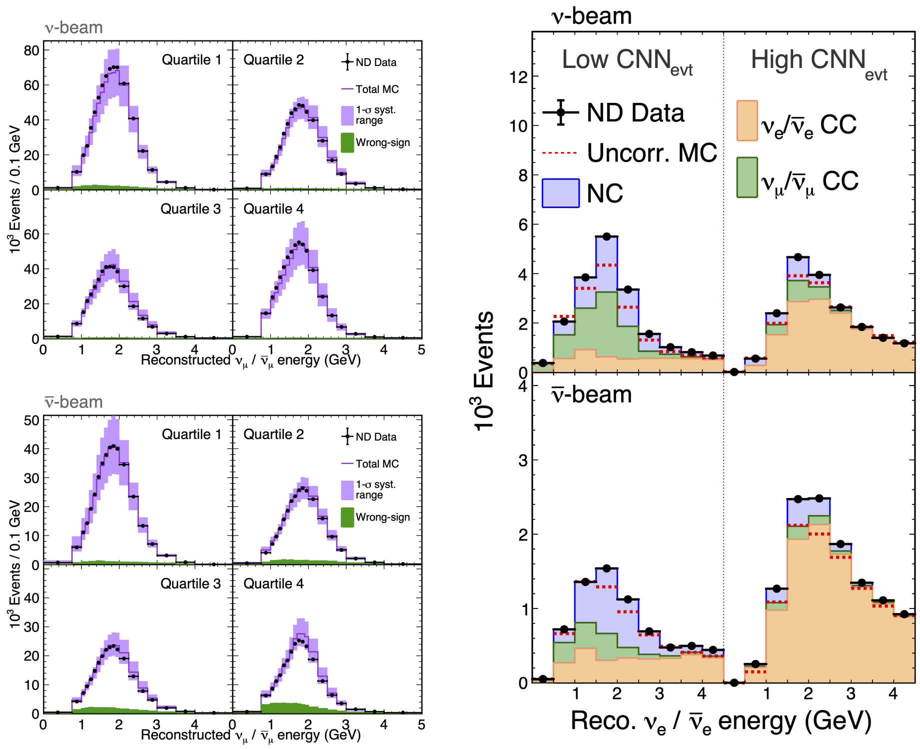

(a) ν_μ Spectrum (b) ν_e Spectrum

Figure 2.20: Spectra at the near detector, showing the final analysis binnings for ν_μ (left) and ν_e (right). Note that the ν_e ND spectrum is essentially a measurement of the beam ν_e background, with the high CNN$_{\text{evt}}$ sample having a higher purity of beam ν_e events. Also note that the peripheral sample does not exist in the ND [77].

2.8 Systematic Uncertainties

Simulation will never match reality. As a result, NOvA's simulation will introduce systematic biases that need to be accounted for to produce accurate oscillation results. For each source of systematic uncertainty, we create new output spectra by adjusting the associated model in the base MC.

Some sources of systematic uncertainty, like those corresponding to light levels in the detector, require rerunning the base MC with the model changed. However, most other systematics can be handled much less intensively by reweighting the base simulation.

Each source of systematic uncertainty is associated with one or more parameters, whose value can be interpreted as the deviation from the base simulation. Arbitrary shifts to a systematic parameter are found by interpolating between pre-calculated spectra.

Figure 2.21 shows a summary of systematic errors in NOvA after extrapolation.

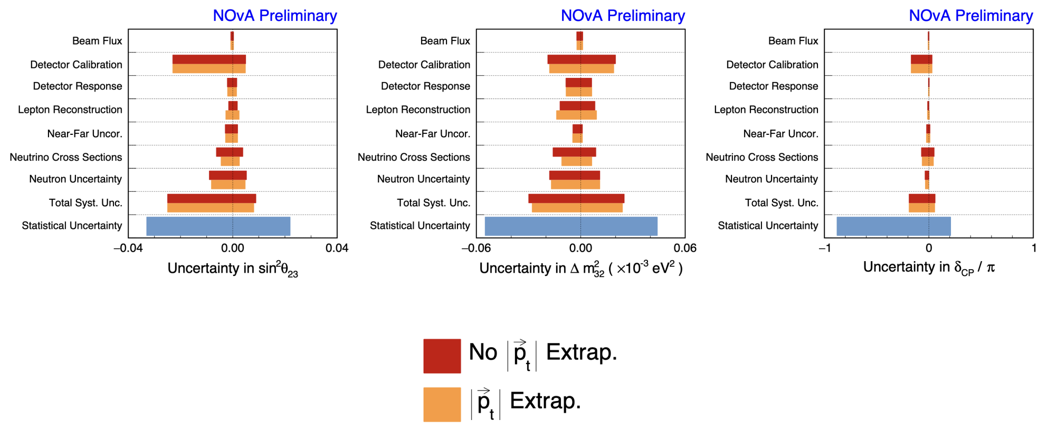

Figure 2.21: Summary of systematic uncertainties in NOvA, using the full extrapolation procedure. NOvA uses 67 individual systematic parameters, which are grouped into categories and added in quadrature for display above [83]. Note that the p_T extrapolation, depicted in orange, is used in oscillation analyses.

The largest individual systematic uncertainty is the 5% energy scale calibration uncertainty. Other detector-related systematics include those related to NOvA's modeling of light and detector readout (see Section 2.5).

While the extrapolation procedure does greatly reduce cross-section and detector systematics which are correlated between the detectors, it does introduce error associated with differences between the detectors. This can be seen in Figure 2.17, where the systematic groups listed as "Near-Far Uncor." and "Lepton Reconstruction" show an increase in uncertainty after extrapolation.

Neutrons also lead to an important uncertainty. Because they are neutral, energy associated with a daughter neutron from a neutrino interaction may be displaced from the interaction vertex. Difficulties in modeling neutron scattering in the detector lead to an associated systematic uncertainty.

The NOvA Test Beam program was designed to help improve our detector response models. For more on this, see Appendix A.

2.9 Neutrino Interaction Uncertainties

Finally, neutrino interaction modeling forms the largest remaining source of uncertainty in the experiment. NOvA considers four main ways that neutrinos can undergo charged-current (CC) interactions with hadronic matter, from lowest- to highest-energy:

- **Quasielastic Scattering (QE)**: the incident neutrino scatters mostly elastically via $W^{\pm}$ exchange. Inelasticity is due to the difference in masses between the parent and daughter nucleon and the charged lepton.
- **Two-Particle-Two-Hole Scattering (2p2h)**: the incident neutrino scatters off a pair of nucleons instead of a single nucleon via the Meson Exchange Current (MEC) in an otherwise QE event. This ejects both nucleons from the nucleus. Such scattering evnts are important to model for NOvA, especially as GENIE does not include 2p2h events by default.
- **Resonant Scattering (RES)**: the incident neutrino causes a struck nucleon to enter an excited state, leading to additional final state daughter mesons.
- **Deep Inelastic Scattering (DIS)**: the incident neutrino scatters inelastically off of the partons in a nucleon, creating a broad spectrum of daughter hadrons.

In addition to these processes, very low-energy neutrinos can scatter coherently off of the entire nucleus. Such coherent scattering (COH) events are rare and can be ignored in this discussion.

Neutral-current (NC) neutrino interactions do not produce a charged lepton to detect. This makes it difficult to identify the event's energy, or the type of neutrino-nucleus interaction. NOvA's measurements are therefore sensitive to NC modeling uncertainty.

NOvA adjusts the GENIE 2.12.2 model for each of the above interaction channels. NOvA's systematic uncertainties related to cross-section models also require some modification beyond the GENIE default. This introduces additional uncertainties related to changes to the QE nuclear model, suppression of RES events at low Q^2, and NOvA's custom 2p2h model, which is based on the Empirical MEC model packaged with GENIE [84]. The result of NOvA's GENIE tune is shown in Figure 2.22.

The net result of these uncertainties is displayed in Figure 2.23. Chapter 5 will discuss a method for replacing these systematics in a parametrization-agnostic way.

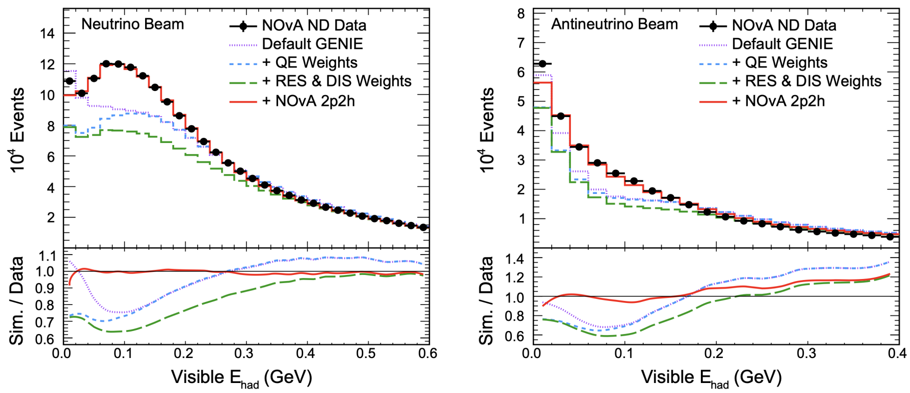

Figure 2.22: NOvA applies weights to GENIE to correct for the inaccuracies for both FHC (left) and RHC (right) beams [84].

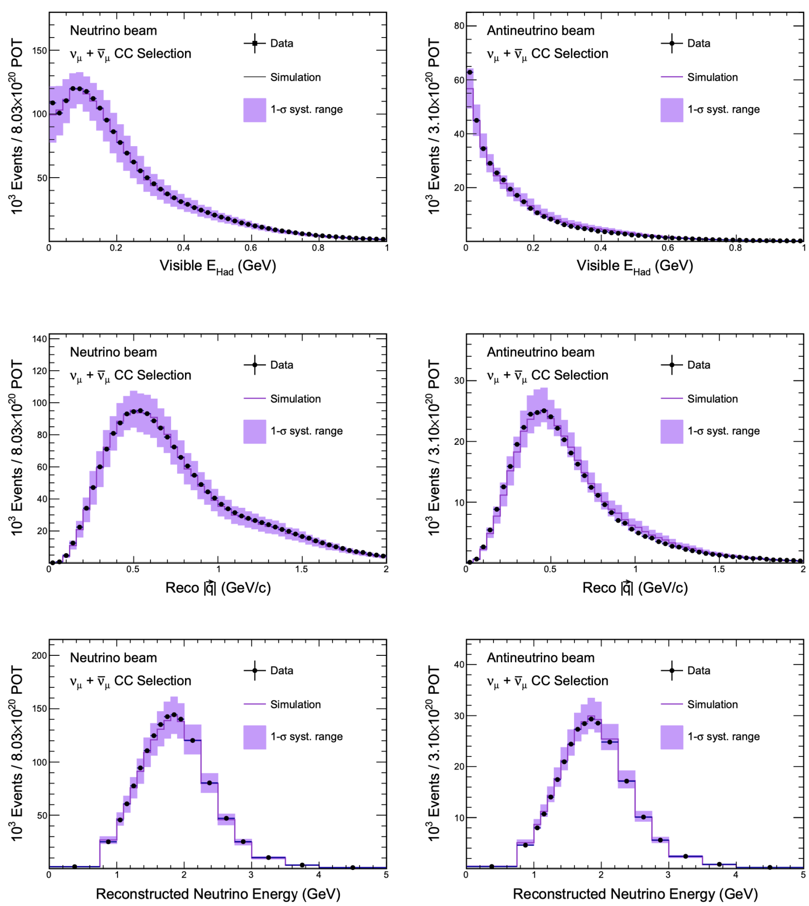

Figure 2.23: The overall impact of NOvA's cross section modeling uncertainties. The shaded band shows $\pm 1\sigma$ relative to the central value of simulation, with each GENIE parameter's contribution added in quadrature. In multiple kinematic variables, and for both FHC (left column) and RHC (right column) beams, NOvA's ND data falls within the 1σ band [84].

2.10 Fitting

NOvA runs both frequentist and Bayesian fits to FD data. Both fits find the χ^2 surface describing the difference between the data-driven predictions and FD data by adjusting the free parameters: the core physics parameters governing neutrino oscillation, and the 67 systematic parameters. In this thesis, we will be solely reporting results from the Bayesian fit.

There are six main oscillation parameters (see Section 1.2). We use external values for the solar parameters Δm^2_{21} and θ_{12}, which leaves Δm^2_{32}, θ_{23}, δ_{CP}, and θ_{13}. Of these four, reactor neutrino experiments have made much stronger measurements on θ_{13} than NOvA can. As a result, we run fits of all four free oscillation parameters, as well as with an external constraint on θ_{13} informed by the latest reactor neutrino measurements.

We punish the fitter from moving the systematic parameters too far from their central values by introducing a penalty term into the overall likelihood function.

NOvA runs two separate Bayesian fitters, both variants of Markov Chain Monte Carlo (MCMC): Aria and Stan. Both fitters produce nearly identical results.

The Aria fitter was developed by NOvA to implement the Metropolis-Rosenbluth-Rosenbluth-Teller-Teller ($\mathrm{MR^2T^2}$) MCMC sampling procedure. For a more detailed explanation of MCMC, and specifically the Aria Fitter used by NOvA for the NOvA-T2K joint fit, see Appendix B.

The other MCMC fitter leverages the Stan library for Hamiltonian MCMC. Hamiltonian MCMC uses a physics-based metaphor to sample the posterior distribution. The posterior distrubtion is treated like a physical surface, and sampling is done by treating a proposed step direction as a particle experiencing gravity while sliding along this surface. The dynamics of this particle are computed by numerically integrating Hamiton's equations. While Hamiltonian MCMC requires much more computation per step, each step is much more uncorrelated than in $\mathrm{MR^2T^2}$ MCMC, requiring fewer steps to properly sample the posterior distribution.

Unlike frequentist fits, Bayesian fits require a prior as input to create posterior probability distributions. In general, we choose flat priors for physics parameters. However, there can be some subtleties with physics quantities like δ_{CP}. For example, flat prior in δ_{CP} would not be flat in $\sin\delta_{CP}$, and vice versa. As a result, for quantities like this, we try both priors to ensure our final credible intervals are not sensitive to the choice of prior. This especially affects measurements of the Jarlskog invariant, as it is nonlinearly dependent on multiple oscillation parameters.

2.11 Recent NOvA 3-Flavor Results

As of the 2020 NOvA analysis the FD saw 211 ν_μ CC and 82 ν_e CC events in an FHC beam; and 105 $\overline{\nu}_\mu$ CC and 33 $\overline{\nu}_e$ CC events in an RHC beam. The ν_e and $\overline{\nu}_e$ event counts are well above the estimated background rates over the same time of $26.8^{+1.6}_{-1.7}$ ν_e events and $14.0^{+0.9}_{-1.0}$ $\overline{\nu}_e$ events, showing clear evidence of oscillation.

Observed FD spectra, along with the frequentist best-fit FD prediction are depicted in Figure 2.24.

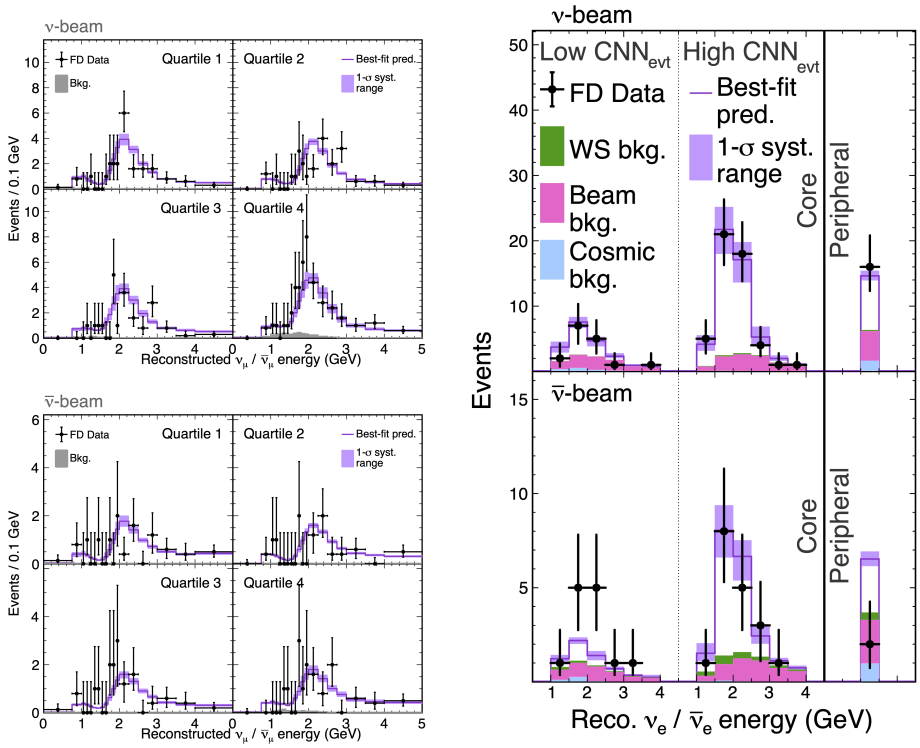

(a) ν_μ Spectrum (b) ν_e Spectrum

Figure 2.24: Spectra at the Far Detector. The best-fit prediction from the frequentist analysis is depicted in purple, with the shaded band showing 1σ in systematic uncertainty. FD Data is overlaid in black [77].

We can then run physics fits on this data. Figure 2.25 show the results of our Bayesian fit, using the reactor constraint. The credible intervals were created with Stan, although Aria contours look virtually identical [85].

Note that "1σ" credible intervals and contours are defined to contain the same amount of posterior as a $\pm 1\sigma$ interval does for a Gaussian distribution (namely ~68%).

When we run fits, we allow Δm^2_{32} to be either positive or negative, which corresponds to the Normal and Inverted Orderings (NO and IO) respectively. The total 1σ interval would therefore be the union of the NO and IO 1σ intervals depicted.

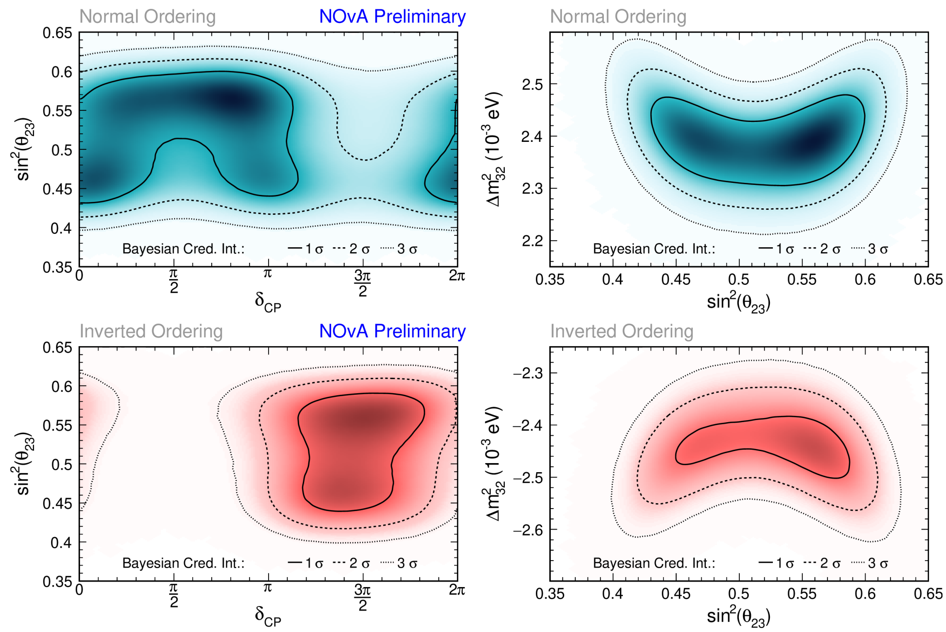

Figure 2.25: Contours for δ_{CP} vs $\sin^2\theta_{23}$ (left) and $\sin^2\theta_{23}$ vs Δm^2_{32} (right) [85].

We can evaluate physics conclusions by seeing how the posterior probability is distributed. The final fit shows a that 68% of the posterior probability lies in the Normal Ordering, indicating a slight preference. We also see a slight preference for the upper octant of θ_{23}, which holds 63% of the posterior probability. Another way of reporting this preference is via a Bayes factor: the ratio of posterior probabilities for two competing hypotheses. If the Bayes factor is greater than 1, then the posterior shows a preference for A over B. Thus, the Bayes factors for the NO and upper octant are 2.08 and 1.67 respectively.

Conclusions on CP violation can be found by looking at the Jarlskog invariant J:

$$J = \cos(\theta_{12})\cos^2(\theta_{13})\cos(\theta_{23})\sin(\theta_{12})\sin(\theta_{13})\sin(\theta_{23})\sin(\delta_{CP})$$

As shown in Figure 2.26, CP-conservation is accepted within a 3σ credible interval, regardless of choice of prior. We can again look at Bayes factors, although here we are comparing the two hypotheses of $J = 0$ and $J \neq 0$. It is not possible to directly calculate a nonzero posterior probability for the point hypothesis $J = 0$, so NOvA uses the Savage-Dickey method to calculate Bayes factors [86]. Using this approach, we see that CP violation is slightly preferred with a Bayes factors of 1.0-1.2 (depending on the choice of δ_{CP} prior) in the NO and 3.4-3.8 in the IO. This indicates that there is a slightly stronger preference for CP violation in the IO compared to in the NO.

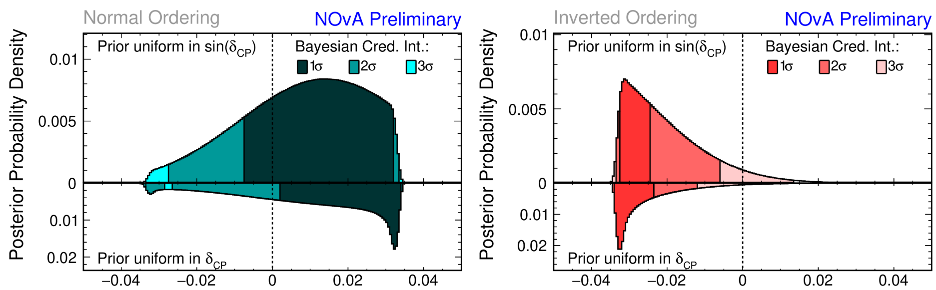

Figure 2.26: The Jarlskog invariant J with two different choices for prior. One prior is flat in δ_{CP}, whereas the other is flat in $\sin \delta_{CP}$. In both cases, the CP-conserving value of $J = 0$ is accepted within the 3σ credible interval [85].

As discussed in the previous section, reactor experiments provide tight constraints on θ_{13}. Figure 2.27 shows the posterior distribution in θ_{13} and θ_{23}, with the reactor measurement overlaid. Especially in the NO, we see that the slight preference for the upper octant in θ_{23} is due to the reactor constraint breaking the degeneracy we see in octant measurements.

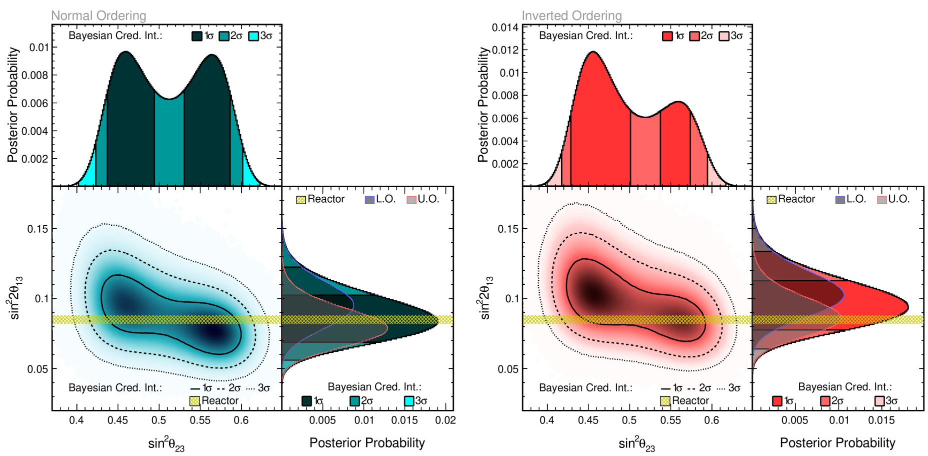

Figure 2.27: $\sin^2\theta_{23}$ and $\sin^2\theta_{13}$ posterior distributions, with the state-of-the-art reactor neutrino measurements of θ_{13} overlaid. Imposing the reactor constraint can break the degeneracy in θ_{23} octant, and causes NOvA's measurements to slightly prefer the upper octant [85].

Chapter 3

T2K AND THE JOINT FIT

The Tokai-to-Kamioka Experiment (T2K) is a similar long-baseline accelerator neutrino oscillation experiment to NOvA. NOvA and T2K have complementary abilities to measure neutrino oscillations, leading to efforts to combine measurements from both experiments. Ideally, a joint fit would provide tighter constraints on oscillation parameters than either experiment could produce individually. This chapter will discuss T2K and the motivation for the joint fit that is the result of this thesis.

3.1 T2K

The T2K experiment is a long-baseline oscillation experiment in Japan. Neutrinos are produced at the J-PARC accelerator facility in Tokai, Ibaraki Prefecture. The beam then travels 295 km to the Super-Kamiokande detector nestled deep in the Japanese Alps. Similar to NuMI, the J-PARC neutrino beam is created by accelerating protons into a rod-shaped graphite target. Daughter mesons are focused via magnetic horns into a beam, and decay to produce neutrinos. By flipping the polarity of the magnetic horns, the beam can be made primarily of either ν_μ or $\overline{\nu}_\mu$.

Figure 3.1: The T2K neutrino beam runs 295 from the J-PARC accelerator facility in Tokai, Ibaraki Prefecture, to the Super-Kamiokande detector in Gifu Prefecture. Maps Data: Google ©2021, Data SIO, NOAA, U.S. Navy, NGA, GEBCO, Landsat/Copernicus, Data Japan Hydrographic Association, Data LDEO-Columbia, NSF, NOAA, TMap Mobility.

The J-PARC neutrino beam is oriented such that Super-Kamiokande is off the beam axis, like NuMI. However, it is at a sharper angle of 2.5° off-axis (compared to NOvA's 14.6 mrad off-axis angle). This results in the unoscillated spectrum expected at SK to be peaked at a lower energy than NuMI, around 600 MeV (see Figure3.2). The lower energy and lower baseline mean that T2K is sensitive to Δm^2_{32} oscillations, like NOvA.

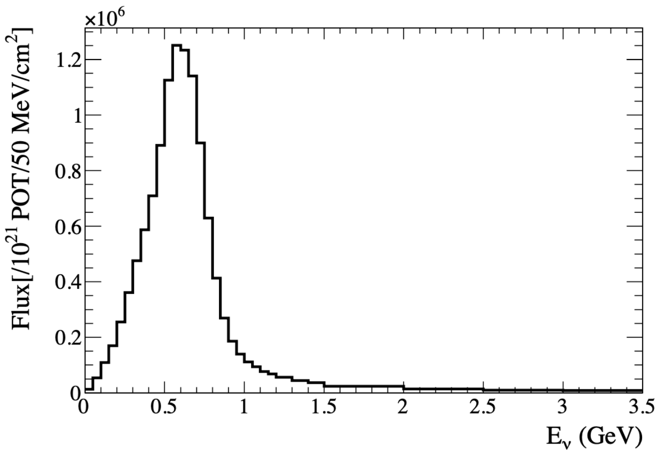

Figure 3.2: Simulated neutrino spectrum at SK due to the 2.5° off-axis T2K beamline, assuming no oscillations [87].

Located 280 meters downstream of the beam source is the ND280 complex of near detectors (Figure 3.3). The complex consists of an on-axis beam monitoring component and a magnetized off-axis neutrino detector.

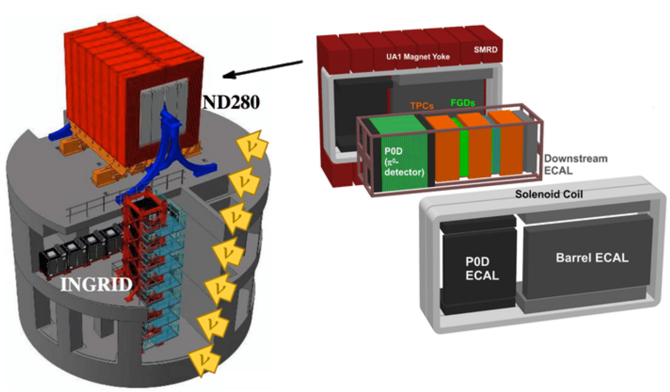

Figure 3.3: The ND280 detector complex. INGRID, the on-axis beam monitor is located in the lower floors, whereas the magnetized off-axis detector suite is located above [88].

The INGRID (Interactive Neutrino GRID) detector provides on-axis beam monitoring. It consists of several identical square modules arranged in a cross. Each module is made of alternating planes of iron and scintillating bars. INGRID's main purpose is as a measurement of beam direction, intensity and symmetry.

The off-axis detectors are responsible for constraining beam flux, energy, and flavor composition. These consist of the Pi-Zero Detector (P∅D), three Time Projection Chambers (TPCs), and two Fine Grained Detectors (FGDs), all surrounded by an Electromagnetic Calorimeter (ECAL). All these detectors in turn are located inside a magnet capable of creating a dipole field with a strength of 2 T. The P∅D consists of drainable bags of water, which serve as target, instrumented with strips of plastic scintillator to track charged particle motion. The TPCs use gaseous Argon as their detecting medium, and the cathode and anode planes are oriented such that the drift electric field is aligned with the external magnetic field. The FGDs are made of strips of plastic scintillator arranged in alternating planes (similar to NOvA's scintillator cells) which allow for 3D track reconstruction within the FGD volume. Finally the surrounding ECAL consists of layers of plastic scintillator and lead to contain and measure any energy that leaves the inner off-axis detectors.

Like its predecessor experiment K2K (see Section 1.3), T2K does not have a custom far detector, and instead leverages the 50 kton Super-Kamiokande (SK) water Cherenkov detector. When neutrinos undergo CC interactions in SK, the resultant charged lepton will create Cherenkov radiation, which emanates outwards in a cone in the lepton's direction of travel. When the Cherenkov cone intersects with the outer wall of PMTs, it will leave a ring-shaped signal. The size, location, and shape of the ring can be used to estimate the location of the neutrino interaction vertex, the lepton's flavor, and its energy. Electrons scatter frequently in the detector medium, and can create electromagnetic showers, whereas muons scatter less. This allows SK to disambiguate ν_μ and ν_e events as ν_μ events will have "sharp" rings whereas ν_e events will have "fuzzy" rings (Figure 3.4). These rings are also referred to as "muon-like" and "electron-like" respectively.

Final state particles from the hadronic system can cause additional Cherenkov rings if any daughter particles are past the energy threshold. As such, T2K separates samples by the number and type of rings seen. The five samples used in the 2020 analysis are 1Rμ (1 muon-like ring) in ν- and $\overline{\nu}$-mode beams, 1Re (1 electron-like ring) in ν- and $\overline{\nu}$-mode beams, and 1Re1de (1 electron-like ring associated with a ν_e CC event and another electron-like ring created from the final state hadronic system) in ν-mode beam.

Backgrounds to the oscillation analysis include beam ν_e's, wrong-sign neutrinos, and

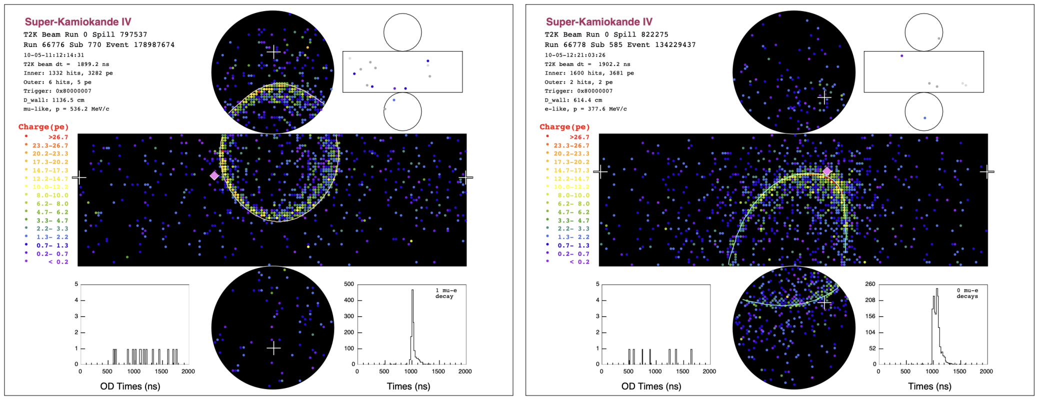

Figure 3.4: Event Displays for muon-like (left) and electron-like (right) events at SK. Reconstructed event vertex coordinates are depicted with white crosses, and the pink diamond indicates the incident neutrino beam direction [87].

NC events which produce a single pion. In NC π^0 events, if the two decay photons of the neutral pion are sufficiently aligned, the resultant electromagnetic shower can look like a single fuzzy Cherenkov ring. Meanwhile, NC $\pi^{\pm}$ events at low momentum can produce Cherenkov rings that look muon-like.

As of 2023, T2K has analyzed ten data-taking runs. Over this time, it has accumulated 318 1Rμ events in ν-mode beam, 138 1Rμ events in $\overline{\nu}$-mode beam, 94 1Re events in ν-mode beam, 16 1Re events in $\overline{\nu}$-mode beam, and 14 1Rede events in ν-mode beam.

3.2 Latest T2K Results

Unlike NOvA, which is able to use extrapolation to build data-driven FD predictions, T2K has to fit their simulation to both ND280 and SK data. T2K's analysis requires many more nuisance parameters in order to give the simulation enough degrees of freedom to properly fit data. In total, T2K ends up needing to run fits on well over 700 parameters, of which over 500 are normalization parameters for regions of phase space in ND280. This makes fits challenging. T2K accomplishes these fits in two different ways in parallel analyses: one frequentist and one Bayesian [89]. As with NOvA, we will only report the results of the Bayesian analysis.

The fits to ND280 and SK data can be done either sequentially or simultaneously. A separate ND280 fit can constrain uncertainties before a final fit. This final fit is done with an MCMC fitter called MaCh3, which uses a slightly modified version of the MR^{2}T^2 algorithm due to Hastings [90]. Because MCMC is capable of running fits with large numbers of nuisance parameters, MaCh3 is able to manage the simultaneous fit.

Figure 3.5 shows the spectra observed at SK. Good data-Monte Carlo agreement shows clear evidence of $\overset{(-)}{\nu}_\mu$ disappearance and $\overset{(-)}{\nu}_e$ appearance.

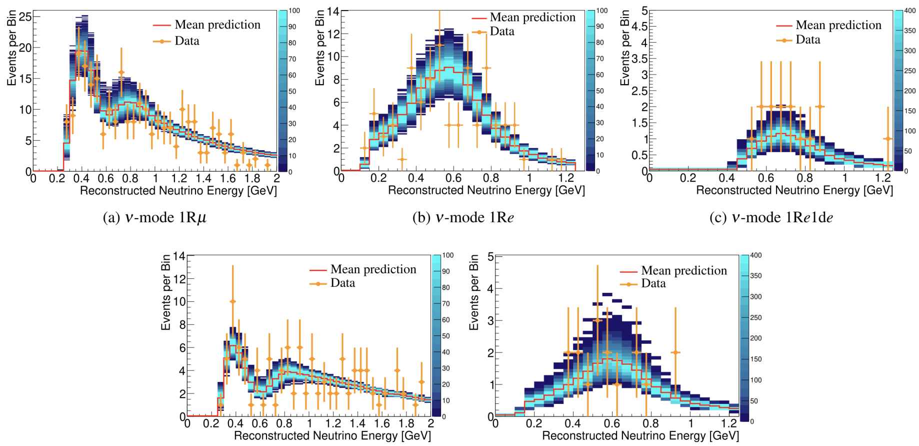

(a) ν-mode 1Rμ (b) ν-mode 1Re (c) ν-mode 1Re1de

(d) $\overline{\nu}$-mode 1Rμ (e) $\overline{\nu}$-mode 1Re

Figure 3.5: Neutrino energy spectra for the five SK samples. For each sample, the orange points represent data with Poisson error bars, and the colored bands correspond to the posterior density of predictions produced by the Bayesian fit. The 1Rμ samples show a clear dip corresponding with oscillations [91].

Results from the most recent T2K oscillation fit are shown in 3.6, both with and without external constraints on θ_{13} from reactor neutrino experiments.

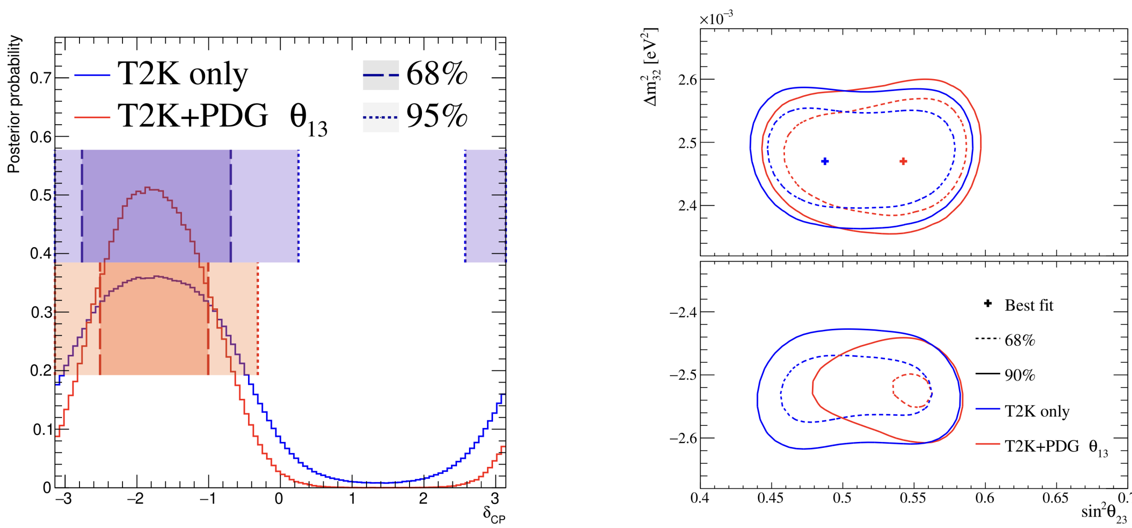

(a) δ_{CP} posterior probability (b) Δm^2_{32} and $\sin^2\theta_{23}$ credible intervals

Figure 3.6: Posterior probability and credible intervals for (a) δ_{CP} and (b) Δm^2_{32} and $\sin^2\theta_{23}$. The intervals and contours are presented both with (red) and without (blue) θ_{13} constraints from the Particle Data Group [91].

In summary, T2K finds that CP-conservation is included in the 95% (2σ) credible interval. T2K has a mild preference for the normal mass ordering, and a mild preference for the upper octant, but these are not statistically significant.

3.3 NOvA vs T2K

Given that NOvA and T2K are both probing similar physics, it is worth comparing and contrasting their measurements.

A useful tool for comparing NOvA and T2K is a "bi-event plot". These plots show how changes to oscillation parameters affect the $\overset{(-)}{\nu}_e$ appearance signals. In Figure 3.7, δ_{CP} and the sign of Δm^2_{32} are varied across all possible values, and all other oscillation parameters are fixed in value.

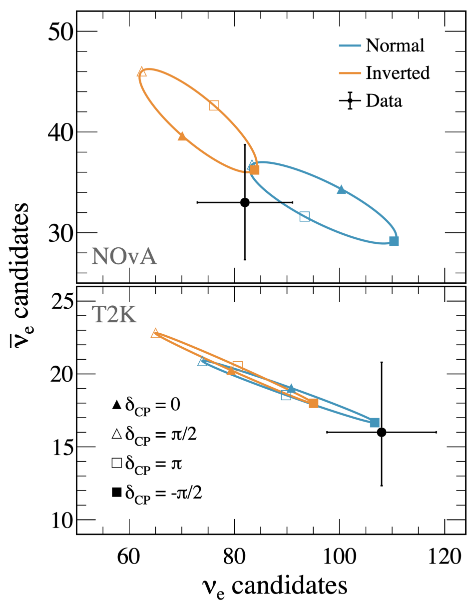

Figure 3.7: Bi-event plot for NOvA and T2K, depicting ν_e and $\overline{\nu}_e$ candidate events at the resepective far detectors. δ_{CP} is allowed to vary from $-\pi$ to π, and all other oscillation parameters are set to the maximum posterior probability values listed in Chapter 4 [92].

In the absence of matter effects, CP violation would manifest as a change in the appearance probabilities for ν_e and $\overline{\nu}_e$. For example, if $\delta_{CP} = +\frac{\pi}{2}$, then we would expect an excess of ν_e appearance and a deficit of $\overline{\nu}_e$ appearance relative to oscillations without CP violation. The opposite (a deficit of ν_e and an excess of $\overline{\nu}_e$) would be expected if $\delta_{CP} = -\frac{\pi}{2}$.

The introduction of matter effects allows for the disambiguation of the sign of Δm^2_{32}. In the normal ordering, we expect an excess of ν_e and deficit of $\overline{\nu}_e$ events relative to no CP violation and no matter effects, with the opposite being true for the inverted ordering.

NOvA's baseline is around thrice that of T2K. Due to the experiments having similar L/E, NuMI neutrinos therefore carry three times the energy as J-PARC neutrinos, leading to stronger matter effects. This means NOvA has a greater ability to disambiguate the neutrino mass ordering. However, this can actually make it more difficult to disambiguate δ_{CP}, as values of δ_{CP} around $+\frac{\pi}{2}$ in the normal ordering and $-\frac{\pi}{2}$ in the inverted ordering can produce similar $\overset{(-)}{\nu}_e$ appearance signals.

Thus, while the two experiments are probing the same physics, the different baselines make their measurements complementary. NOvA has greater sensitivity to the mass ordering, while T2K's measurement of δ_{CP} is uncorrelated from the mass ordering.

The actual results from NOvA and T2K are also compellingly different. T2K's data has shown an excess of ν_e-like and a deficit of $\overline{\nu}_e$-like events at SK (relative to oscillations without matter effects). In comparison, NOvA's FD data lies directly between the predicted event counts for normal and inverted mass orderings. T2K therefore has a stronger preference for CP violation than NOvA, although as seen in Section 3.2, it is still not statstically significant enough for a definitive statement.

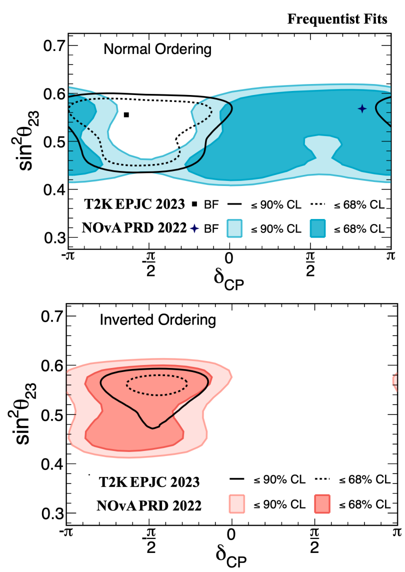

Figure 3.8: δ_{CP} vs $\sin^2\theta_{23}$ contours for NOvA (shaded) and T2K (solid black). Note that in the normal mass ordering, NOvA and T2K's confidence intervals do not greatly overlap, whereas there is more overlap in the inverted ordering [93].

Beyond the bi-event plot, we can look at the confidence intervals produced by the full oscillation analyses. In Figure 3.8, we see that there are few regions of the $\sin^2\theta_{23}$-δ_{CP} confidence intervals with significant overlap in the normal ordering. This "tension" has led some theorists to speculate that there may be new physics behind the difference

[94, 95]. However, there are still significant overlaps in the 95% confidence intervals, especially in the inverted ordering.

3.4 Why Do a Joint Fit?

Given the two experiments' complementarity, it is natural to consider combining the measurements. Qualitatively, Figure 3.8 already points to the idea that the inverted ordering would be preferred in a joint fit.

Approximate global fits based on the publicly available log-likelihood projections exist [15, 96]. However, these analyses are unable to use the full detector response models, near and far detector datasets, energy estimators, systematic uncertainty treatments, and full likelihood calculations over the full high-dimensional oscillation parameter space. The joint fit presented in Chapter 4 incorporates all of these features. A joint fit between the two experiments also provides an opportunity for cross-pollination before the next generation of oscillation experiments. Members of each collaboration would be provided opportunities to learn how the other collaboration conducts their oscillation analysis. In particular, the joint fit allows both collaborations to compare and contrast ways to incorporate near detector data and reduce the impact of systematic errors. Additionally, the experiments would be able to understand each others' neutrino interaction models. This would even provide opportunities to correlate neutrino interaction systematics where possible [97].

As a result, NOvA and T2K began pursuing a joint fit produced directly by working groups from the two collaborations.

Chapter 4

THE NOvA-T2K JOINT FIT

4.1 Preparing the Joint Fit

Between such large and complicated analyses, there were several architectural and technical challenges in order to produce a joint result. Comparing Figures 4.1a and 4.1b, we see that there are substantial differences between the way the two experiments structure their analyses. For technical reasons, it is not feasible to share code or merge the two code bases. Additionally, to ensure robustness of the result, we want both T2K and NOvA to be able to run the full joint fit and reproduce the final results.

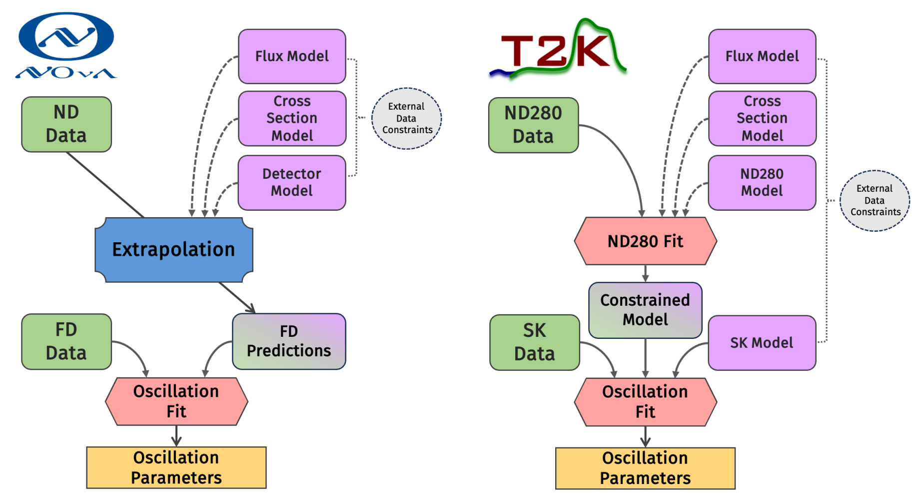

(a) The NOvA analysis (b) The T2K Analysis

Figure 4.1: Comparison of the NOvA and T2K analysis flows. While NOvA uses extrapolation to create data-driven FD predictions that are fit to FD Data, T2K first fits to ND280 data to constrain their model before final SK fits [93].

The solution is to containerize each experiment's likelihood function. Each experiment creates an Apptainer container which holds a streamlined version of the entire analysis. A fitter should be able to pass oscillation and systematic parameter settings into the container, and receive a log-likelihood. This required the creation of the Bifrost library to facilitate communication into and out of the container and the DummyLLH library to standardize the data format used for the fit parameters and log-likelihood [98, 99]. This architecture is sketched in Figure 4.2. Each experiment is able to marginalize and assign penalty terms to systematic parameters in order to construct a joint likelihood.

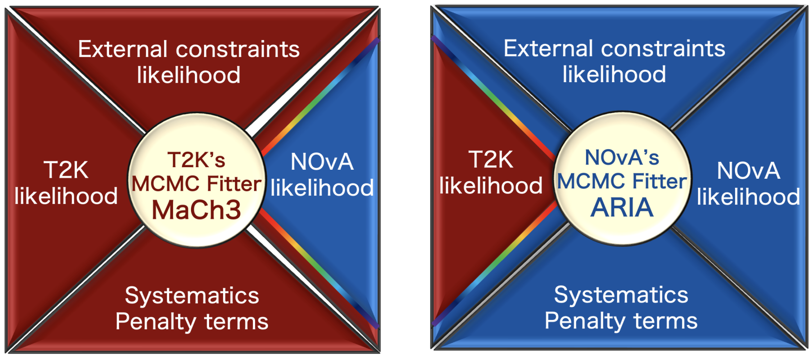

Figure 4.2: Schematic depicting the intended code structure for the joint fit. Red indicates NOvA code, and Blue indicates T2K code. The rainbow-shaded lines demarcate the boundaries of each container, so each experiment is unable to read the code of the other [93].

Because the final fits will end up including ~100 T2K parameters and all ~100 of NOvA's, MCMC fitters are perfect tools for this. Both collaborations have MR^2T^2 MCMC fitters: MaCh3 from T2K and Aria for NOvA, which are used for the final joint fits.

4.2 Asimov Data Points

As we test the joint analysis, we run fits to Asimov fake data [100]. This is fake data created by fixing the unknown physics parameters to set values, and without introducing any statistical fluctuations. In this thesis, I will be showing fits done at the following three Asimov points:

- **Asimov 0**: The NOvA best-fit point
- **Asimov 1**: The T2K best-fit point
- **Asimov 4**: Similar to the NuFIT global best fit (Inverted Ordering)

Asimov 0 and 1 are useful as it allows us to force each experiment to fit to data that looks like what the other saw. Finally Asimov 4 lets us test fits in the region of phase space where the final result may lie.

4.3 Should We Introduce Correlations?

Ultimately, NOvA and T2K are both trying to measure neutrino ocillations. As the underlying physics are the same, it's worth considering when and how to correlate systematic parameters between the two experiment, and whether an improper treatment of these would bias the final results. We can break these potential sources of correlation into three major sections: flux models, detector response models, and cross-section models. In each we need to consider which models/systematics are significant, if correlations can be drawn between the two experiments, and whether these correlations themselves are significant.

4.4 Flux Models

The two experiments see fluxes at very different energies and use different simulation software to predict fluxes. In addition, the two experiments use different data sets to tune their flux models.

NOvA uses 120 GeV protons from the NuMI beamline [62]. NOvA's flux model uses thin target data from the NA49 experiment, which collected data with a beam at a slightly higher energy of 158 GeV [101]. This is scaled to NOvA energies.

T2K's beam is at 31 GeV. They use NA61/SHINE data which operated at J-PARC beam energy and also uses a replica target of the T2K beam, in addition to thin target data. Figure 4.3 shows that the two hadronic production datasets are similar enough for the joint fit.

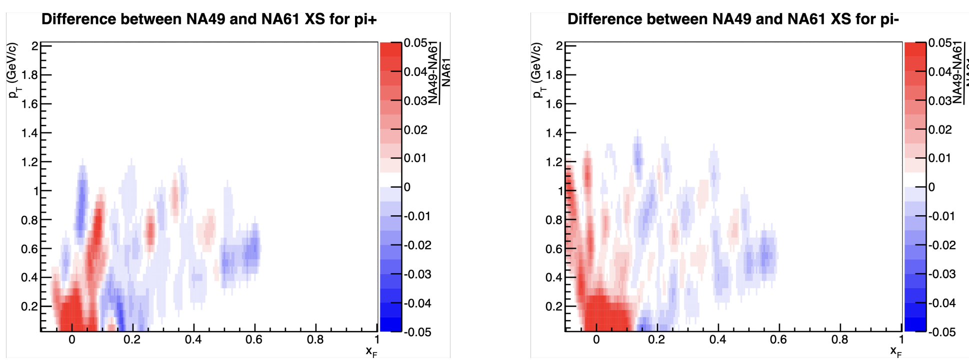

Figure 4.3: Comparison of hadronic production data from NA49 (used by NOvA) and NA61/SHINE (used by T2K) [102]. The left figure shows π^+ production, with π^- production on the right.

Studies have shown that there is a negligible difference between the two model tunings (see Figure 4.4).

T2K also did a study to verify that results are insensitive to the choice of Hadronic interaction model. NOvA uses GEANT4 v9.2.p03 with the FTFP_BERT model while T2K uses FLUKA 2011.2x and GEANT3. T2K shows that GEANT4 does not strongly affect the main beam region (see Figure 4.5).

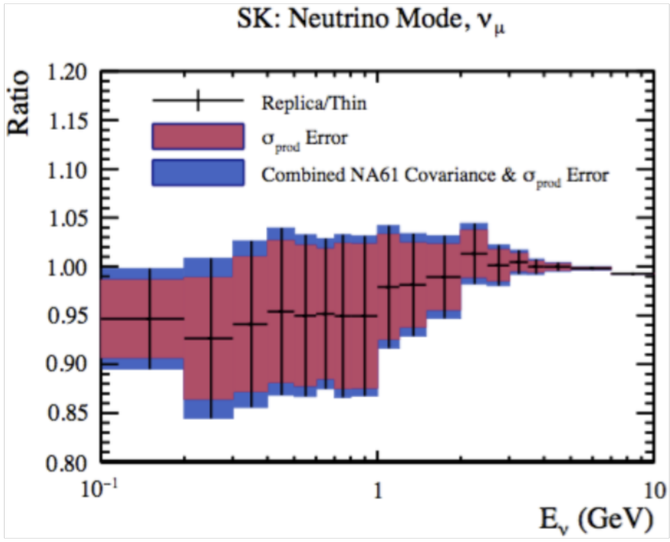

Figure 4.4: Comparison of T2K flux predictions using only thin-target data instead of replica target data from NA61/SHINE.

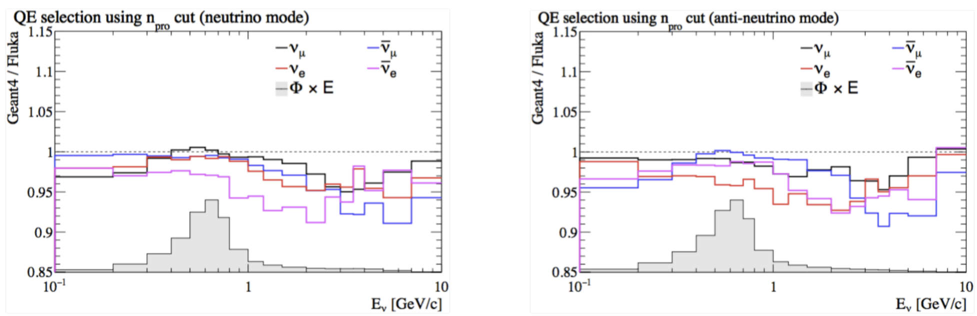

Figure 4.5: Comparison of T2K predicted QE selected fluxes using GEANT4 and FLUKA. ν beam on the left, and $\overline{\nu}$ beam on the right.

Finally, both experiments sought to find sources of of possible correlations in flux uncertainties. As many flux uncertainties are based on the specific source data and the specific experimental apparatus, the only possible place correlations can be introduced are in the hadronic interaction models. However, the data and tuning methods used by NOvA and T2K are quite different. Even more importantly, beam uncertainties have some of the smallest impacts on oscillation measurements (see Figure 2.21). As a result, the joint fit does not incorporate any flux correlations.

4.5 Detector Response Models

Detector calibration provides the largest single source of systematic uncertainty for both experiments. NOvA's calibration systematic is the largest source of error (Figure 2.21). On T2K, SK calibration error is the largest individual systematic.

NOvA's energy estimator depends on the material properties of the NOvA detector for calorimetry or to range out muons. This has no direct analog in SK. Each experiment also uses primarily internal data for calibration. As such, no correlations can be drawn from the between the two experiments' calibration systematics [103].

T2K handles pion multiplicity as separate selections, so secondary interactions lead to a powerful source of error. NOvA makes no such distinction. Because overall hadronic energy is estimated calorimetrically, secondary interactions from pions do not introduce significant error. Thus, even if it were possible to correlate T2K's pion selection systematics to NOvA, there would be little to no effect.

Inversely, low energy neutrons degrade NOvA's calorimetric energy estimation. This results in a large source of systematic error on NOvA. Since almost all interactions at T2K energies are quasielastic, neutrino energy estimation is done exclusively via lepton reconstruction. As a result, these low-energy neutrons do not have a large effect on T2K's oscillation results. Thus, correlations on neutron uncertainties can also be ignored.

4.6 Cross Section Models

The two experiments use completely different interaction simulation software. NOvA uses GENIE to simulate neutrino interactions, whereas T2K uses NEUT. The two analyses are so tighty tied to the choice of event generator that it is infeasible for the experiments to be able to switch to using the same software for the joint fit. The upshot is that it makes it very difficult to draw correlations between the two experiments. For a deeper dive into a possible avenue to accomplish this task, see Chapter 5.

First, the joint fit group sought to bracket the possible effects of correlating interaction systematics. Afterwards, we would look for systematics that can be correlated across the two experiments with minimal technical overhead. Towards this, we started by looking for the largest individual interaction systematics in each experiment.

A one-to-one mapping of systematics is impossible between NOvA and T2K. Table 4.1 shows a list of neutrino interaction systematics, broken down into broad categories.

Category	NOvA Parameters	T2K Parameters
CCQE	ZNormCCQE ZExpAxialFFSyst2020_EV1 ZExpAxialFFSyst2020_EV2 ZExpAxialFFSyst2020_EV3 ZExpAxialFFSyst2020_EV4 RPAShapeenh2020 RPAShapesupp2020	M_A QE Q2_norm_0 Q2_norm_1 Q2_norm_2 Q2_norm_3 Q2_norm_4 Q2_norm_5 Q2_norm_6 Q2_norm_7 EB Dial C nu EB Dial C nubar EB Dial O nu EB Dial O nubar
MEC	MECEnuShape2020Nu MECEnuShape2020AntiNu MECShape2020Nu MECShape2020AntiNu MECInitStateNPFrac2020Nu MECInitStateNPFrac2020AntiNu	2p2h Norm nu 2p2h Norm nubar 2p2h C to O 2p2h Shape C 2p2h Shape O 2p2h Edep low Enu 2p2h Edep high Enu 2p2h Edep low Enubar 2p2h Edep high Enubar
RES	MaCCRES MvCCRES MaNCRES MvNCRES LowQ2RESSupp2020	CA5 M_A RES ISO Bkg Low PPi ISO Bkg
FSI	hNFSI_MFP_2020 hNFSI_FateFracEV1_2020	FEFQE FEFQEH FEFINEL FEFABS FEFCX

Table 4.1: A list of all large cross-section parameters in NOvA (ignoring PCA), and all cross-section parameters in T2K. Parameters are grouped by true interaction mode. Adapted from [104].

T2K conducted internal studies to determine its most important systematics. This was accomplished by individually forcing each systematic parameter away from its central value, and observing how that would change credible intervals. The systematics resulting in the largest changes to credible intervals would be the final results. Of note, none of the systematic parameters were able to change the Δm^2_{32} credible intervals by a noteworthy amount. As a result, T2K created a fake extra-large energy scale systematic to amplify potential effects for these bracketing studies.

On the NOvA side, extrapolation means that the cross-section parametrization only affects measurements at the level of ND-FD differences. Before doing data fits, NOvA tunes the GENIE event generator to ND Data, making modifications where significant Data-MC disagreement arises [84]. See Section 2.9 for a further explanation of the NOvA cross section tune.

Table 4.2 summarizes the final results of the biggest individual systematics for NOvA and T2K.

Oscillation Parameter	Largest NOvA Systematic(s)	Largest T2K Systematic(s)
δ_{CP}	second class currents and radiative correction	$\sigma_{\nu_e}/\sigma_{\nu_\mu}$ and $\sigma_{\bar{\nu}_e}/\sigma_{\bar{\nu}_\mu}$
$\sin^2\theta_{23}$	Neutron visible energy	2p2h C-O scaling
Δm^2_{32}	Calibration	*7% SK energy scale

Table 4.2: The most impactful systematics for each experiment. Note that in T2K's actual analysis, the SK energy scale is only 2% and has been boosted to 7% here to make it impactful.

4.7 Impact of Correlations on the Joint Fit

Because T2K cannot use extrapolation to mitigate the effects of systematic errors, it needs to put very tight constraints on their cross-section model at ND280 in order to accurately predict event rates at SK. By imposing these constraints, interaction uncertainties become comparable to SK detector uncertainties in all selection samples.

If it were technically feasible to correlate the two models, an immediate side effect would be that NOvA's cross-section model would be constrained by T2K's ND280 fit. To bracket the impact of this effect, a "shrink-and-pull" study was done on NOvA, wherein for each systematic parameter, the central value is "pulled" away from the tuned value and the uncertainty is "shrunk". This leaves the fitter unable to return the parameter to the original tuned central value. If the shrunk-and-pulled model returns similar credible intervals, we can be sure that NOvA's oscillation measurements are indeed highly independent of any additional external constraints from ND280 on the cross-section model.

To estimate the magnitude of the "shrink" T2K's ND280 fit could put on NOvA, we observe how T2K interaction uncertainties are reduced by the ND280 fit. These reductions in uncertainty are averaged across the systematics in each interaction category to inform what shrink NOvA should apply for its systematics in that category. Since T2K is at too low an energy to provide any significant statsitical power on DIS events, NOvA's DIS systematics are left unchanged. Table 4.3 shows the effect of this procedure on each category.

Category	Average shrink
CCQE	44%
MEC	58%
RES	58%
DIS	100%
FSI	43%

Table 4.3: Average total magnitude for systematics in each category after the shrink imposed by the T2K ND280 fit in each systematic category.

To finish bracketing the effects of possible correlations, the model's central values are "pulled", since ND280 would suggest some shifting of the best-fit model. This was accomplished by throwing each NOvA systematic on a $[-1\sigma, +1\sigma]$ interval (before shrinking the uncertainties). This was repeated a few times, and fits were done as long as the overall bias on the mean neutrino energy, hadronic energy and muon energy are all less than 2%, 2%, and 5% respectively for both FHC and RHC, to ensure the model remains reasonably compatible with NOvA's own ND data.

The results of this study on NOvA contours can be seen in Figure 4.6. The fact that the contours hardly change indicates that NOvA's extrapolation is robust to potential constraints from the T2K ND280 Fit.

T2K ran its own version of the shrink-and-pull study where it included the NOvA likelihood to test the impact on joint results. After shrinking and pulling each systematic individually, the only systematics which meaningfully affected the contours were those large systematics already identified in Table 4.2.

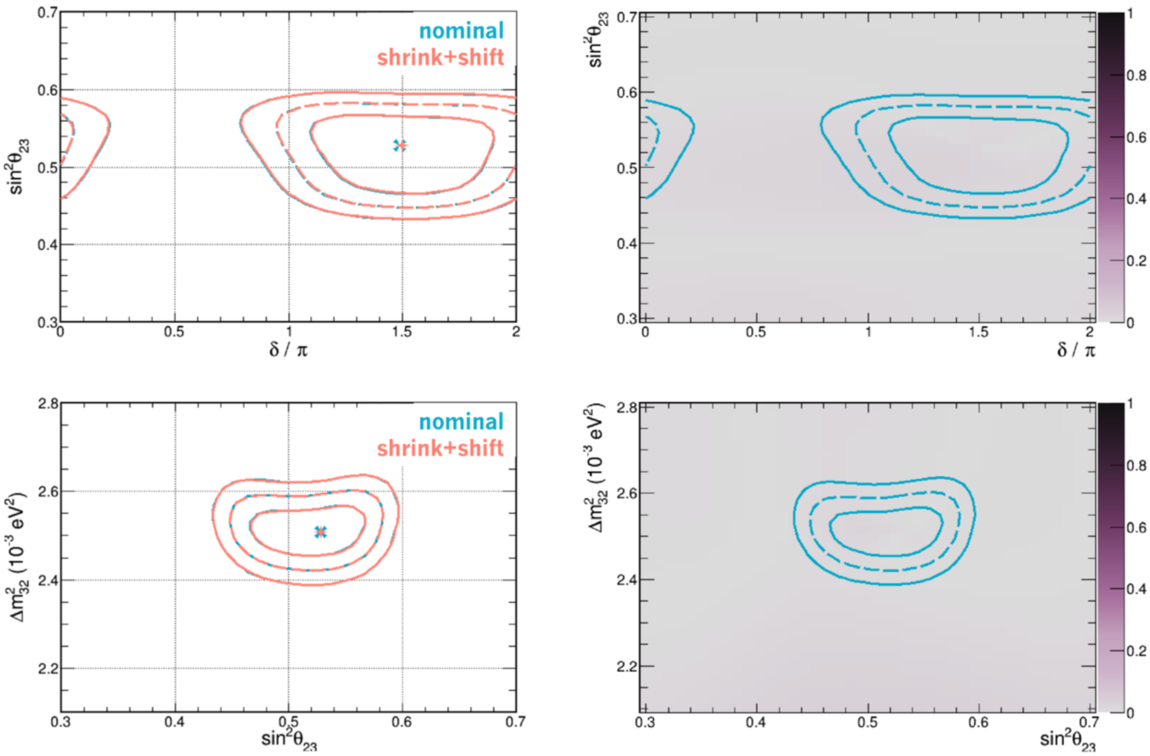

Figure 4.6: Contours in oscillation parameters due to shrink caused by T2K. Left plots are a comparison of the contours in δ_{CP} vs $\sin^2\theta_{23}$ (top) and $\sin^2\theta_{23}$ vs Δm^2_{32} respectively. The right plots show the difference between the two log-likelihood surfaces ("nominal" - "shrink+shift"). These differences are close to 0 everywhere, indicating that NOvA's fit procedure is robust to any strong constraints imposed by the T2K ND Fit. The blue contours on the right plots are only to provide a visual aid of where the most relevant regions in this space lie.

With the list of impactful systematics from Table 4.2 in hand, we can impose correlations or anti-correlations between these systematics and evaluate how that affects the final fit. Figures 4.7, 4.8 and 4.9 show the impacts on δ_{CP}, $\sin^2\theta_{23}$, and Δm^2_{32} of adding these correlations or anticorrelations. There is no appreciable impact of correlations on these credible intervals.

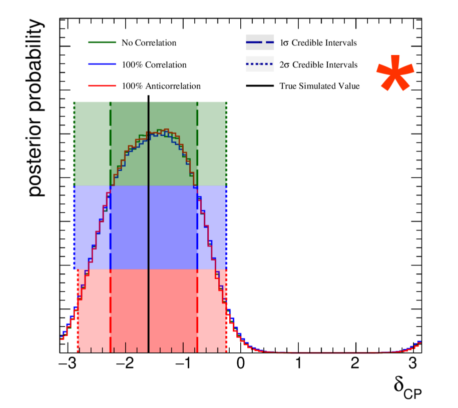

(a) Marginalized posterior in δ_{CP}, assuming correlations most impactful to δ_{CP} (Asimov 1)

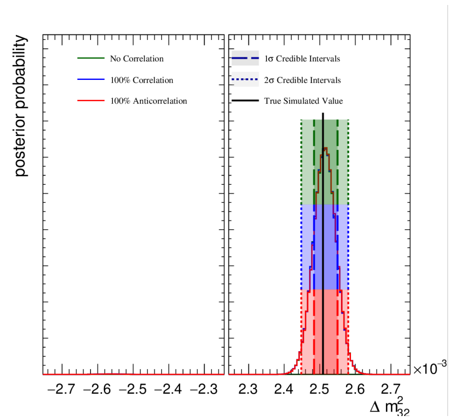

(b) Marginalized posterior in Δm^2_{32}, assuming correlations most impactful to δ_{CP} (Asimov 1)

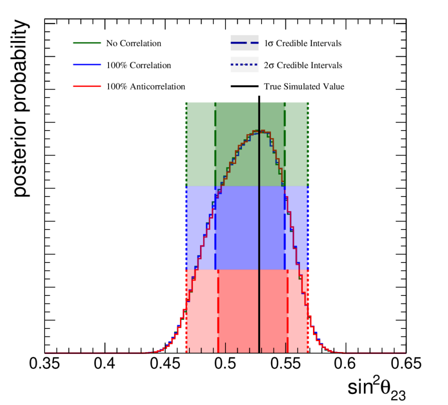

(c) Marginalized posterior in $\sin^2\theta_{23}$, assuming correlations most impactful to δ_{CP} (Asimov 1)

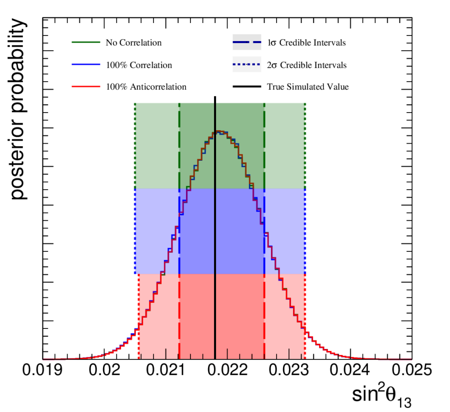

(d) Marginalized posterior in $\sin^2\theta_{13}$, assuming correlations most impactful to δ_{CP} (Asimov 1)

Figure 4.7: Impact of correlating or anticorrelating the systematics most important to δ_{CP} measurements. Namely, T2K's ν_e-ν_μ and $\overline{\nu}_e$-$\overline{\nu}_\mu$ cross section ratios are correlated to NOvA's second class current and radiative correction systematics. δ_{CP}, which should most be affected by correlating these sysetmatics, is marked with an asterisk. Note that there is almost no change in credible intervals due to either correlating or anticorrelating these systematics.

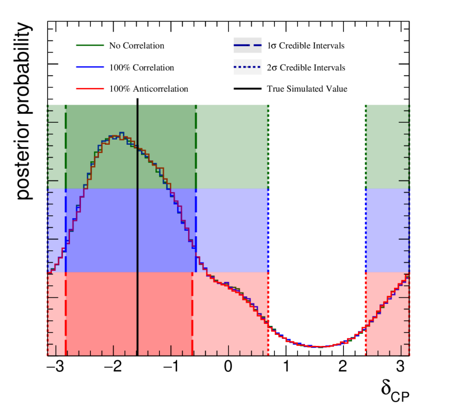

(a) Marginalized posterior in δ_{CP}, assuming correlations most impactful to $\sin^2\theta_{23}$ (Asimov 4)

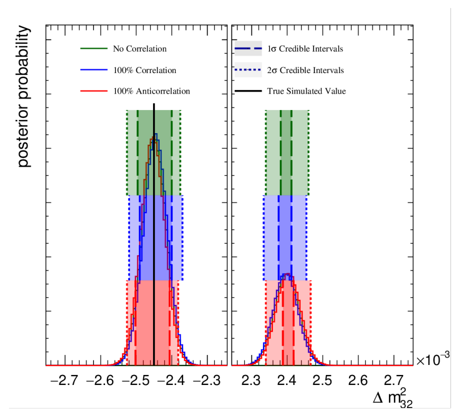

(b) Marginalized posterior in Δm^2_{32}, assuming correlations most impactful to $\sin^2\theta_{23}$ (Asimov 4)

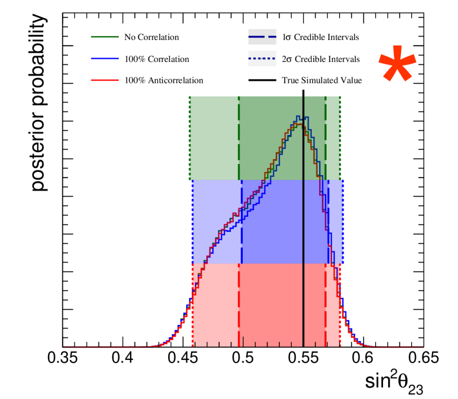

(c) Marginalized posterior in $\sin^2\theta_{23}$, assuming correlations most impactful to $\sin^2\theta_{23}$ (Asimov 4)

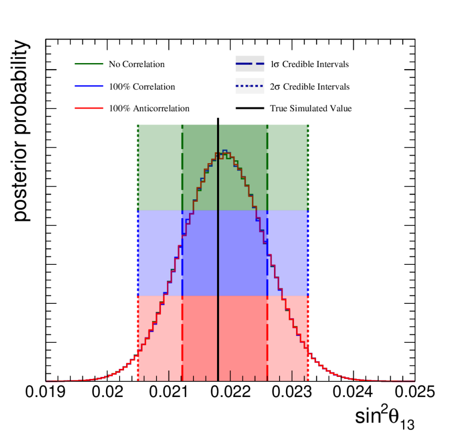

(d) Marginalized posterior in $\sin^2\theta_{13}$, assuming correlations most impactful to $\sin^2\theta_{23}$ (Asimov 4)

Figure 4.8: Impact of correlating or anticorrelating the systematics most important to $\sin^2\theta_{23}$ measurements. Specifically, T2K's 2p2h C-O scaling is correlated with NOvA's neutron visible energy systematic. $\sin^2\theta_{23}$, which should most be affected by correlating these sysetmatics, is marked with an asterisk. Note that there is almost no change in credible intervals due to either correlating or anticorrelating these systematics.

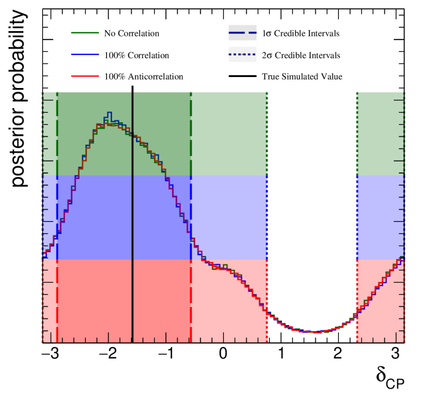

(a) Marginalized posterior in δ_{CP}, assuming correlations most impactful to Δm^2_{32} (Asimov 4)

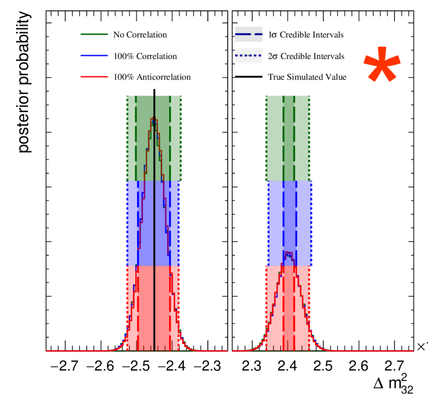

(b) Marginalized posterior in Δm^2_{32}, assuming correlations most impactful to Δm^2_{32} (Asimov 4)

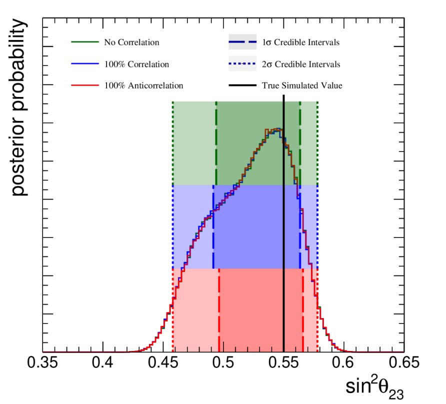

(c) Marginalized posterior in $\sin^2\theta_{23}$, assuming correlations most impactful to Δm^2_{32} (Asimov 4)

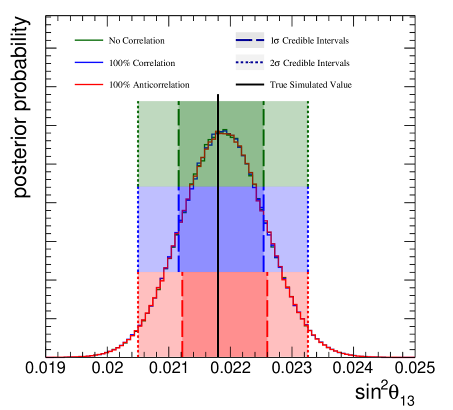

(d) Marginalized posterior in $\sin^2\theta_{13}$, assuming correlations most impactful to Δm^2_{32} (Asimov 4)

Figure 4.9: Impact of correlating or anticorrelating the systematics most important to Δm^2_{32} measurements. Specifically, T2K's 7% SK energy scale is correlated with NOvA's calibration systematics. Δm^2_{32}, which should most be affected by correlating these sysetmatics, is marked with an asterisk. Note that there is almost no change in credible intervals due to either correlating or anticorrelating these systematics.

4.8 The "Nightmare" Studies

The above plots show that correlating individual systematics would have no appreciable effect on the joint fit. What if, instead, correlations and anti-correlations between several systematic parameters somehow conspired to sabotage the joint fit? We can more directly explore how correlations can change our fit results by creating new systematic errors that are on the scale of statistical uncertainty.

T2K identified two so-called "nightmare" paremeters:

- **The Δm^2_{32} nightmare**: Shifting the reconstructed neutrino energy would bias the oscillation maximum region and would affect Δm^2_{32} measurements.
- **The $\sin^2\theta_{23}$ nightmare**: Changing the normalization of reconstructed event counts in the oscillation dip region would affect $\sin^2\theta_{23}$ (and to a lesser extent δ_{CP} measurements).

NOvA identified the following nightmare parameter to complement T2K's nightmares:

- **The NOvA nightmare**: An artificially huge boost (4σ) to the neutron systematic could cause NOvA's FD spectra to vary on the scale of statistical uncertainty.

Given the size of these fake systematic parameters, the impact of correlations should be readily apparent. To test this, we create fake data sets where NOvA and T2K's nightmares are correlated. We then compare likelihood contours generated by doing the fits where we assume no correlation, 100% correlation (correct), or 100% anticorrelation (incorrect). The output contours are shown in Figures 4.10 and 4.11.

Once again, the fits are largely robust to the affect of such large correlations. As long as we do not pick the wrong correlation (e.g. assuming anticorrelation when in reality two parameters are positively correlated), we are safe neglecting the effects of correlations.

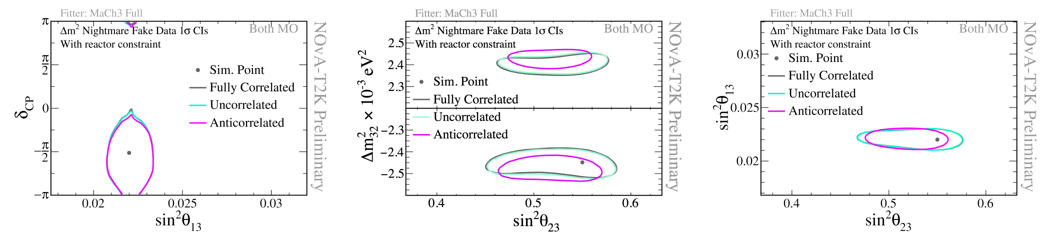

(a) $\sin^2\theta_{13}$ vs δ_{CP} (b) $\sin^2\theta_{23}$ vs Δm^2_{32} (c) $\sin^2\theta_{23}$ vs $\sin^2\theta_{13}$

Figure 4.10: Impact of assuming the wrong correlation with the T2K Δm^2_{32} nightmare and the NOvA neutron nightmare. These nightmare parameters should strongly impact Δm^2_{32} measurements, and indeed we see that incorrectly assuming an anticorrelation between the nightmare parameters leads to markedly different contours. However, assuming no correlation mostly recovers the same contours. Thus, we see that it is safer to ignore correlations between NOvA and T2K systematics than to assume an incorrect correlation.

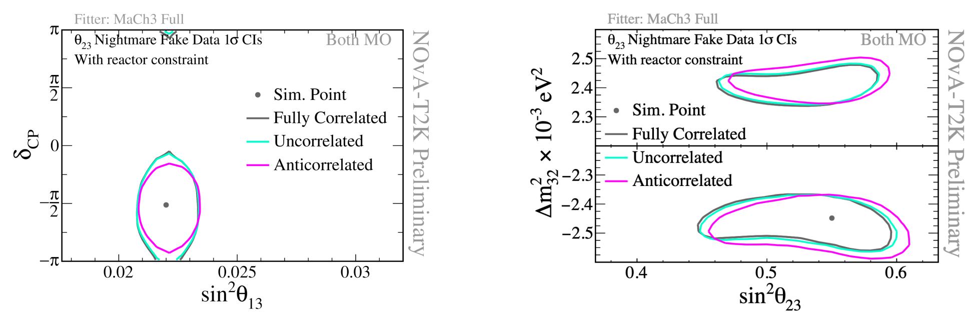

(a) $\sin^2\theta_{13}$ vs δ_{CP} (b) $\sin^2\theta_{23}$ vs Δm^2_{32}

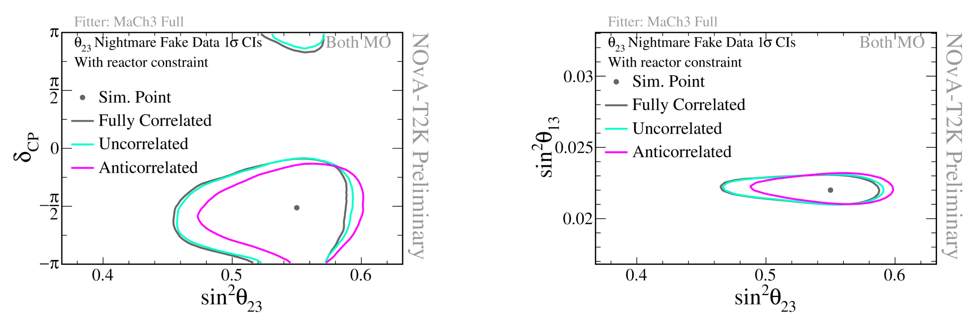

(c) $\sin^2\theta_{23}$ vs δ_{CP} (d) $\sin^2\theta_{23}$ vs $\sin^2\theta_{13}$

Figure 4.11: Impact of assuming the wrong correlation with the T2K $\sin^2\theta_{23}$ nightmare and the NOvA neutron nightmare. These nightmare parameters should strongly impact $\sin^2\theta_{23}$ measurements, as well as δ_{CP} to a lesser extent. Indeed, we see that incorrectly assuming an anticorrelation between the nightmare parameters leads to markedly different contours. However, assuming no correlation mostly recovers the same contours. Thus, we see that it is safer to ignore correlations between NOvA and T2K systematics than to assume an incorrect correlation.

4.9 Correlating δ_{CP} parameters

Both experiments incorporate systematic parameters to account for higher-than-tree-level corrections to $\overset{(-)}{\nu}_e$ cross-sections [105]. These parameters are relevant to δ_{CP} (see Table 4.2).

This makes it very achievable to correlate these parameters between the two experiments. We start with the following correlation matrix, which forces ν_e and $\overline{\nu}_e$'s to behave the same way between the two experiments, and puts the desired anticorrelations between ν_e and $\overline{\nu}_e$'s.

$$
\begin{array}{c}
 \\
\sigma_{\nu_e}^{\text{T2K}} \\
\sigma_{\overline{\nu}_e}^{\text{T2K}} \\
\sigma_{\nu_e}^{\text{NOvA}} \\
\sigma_{\overline{\nu}_e}^{\text{NOvA}}
\end{array}
\begin{array}{c}
\begin{array}{cccc}
\sigma_{\nu_e}^{\text{T2K}} & \sigma_{\overline{\nu}_e}^{\text{T2K}} & \sigma_{\nu_e}^{\text{NOvA}} & \sigma_{\overline{\nu}_e}^{\text{NOvA}}
\end{array} \\
\begin{bmatrix}
1 & -0.5 & 1 & -0.5 \\
-0.5 & 1 & -0.5 & 1 \\
1 & -0.5 & 1 & -0.5 \\
-0.5 & 1 & -0.5 & 1
\end{bmatrix}
\end{array}
$$

Converting the above into a covariance matrix on systematic shifts and taking into account NOvA's normalization on systematic shifts, we get the following:

$$
\begin{array}{c}
 \\
\sigma_{\nu_e}^{\text{T2K}} \\
\sigma_{\overline{\nu}_e}^{\text{T2K}} \\
\sigma_{\nu_e}^{\text{NOvA}} \\
\sigma_{\overline{\nu}_e}^{\text{NOvA}}
\end{array}
\begin{array}{c}
\begin{array}{cccc}
\sigma_{\nu_e}^{\text{T2K}} & \sigma_{\overline{\nu}_e}^{\text{T2K}} & \sigma_{\nu_e}^{\text{NOvA}} & \sigma_{\overline{\nu}_e}^{\text{NOvA}}
\end{array} \\
\begin{bmatrix}
8\times10^{-4} & -4\times10^{-4} & 0.04 & -0.02 \\
-4\times10^{-4} & 8\times10^{-4} & -0.02 & 0.04 \\
0.04 & -0.02 & 2 & -1 \\
-0.02 & 0.04 & -1 & 2
\end{bmatrix}
\end{array}
$$

By including this covariance matrix in the fit, we are able to cleanly account for the correlations in these specific uncertainties.

4.10 Fake Data Fits

T2K's fit procedure (see Figure 4.1b) requires a fit of ND280 Data with a very sophisticated cross-section model to provide constraints for the final SK fit. To make sure their analysis still has the flexibility to fit data produced from out-of-model variations, T2K runs several fits using fake data generated with alternative models. Ideally, T2K should be able to recover very similar results using these fake data. These so-called Fake Data Studies (FDS) are the primary method by which T2K can stress-test its fit procedures. These are very similar to the "nightmare" studies mentioned in Section 4.8. The joint fit thus sought to conduct our own fake data studies.

Two of these FDS were to see how the joint fits would be affected by effectively only using one experiment's cross-section tune. This is of course complicated by the fact that NOvA and T2K do not tune their cross-section model to the other experiment's energy range. As such these are not true FDS, so these results were not used to assess the bias of the fits and are instead used as a sanity check.

Two other fake data studies were based on models that proved to be the most impactful to the T2K-only 2020 analysis. A final FDS was to address inconsistencies in pion secondary interactions from the T2K 2020 analysis. In summary, the sive FDS are:

- **T2K-like reweight**: (Sanity check, not used to assess joint fit bias) A reweight of NOvA's GENIE cross-section model to resemble the NEUT model used by T2K after the ND280 fit.
- **NOvA-like reweight**: (Sanity check, not uesd to assess joint fit bias) A reweight of T2K's NEUT model to resemble the GENIE model used by NOvA.
- **Minerva 1π**: Significant in the 2020 T2K Analysis. Motivated by the low-Q^2 suppression of pion production observed in MINERvA and in bubble chamber experiments at Argonne and Brookhaven relative to GENIE v2 [36]. This FDS implements this low-Q^2 suppression.
- **NonQE**: Significant in the 2020 T2K Analysis. Motivated by the fact that CC0π events are underpredicted. The resulting pulls on the QE model are quite large, so this FDS instead accounts for this CC0π excess using non-QE parameters.
- **Pion SI**: Confirming the fix of a bug in T2K's 2020 analysis. This bug centered around a mismatch between NEUT and GEANT4, where secondary interactions from pions were not getting properly propagated between the two models.

These analyses are repeated at each of the three Asimov points described in Section 4.2.

Bias Metrics

Before embarking on the FDS, we need to decide on metrics to determine whether the results of an FDS are acceptable or show evidence of significant bias.

The final requirements are:

$$\Delta_{\text{mid}} \leq 0.25 \cdot 2\sigma_{\text{syst.}}$$
$$\Delta\sigma \leq 0.1 \cdot \sigma_{\text{Asimov}}$$

where:

- Δ_{mid} is the change in the midpoint of the credible interval;
- $\sigma_{\text{syst.}} = \sqrt{\sigma^2_{\text{Asimov}} - \sigma^2_{\text{stat.}}}$ is the systematic uncertainty;
- $\sigma_{\text{stat.}}$ is statistical uncertainty;
- $\Delta\sigma = |\sigma_{\text{Asimov}} - \sigma_{\text{FDS}}|$ is the change in uncertainty width due to this FDS;
- σ_{Asimov} and σ_{FDS} are the overall fit uncertainties for Asimov data and Fake Data, respectively.

Put in words: the midpoint of the credible interval should not be shifted too far relative to the size of systematic uncertainty, and the change in uncertainty should not be too large relative to the overall uncertainty from fits to Asimov data.

FDS Results

Figure 4.12 shows the result of shifts in an example FDS (Minerva 1π). While shifting the model causes ND spectra to change, extrapolation means that FD spectra are resilient to this change (Figure 4.13).

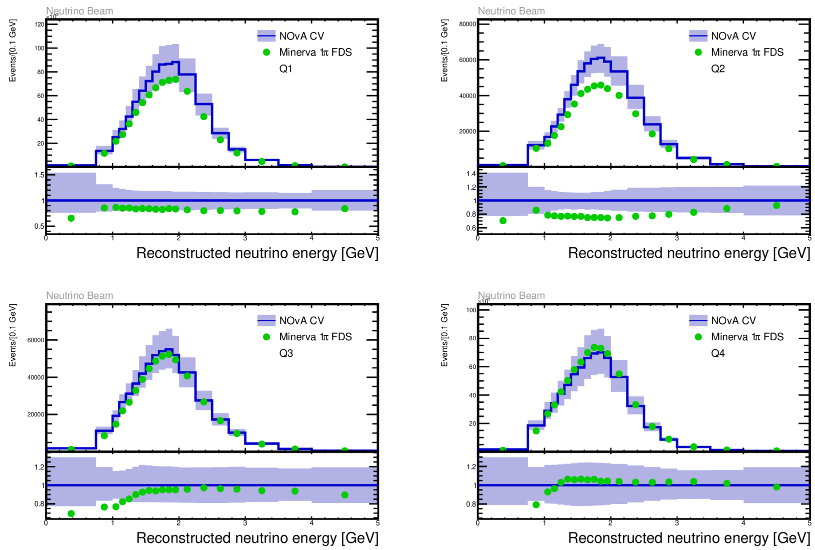

Figure 4.12: Comparisons of ND spectra from the Minerva 1π FDS. This alternative model can substantially affect ND spectra.

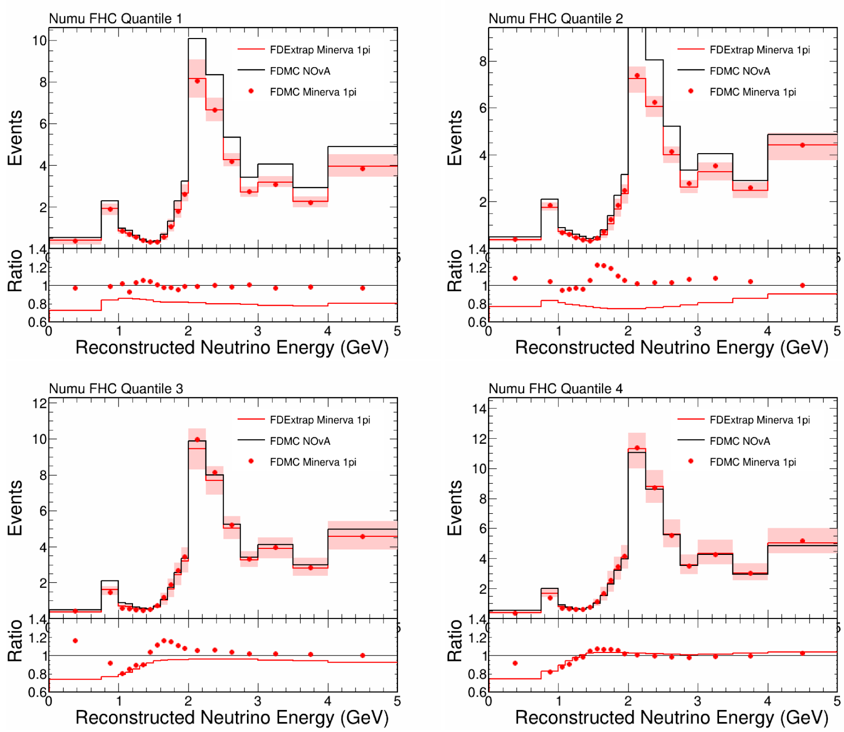

Figure 4.13: Comparisons of FD spectra from the Minerva 1π FDS. The black curve shows the initial FD MC spectrum. The red dots depict the shifted fake data, and the red curve with the drawn error band shows the result of extrapolation. NOvA's extrapolation procedure is resilient to the effect of this FDS. This is reflected in table 4.4, where we see that none of the FDS are sufficient to cause issues.

Table 4.4 shows a summary of the bias metrics. All FDS pass.

Oscillation Parameter	FDS	$\max \frac{\Delta_{\text{mid}}}{2\sigma_{\text{syst.}}}$ (Asimov) Target: $< 25\%$		$\max \frac{\Delta\sigma}{\sigma_{\text{Asimov}}}$ (Asimov) Target: $< 10\%$	
$\sin^2\theta_{23}$	MINERvA 1π	14.38%	(0)	9.6%	(1)
	non-QE	6.93%	(1)	4.7%	(4)
	pion SI	2.39%	(0)	4.7%	(4)
	T2K-like	0%	(1)	4.7%	(4)
	NOvA-like	3.47%	(1)	4.7%	(4)
Δm^2_{32}	MINERvA 1π	3.47%	(4)	3.6%	(4)
	non-QE	19.46%	(0)	7.1%	(0)
	pion SI	10.21%	(4)	3.7%	(1)
	T2K-like	12.13%	(4)	4.0%	(1)
	NOvA-like	4.81%	(0)	4.0%	(1)

Table 4.4: Summary table of the bias metric results. The Asimov point where bias metrics are maximized is noted in parentheses. These bias metrics were calculated without accounting for the fact that some of these modified Asimov "universes" lend themselves to intrinsically different oscillation sensitivity. This means that the above bias metrics are actually somewhat inflated, but they all pass anyway. Note that for some FDS, statistical credible intervals are discontinuous, which makes it impossible to calculate bias metrics as described in the text. In these cases, the changes in credible intervals and contours are inspected visually. In summary, no FDS shows any strong bias, including the T2K-like and NOvA-like sanity-check studies. Adapted from [106].

4.11 Data Fits

Between the two experiments, there are 1,010 neutrino events, distributed as shown in Table 4.5.

Channel	NOvA	T2K
ν_e	82	94 (ν_e) 14 (ν_e 1π)
$\bar{\nu}_e$	33	16
ν_μ	211	318
$\bar{\nu}_\mu$	105	137

Table 4.5: Number of neutrino events collected by each experiment up to 2020. There are a total of 1,010 events between the two experiments, with roughly equal statistics across the two.

Figures 4.14 and 4.15 show NOvA and T2K data with predictions drawn from the joint-fit posterior distribution overlaid.

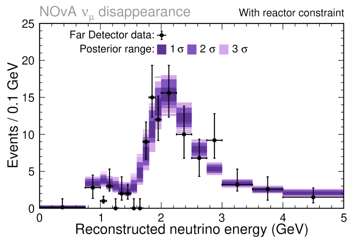

(a) ν_μ sample

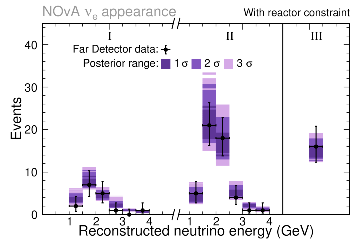

(b) ν_e sample

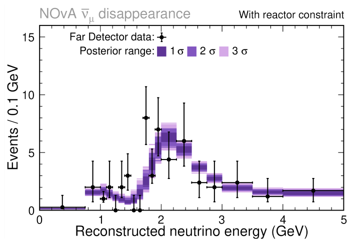

(c) $\overline{\nu}_\mu$ sample

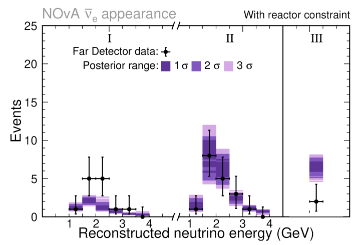

(d) $\overline{\nu}_e$ sample

Figure 4.14: Posterior predicted spectra from the NOvA-T2K joint fit for NOvA. Note that here the $\overset{(-)}{\nu}_\mu$ samples are integrated over the four hadronic energy fraction quantiles [107].

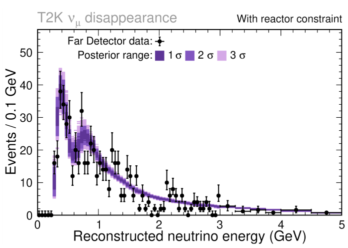

(a) ν_μ sample

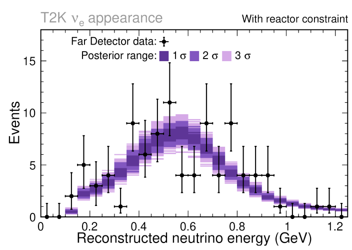

(b) ν_e 0π sample

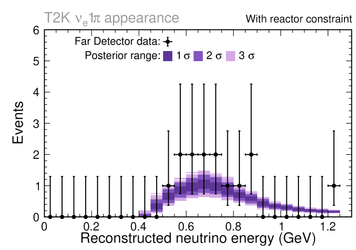

(c) ν_e 1π sample

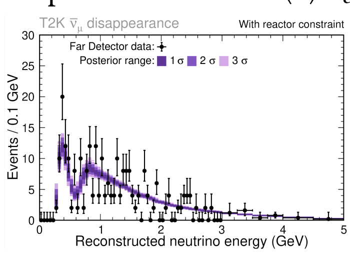

(d) $\overline{\nu}_\mu$ sample

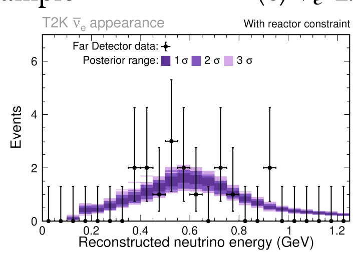

(e) $\overline{\nu}_e$ sample

Figure 4.15: Posterior predicted spectra from the NOvA-T2K joint fit for the five T2K samples [107].

We can test the compatibility of the datasets with the posterior distribution via the posterior predictive p-values (PPP) of the experiments [108]. These values show that our data is described well by the joint fit (Table 4.6).

NOvA and T2K are not as sensitive to θ_{13} as reactor neutrino experiments. As a result, we can stand to benefit by using reactor measurements as a constraint on the

Channel	NOvA	T2K	Combined
ν_e	0.90	0.19 (ν_e) 0.79 (ν_e 1π)	0.62
$\overline{\nu}_e$	0.21	0.67	0.40
ν_μ	0.68	0.48	0.62
$\overline{\nu}_\mu$	0.38	0.87	0.72
Total	0.64	0.72	0.75

Table 4.6: Posterior predictive p-values for each dataset. All PPP values show good agreement with data (an ideal PPP value is 0.5).

posterior distributions.

All of the following plots are made both with and without a reactor constraint. Unless otherwise noted, all of the following plots show the posterior distribution without assumption of the neutrino mass ordering (Both MO), the posterior distribution assuming the NO (NO Conditional), and the posterior distribution assuming the IO (IO Conditional). Finally, all of the following plots presented here were made with NOvA's Aria fitter.

Looking at the δ_{CP} posteriors we can draw a few conclusions (Figure 4.16). First, we see that the NOvA-T2K joint fit can exclude $\delta_{CP} = \frac{\pi}{2}$ from 3σ credible intervals. Switching to the Jarlskog invariant (Figure 4.17), we see that $J = 0$ (CP conservation) is excluded in the inverted ordering from 3σ credible intervals (Figure 4.17).

We see noteably tight bounds on Δm^2_{32} (Figure 4.18), and as will be discussed in the next section, this measurement is world-leading.

Figure 4.19 shows how NOvA-T2K's θ_{13} measurement gives a less precise measurement than the 2020 PDG reactor angle measurement of $\sin^2 2\theta_{13} = 0.085 \pm 0.0027$. Figure 4.20 shows that reactor experiments help us break the degeneracy in θ_{23} and turn a slight preference for the lower octant into a slight preference for the upper octant [109]. Figure 4.21 shows the joint contours in θ_{23} vs θ_{13}, which demonstrates how this degeneracy breaking comes about.

Finally, Figures 4.22, 4.23, 4.24, and 4.25 show more 2D contours over various pairs of oscillation parameters.

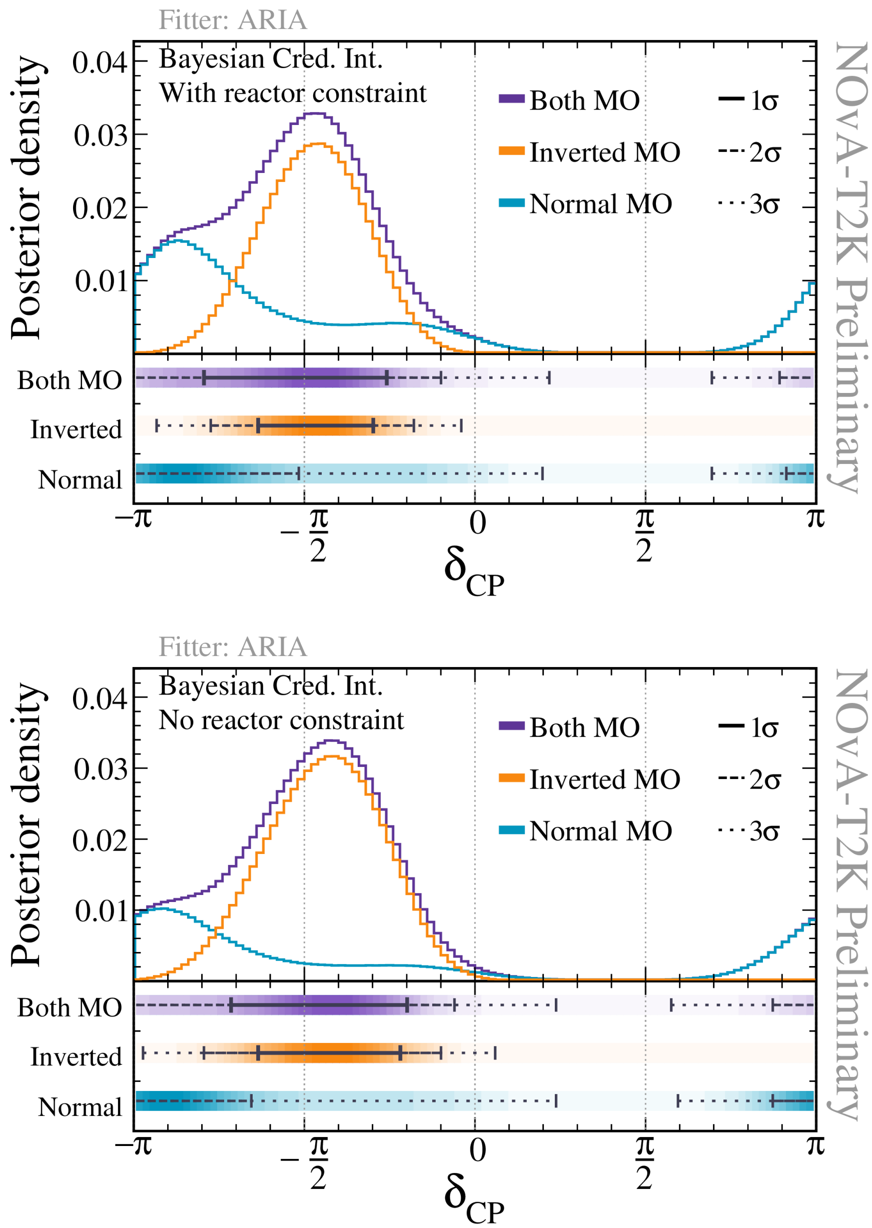

Figure 4.16: 1D posterior distributions of δ_{CP}. Posteriors are made both with (top) and without (bottom) the θ_{13} reactor constraint. In contrast to later figures, the NO and IO curves (blue and orange respectively) are non-conditional, meaning they show how the overall posterior density on δ_{CP} is distributed between the NO and IO.

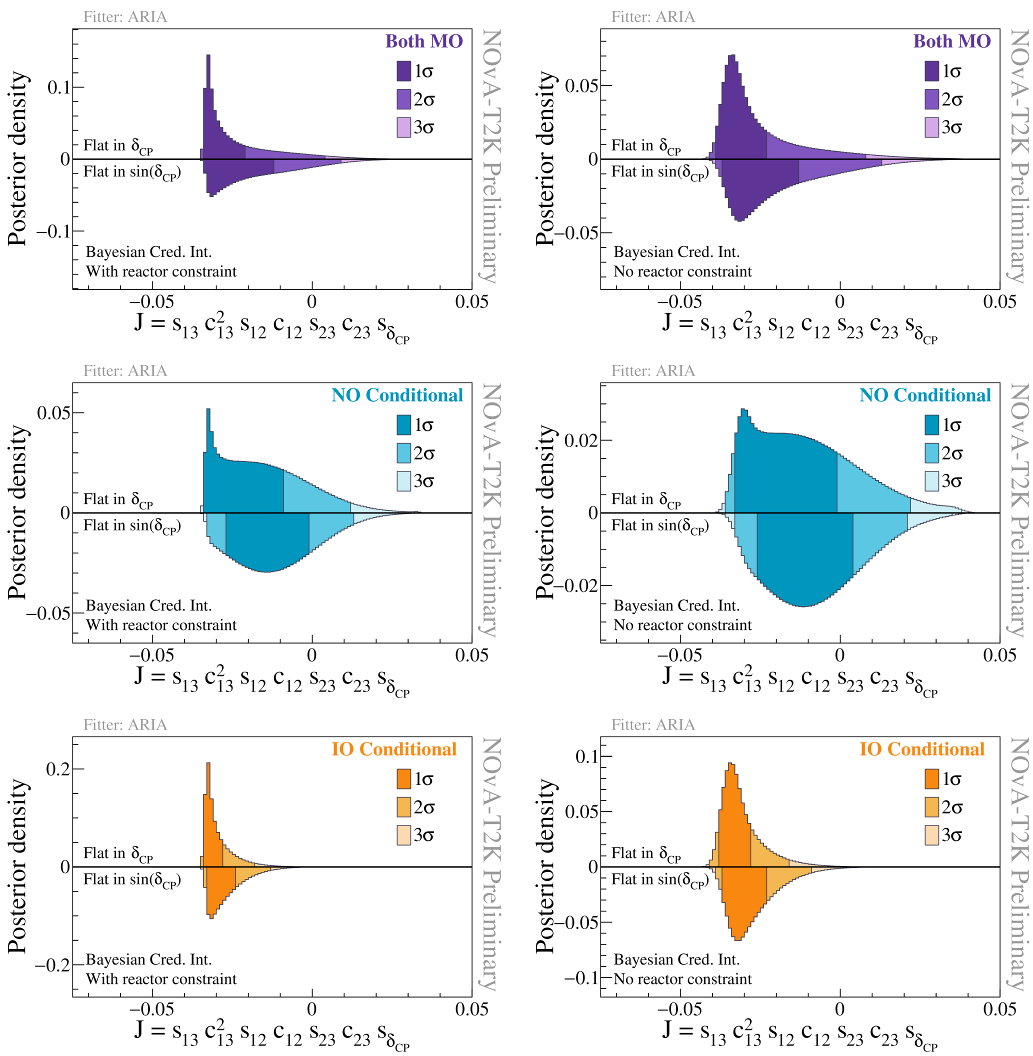

Figure 4.17: 1D posterior distributions of the Jarlskog Invariant J, using two different priors. Posteriors are made over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

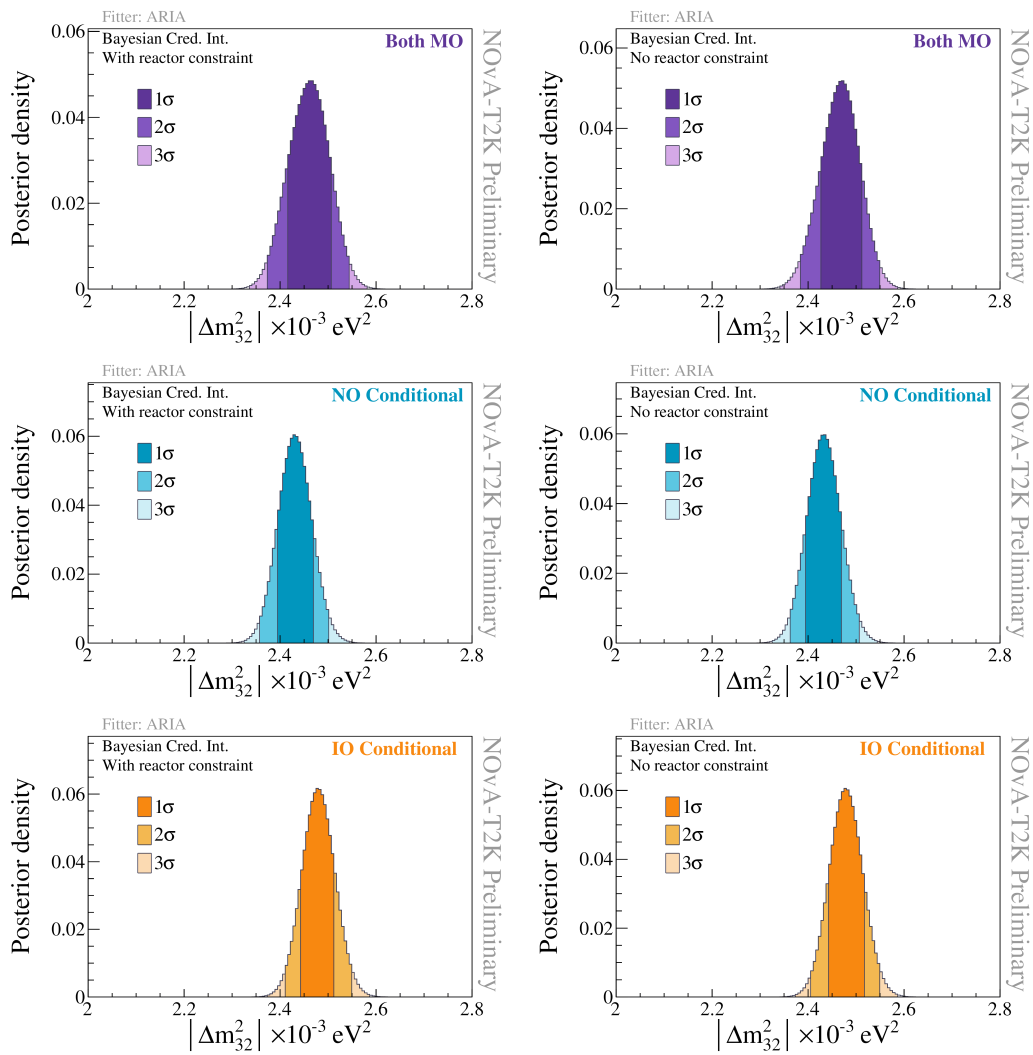

Figure 4.18: 1D posterior distributions of $\left|\Delta m^2_{32}\right|$. Posteriors are made using the Aria fitter, over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

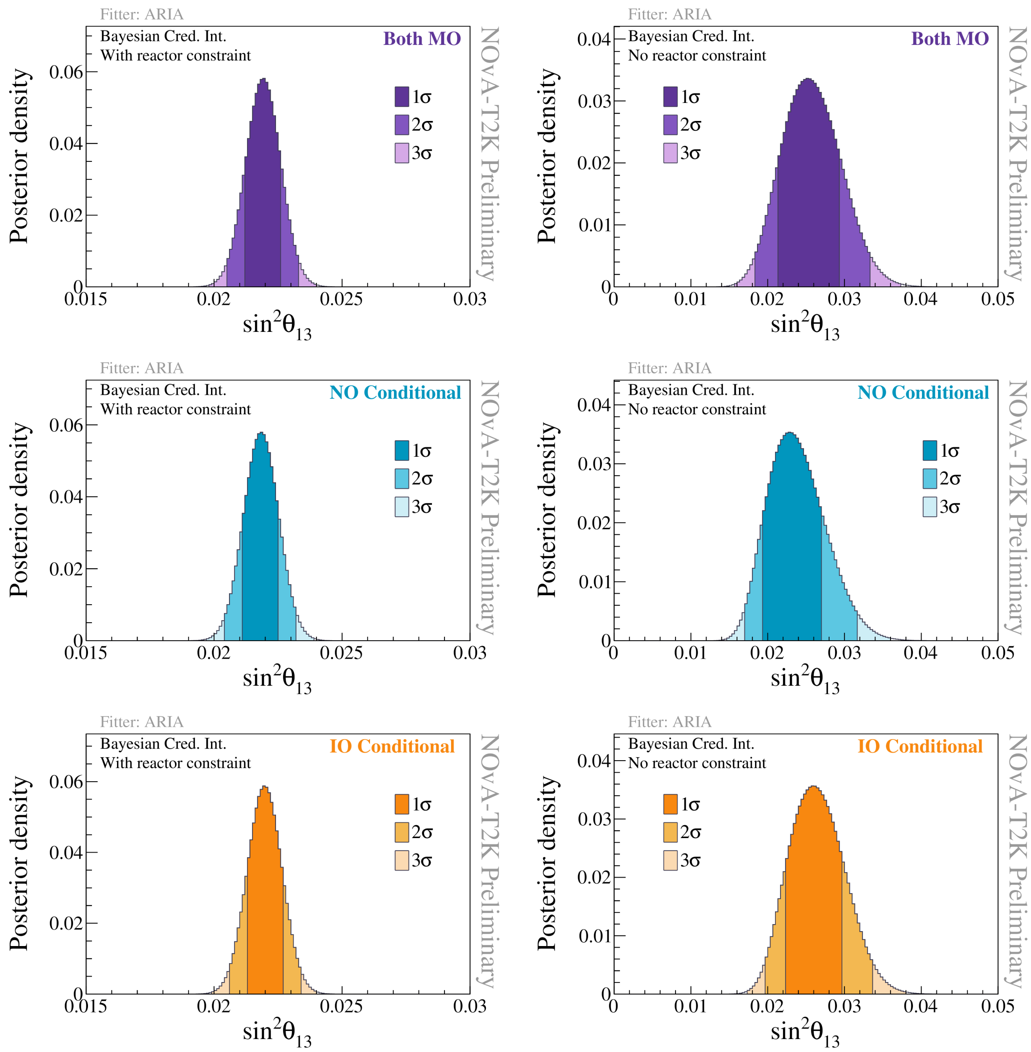

Figure 4.19: 1D posterior distributions of $\sin^2\theta_{13}$. Posteriors are made using the Aria fitter, over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint. Note the change in x-axis when adding the reactor constraint.

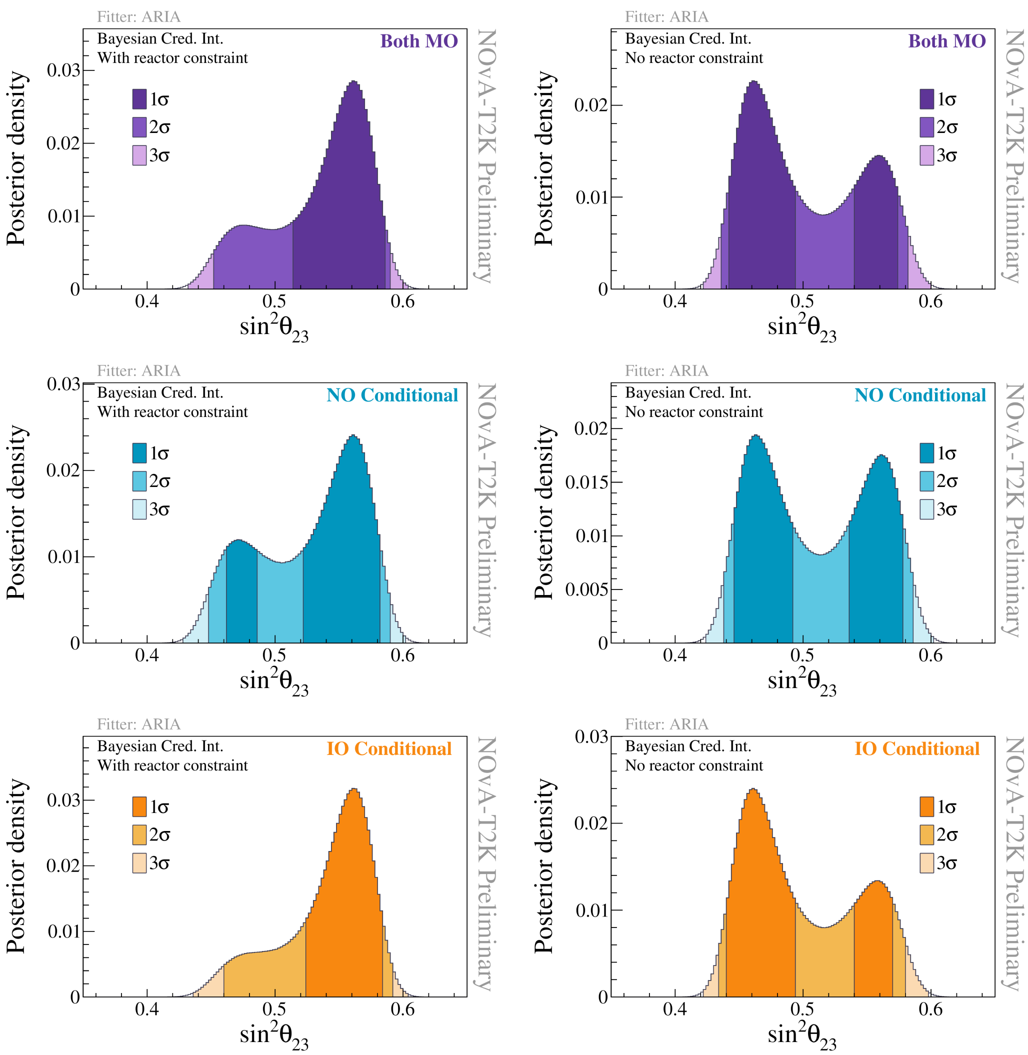

Figure 4.20: 1D posterior distributions of $\sin^2\theta_{23}$. Posteriors are made using the Aria fitter, over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

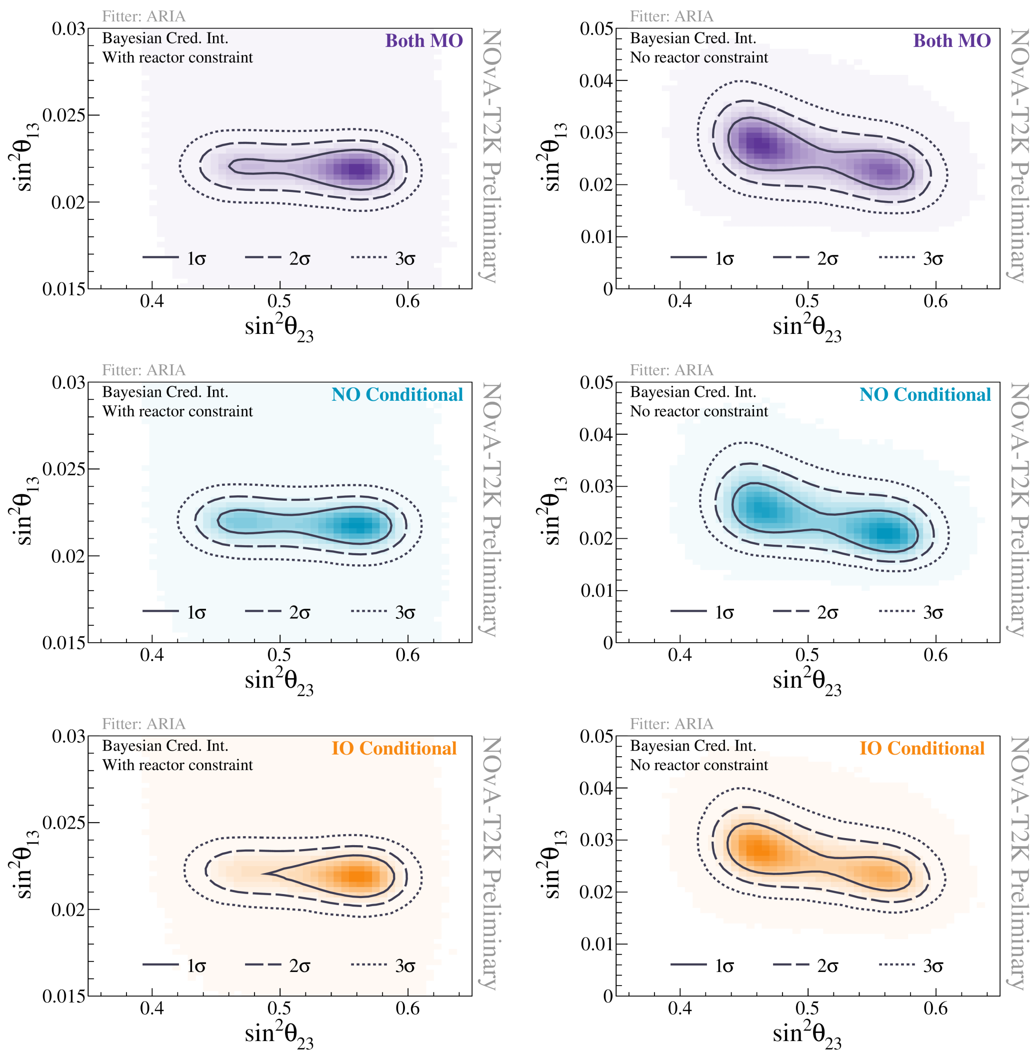

Figure 4.21: 2D posterior distributions of $\sin^2\theta_{23}$ vs $\sin^2\theta_{13}$. Posteriors are made over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint. Note the change in y-axis when adding the reactor constraint.

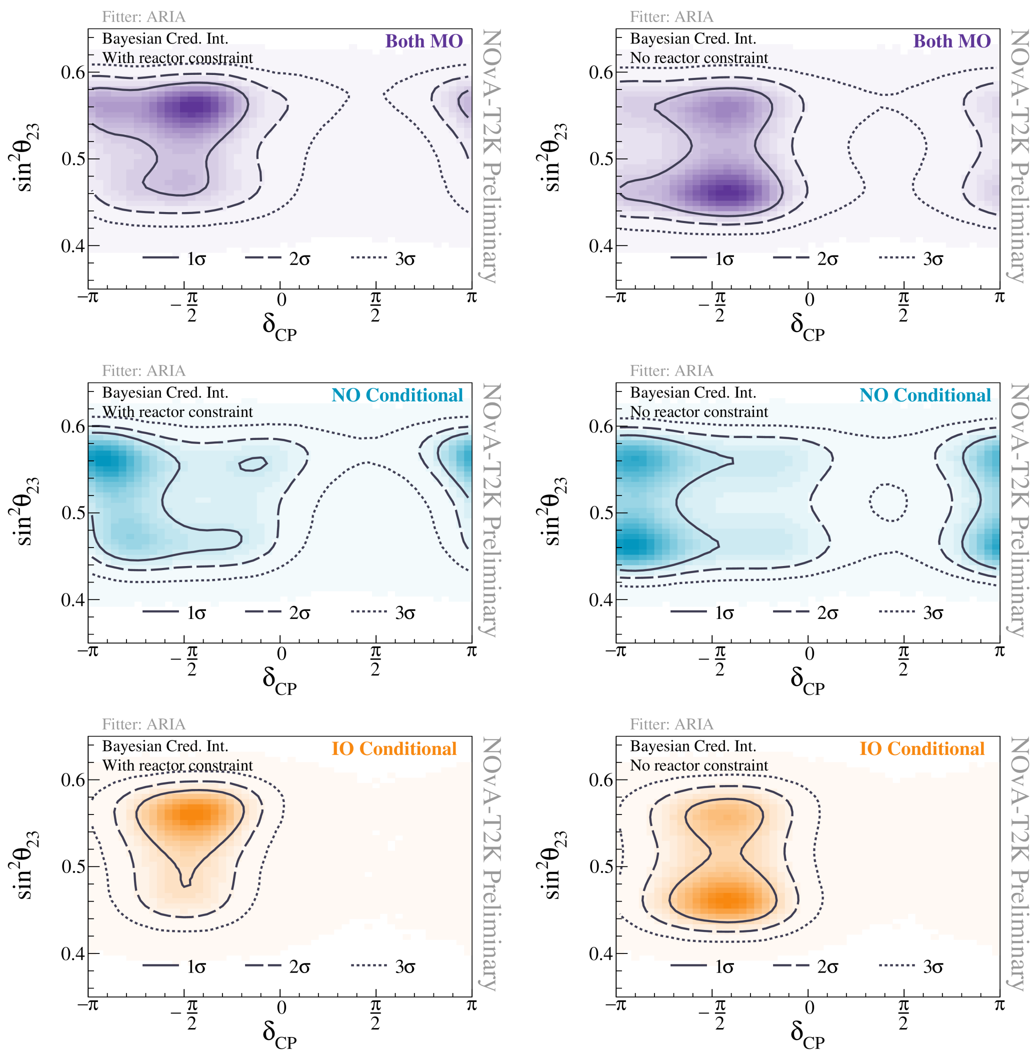

Figure 4.22: 2D posterior distributions of δ_{CP} vs $\sin^2\theta_{23}$. Posteriors are made over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

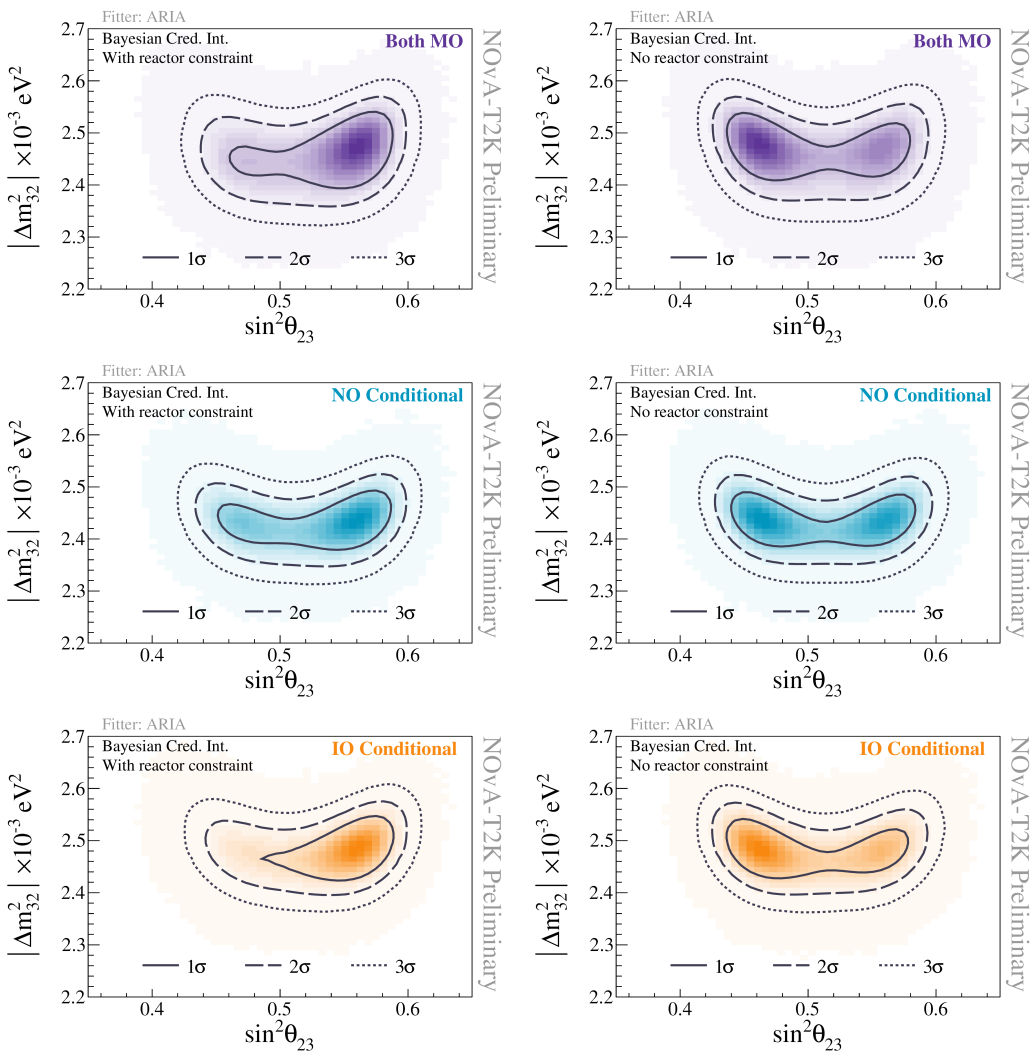

Figure 4.23: 2D posterior distributions of $\sin^2\theta_{23}$ vs $\left|\Delta m^2_{23}\right|$. Posteriors are made using the Aria fitter, over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

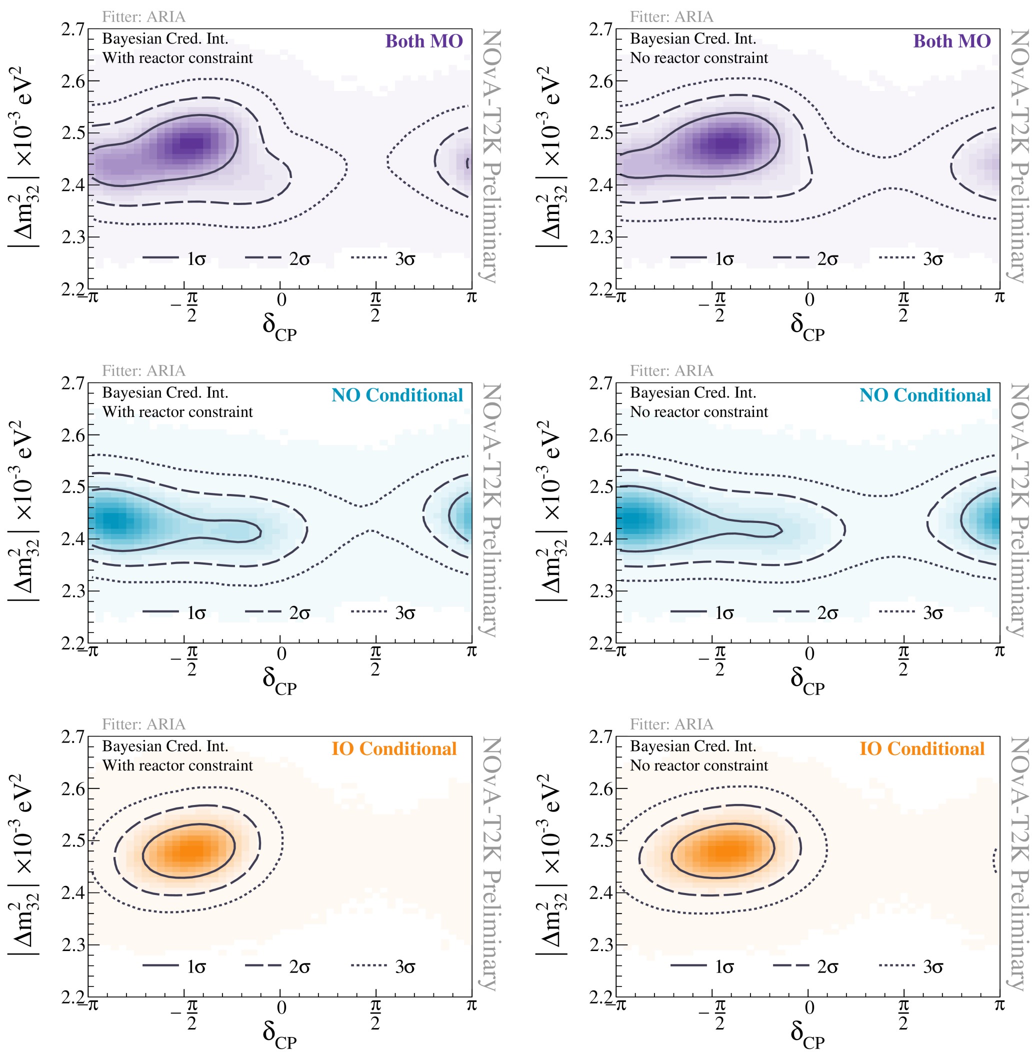

Figure 4.24: 2D posterior distributions of δ_{CP} vs $\left|\Delta m^2_{23}\right|$. Posteriors are made over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

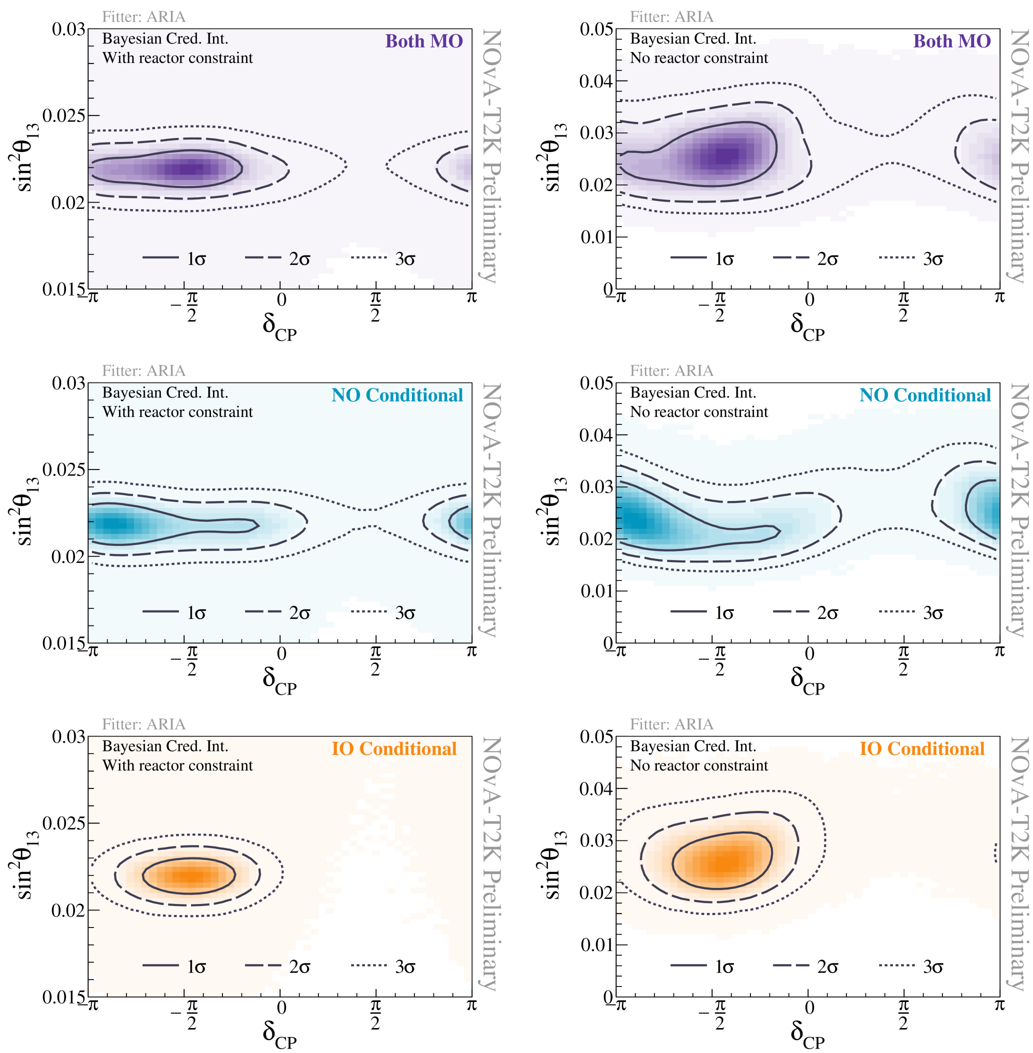

Figure 4.25: 2D posterior distributions of δ_{CP} vs $\sin^2\theta_{13}$. Posteriors are made over both mass orderings (top), normal ordering (center), and inverted ordering (bottom). The left column includes the θ_{13} reactor constraint, while the right column does not include the reactor constraint.

We can compare the above credible intervals with NOvA and T2K individually. The Figure 4.26 shows how the joint fit shows a slight preference for the IO and provides tighter bounds on Δm^2_{32} than either parent experiment. The δ_{CP} vs $\sin^2\theta_{23}$ contours (Figure 4.27) show that the contours lie at the overlap between the NOvA and T2K. Additionally, we see that the highest posterior density occurs in the inverted ordering, which is the part of the phase space where NOvA and T2K have strongest overlap.

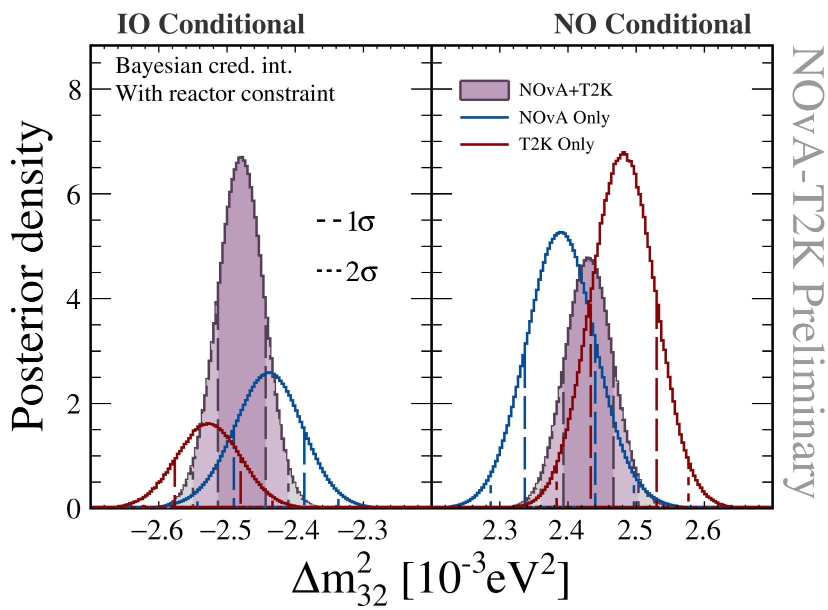

Figure 4.26: Comparison of Δm^2_{32} between the NOvA-only, T2K-only and NOvA-T2K posteriors. The joint fit shows greater posterior density in the IO, despite the parent experiments preferring the NO.

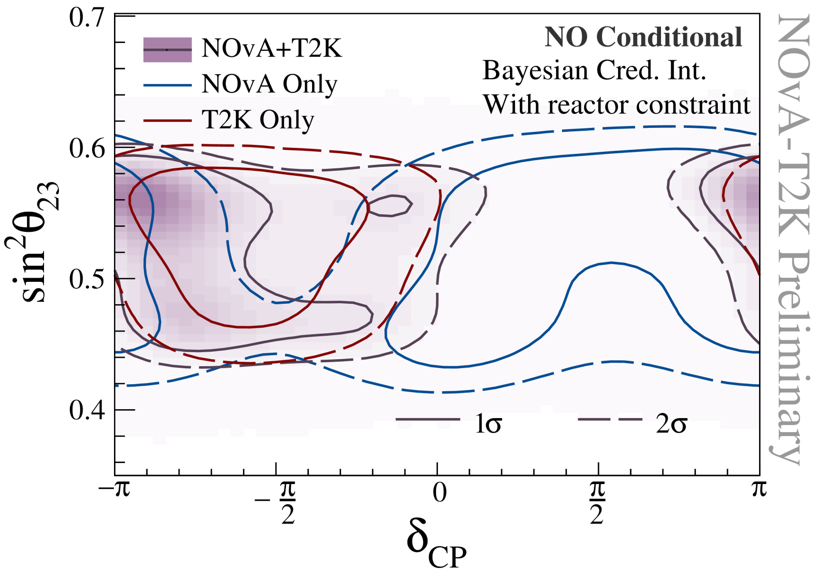

(a) Normal Ordering

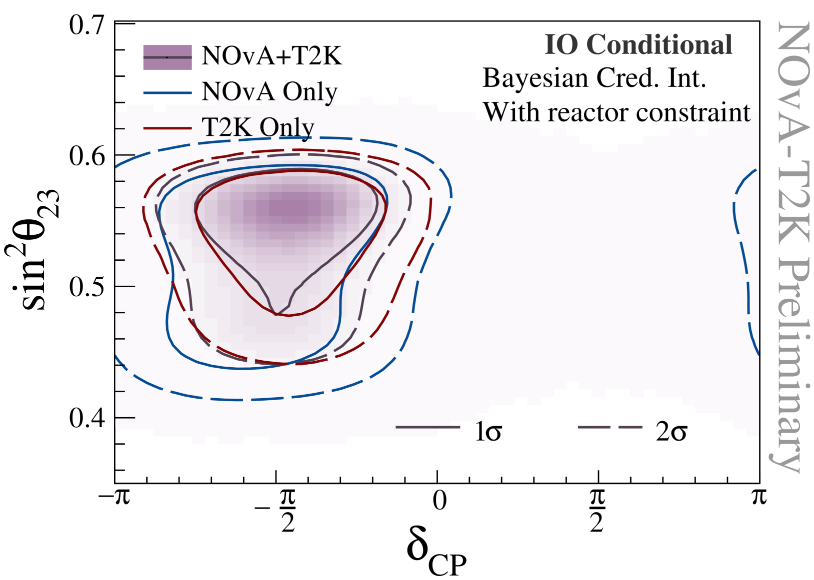

(b) Inverted Ordering

Figure 4.27: Comparison of δ_{CP} vs $\sin^2\theta_{23}$ between the NOvA-only, T2K-only and NOvA-T2K. As expected, the joint contour lies between the NOvA-only and T2K-only contours, and a slight majority of the posterior distribution lies in the inverted ordering.

Comparing across other experiments we can notice a few things:

Our δ_{CP} and θ_{23} measurements are consistent with all other experiments (Figures 4.28 and 4.29). Additionally, θ_{13} measurements are consistent with other experiments, although reactor experiments have more precision, particularly Daya Bay (Figure 4.30).

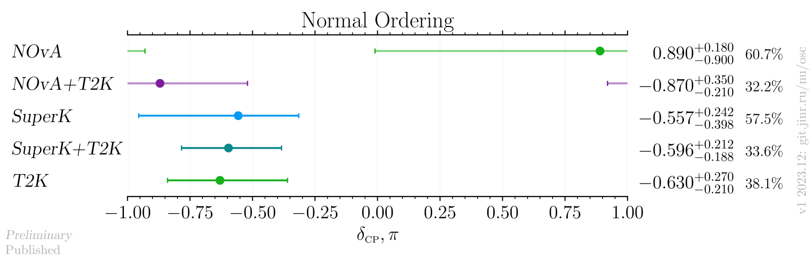

(a) Normal Ordering

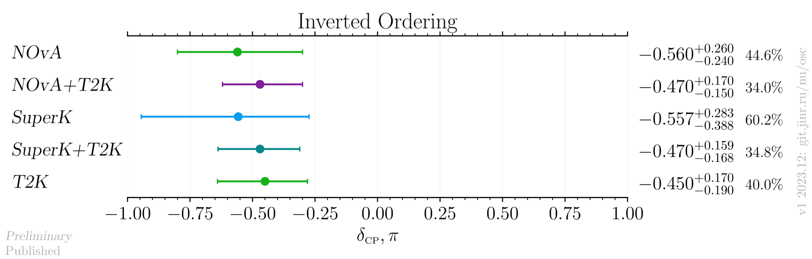

(b) Inverted Ordering

Figure 4.28: Comparison of δ_{CP} measurement with other experiments.

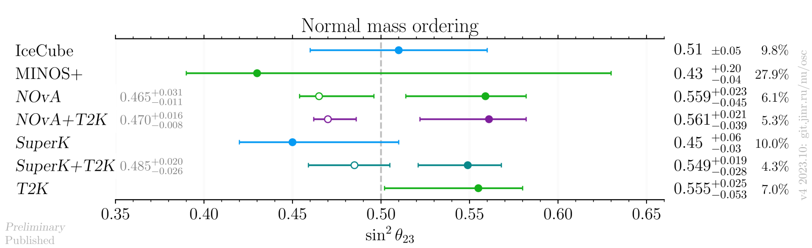

(a) Normal Ordering

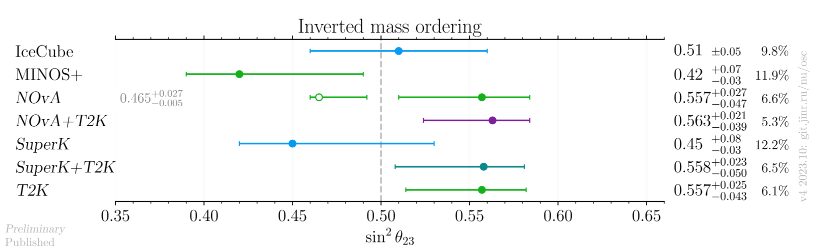

(b) Inverted Ordering

Figure 4.29: Comparison of θ_{23} measurement with other experiments.

Finally, we see that the NOvA-T2K joint fit has provided the tightest constraint on $|\Delta m^2_{23}|$ to date (Figure 4.31).

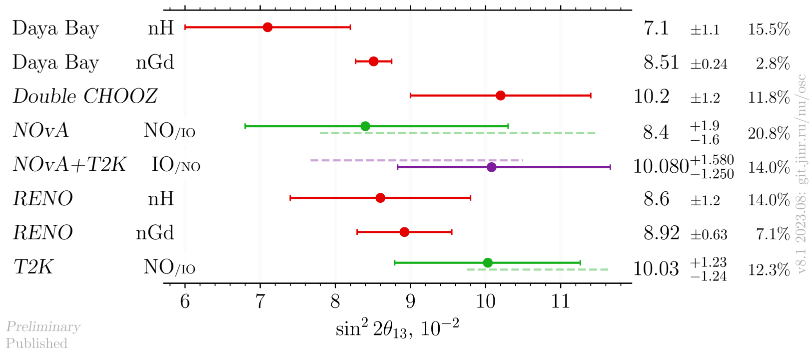

Figure 4.30: Comparison of $\sin^2\theta_{13}$ measurement with other experiments.

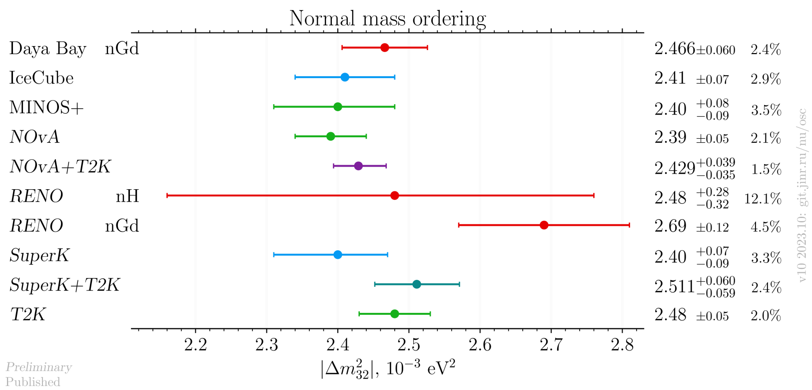

(a) Normal

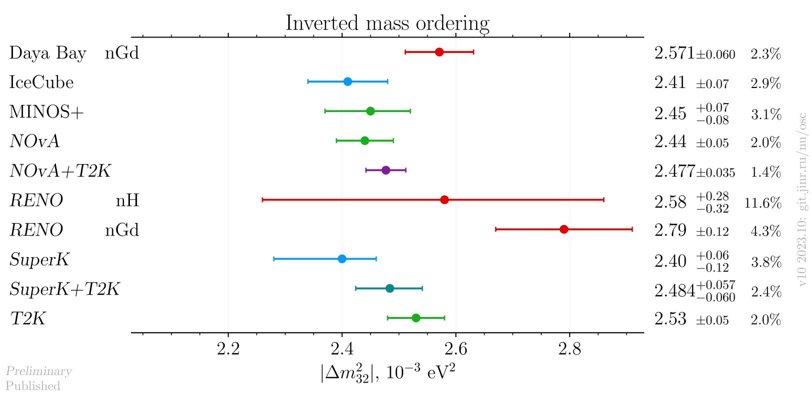

(b) Inverted

Figure 4.31: Comparison of Δm^2_{32} measurement with other experiments.

4.12 2D Reactor Constraints

The high precision we have achieved in measuring $|\Delta m^2_{32}|$ can lead us towards a better way to interface with reactor neutrino experiments in determining the mass ordering. As Nunokawa et al. show, reactor neutrino and accelerator neutrinos are expected to show slight tension in $|\Delta m^2_{32}|$ when assuming the incorrect mass ordering and more agreement in the correct mass ordering [110]. Looking at a comparison between

Daya Bay and the NOvA-T2K result 4.32, we see that reactor neutrino experiments and NOvA-T2K differ more in the inverted than in the normal ordering. This would suggest that if we switched to a 2D reactor constraint by including Daya Bay's $|\Delta m^2_{32}|$ measurements, we might see a preference for the normal ordering vs inverted ordering. Indeed, we see that the ordering preference flips back to the normal ordering when we use a 2D reactor constraint (Figure 4.33). This preference is still statistically insignificant.

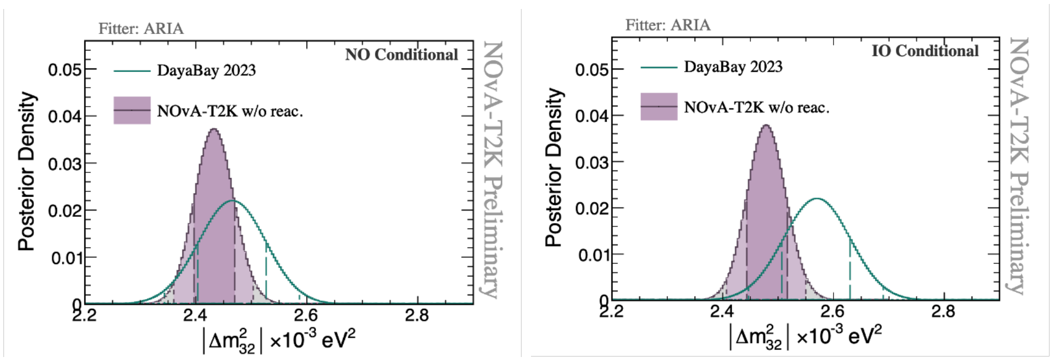

(a) Normal Ordering (b) Inverted Ordering

Figure 4.32: Comparison of Δm^2_{32} posteriors for NOvA vs Daya Bay. The agreement between NOvA+T2K and Daya Bay is stronger in the NO than in the IO.

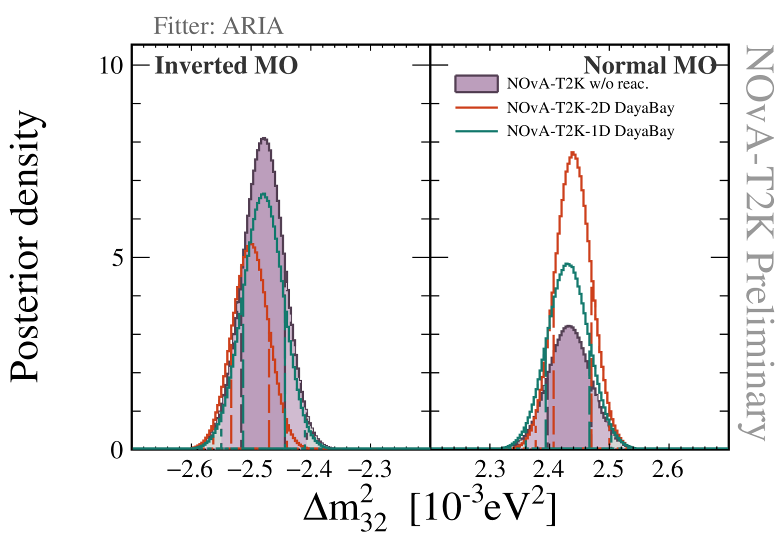

Figure 4.33: Comparison of Δm^2_{32} posteriors without reactor constraint (purple, shaded), with a θ_{13} reactor constraint (green), and with a 2D reactor constraint (red). Note that using the 2D constraint flips the joint fit's mass ordering preference to NO.

4.13 Summary of Results

All told, we report the following exciting results:

- Measurement of Δm^2_{32} at world-leading precision.
- Exclusion of $\delta_{CP} = +\frac{\pi}{2}$ from 3σ credible intervals, with no assumptions about other oscilation parameters.
- A slight preference for the inverted ordering.
- Assuming the IO, evidence of CP violation at 3σ.

These results are summarized in Tables 4.7 and 4.8.

	Preference	Bayes Factor
Mass Ordering	Inverted Ordering	1.3
θ_{23} Octant	Upper Octant	3.5

Table 4.7: Bayes factors for the joint fit's preferences for the neutrino mass ordering and the octant of θ_{23}. Neither Bayes factor is significant, so our results indicate only slight preferences for the Inverted Ordering and the Upper Octant.

Parameter	Value (NO Conditional)	Value (IO Conditional)
δ_{CP} $[\pi]$	$-0.870^{+0.350}_{-0.210}$	$-0.470^{+0.170}_{-0.150}$
$\sin^2\theta_{23}$	$[0.462, 0.486] \cup [0.522, 0.582]$	$0.563^{+0.021}_{-0.039}$
Δm^2_{32} $\left[10^{-3}\text{eV}^2\right]$	$2.429^{+0.039}_{-0.035}$	-2.477 ± 0.035

Table 4.8: Measured values from the NOvA-T2K joint fit of δ_{CP}, $\sin^2\theta_{23}$, and Δm^2_{32}.

Chapter 5

GENIE IN A BOTTLE

One of the great challenges of the joint fit, as explained in Section 4.6 is the different parametrization of cross-section models. Table 4.1 provides a compelling depiction of how different the cross-section parametrization is between NOvA and T2K.

5.1 Motivation

It is not possible to directly port GENIE's knobs into NEUT. The fake data studies described in Section 4.10 describe a way to reweight data to create fake data that resembles the central value tune of the other experiment. This does not allow one to actually adopt the other experiment's cross-section model parametrization, as fits still take place in the original parametrization used by the experiment.

This chapter will describe a method that may make it possible for future joint fits to efficiently correlate systematic parameters between two experiments, regardless of how disparate the underlying model parametrizations are.

5.2 Premise

Consider an analysis of a single long baseline neutrino experiment. The basic premise behind a fit is that systematic errors on any simulations or models used are treated as nuisance parameters in service of oscillation fits. To do this, experiments will typically make a tunable simulation which will be fit to data.

Simulation begins by creating neutrino interactions with detector material using a neutrino event generator like GENIE or NEUT. The resultant daughter particles propagate through the detector material in physics simulation software like GEANT4, depositing energy. The MC then uses a detector simulation to produce simulated output for the event.

These simulated readouts can be fed into the usual selection and reconstruction process, allowing us to get MC spectra in reconstructed variables that can be compared directly with data.

Almost every step in the above process differs between two experiments like NOvA and T2K. Specifically, the neutrino event generators used, GENIE and NEUT respectively, use completely different parametrizations to each other. There is no easy translation between the two generators. Additionally, the fact that NOvA and T2K have such radically different processes mean that all downstream components of the simulation pipeline are necessarily different to each other too. Given how tightly integrated the tuning of event generators is to each experiment's analysis, it is infeasible for both experiments to share a single framework.

However, there is one transient part of the process which is necessarily common between the two experiments: the outputs from the event generators. These outputs consist of truth information about the simulated events' energies, momenta and topologies. These simulated events (henceforth described as "true events") are a common platform that can be used to connect, and possibly correlate the two models.

5.3 Truth Bin Systematics

Shifting a cross-section systematic parameter will ultimately adjust the rate of events as a function of the event's true final state variables. The Truth Bin Systematic is a construction to take advantage of this truth information. By binning up this final state space, we can let each bin correspond to a Truth Bin Systematic. Shifting this systematic parameter will correspond to increasing or decreasing the weight of events in its bin of phase space.

A sufficiently dense binning would provide a similar degree of control over true spectra as the underlying cross-section model. However, we obviously would have removed all the physics by doing this. However, we can bring this physics back by introducing correlations between Truth Bin Systematics. By calculating the correlated effects of the cross-section model on truth bins, we can effectively encode the uncertainties of the model into a covariance matrix on truth bins. In essence, we have put "GENIE in a Bottle".

Figure 5.1 walks through a toy example. Consider some spectrum we need to fit to data. Systematic parameters will adjust this spectrum according to the underlying cross-section model.

Since we fit histograms and not continuous spectra, we can consider having to fit one histogram to another. The idea behind Truth Bin Systematics is to translate the action of a single model-based systematic parameter into the action of correlated shifts to each bin of the underlying histogram in truth space.

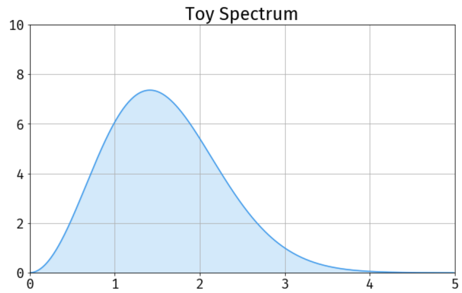

(a) A toy spectrum

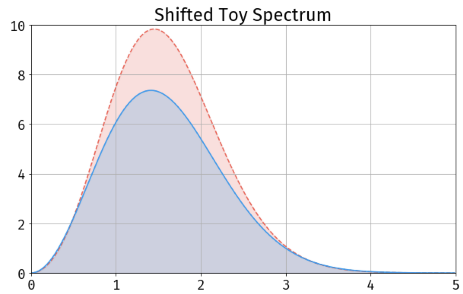

(b) Shifted spectrum

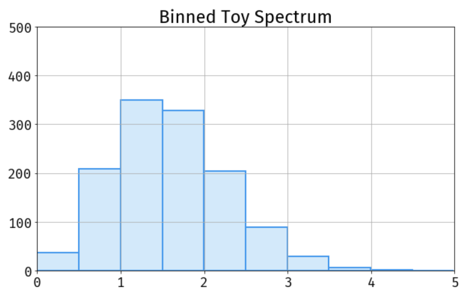

(c) Toy histogram

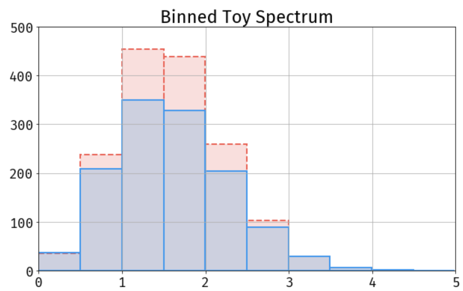

(d) Shifted histogram

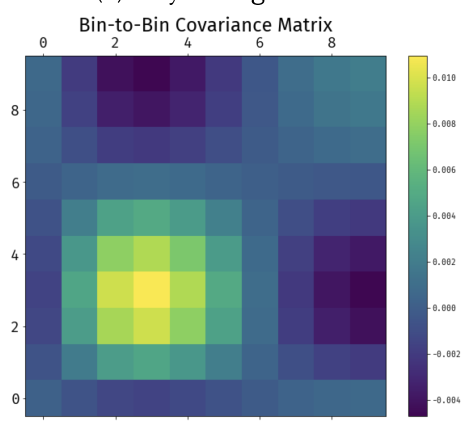

(e) Covariance Matrix across truth bins

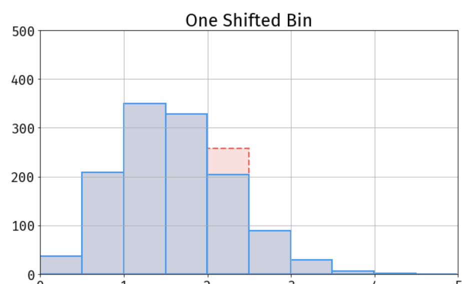

(f) Shifting an individual bin as a Truth Bin Systematic

Figure 5.1: A walkthrough of a toy model of Truth Bin Systematics. in Figure 5.1a, we see a sample true spectrum of some physics parameter which can be adjusted according to a systematic dial (Figure 5.1b). In practice, we bin this spectrum (Figure 5.1c). By seeing how shifts to the systematic parameter affect histograms in this truth variable (Figure 5.1d), we can construct a covariance matrix (Figure 5.1e). The idea behind Truth Bin Systematics is to shift each bin separately (Figure 5.1f), subject to correlations enforced by the covariance matrix. Note that this toy model never leaves the realm of truth variables. In an actual experiment, while Truth Bin Systematics still act on truth information, our final spectra will be in reconstructed variables.

Of course, the individual Truth Bin Systematics should not be treated as independent from one another. To figure this out, we can calculate the way the underlying model shifts correlate individual bins. Methods for doing this are discussed in Section 5.5. An example covariance matrix is depicted in Figure 5.1e.

5.4 Fitting

When fitting to data, cross-section parameters are typically included as part of the set of nuisance parameters. This means that the fitter will adjust these parameters in order to optimize the χ^2 of the data vs MC at a given set of oscillation parameters. Because the cross-section models encode empirical physics, we should include this prior knowledge when fitting these parameters. This can be done via the inclusion of penalty terms to the χ^2 calculation:

$$\chi^2_{\text{penalty}} = \sum_{i \in \text{systs}} z_i^2$$

where z_i represents the fractional shift away from the central value of the ith systematic parameter, expressed in units of σ_i (the standard deviation of systematic i). Adding this term to the nominal χ^2 incentivizes the fitter to keep model parameters close to their central values

The above formulation assumes no correlations between systematics. This assumption is easily adjusted to include correlations, as we need for our Truth Bin Systematics:

$$\chi^2_{\text{penalty}} = \sum_{i,j \in \text{systs}} z_i \left(\rho^{-1}\right)_{ij} z_j$$

where ρ is the correlation matrix across all systematics.

If we start with S uncorrelated physics-based systematics, and convert these into T Truth Bin Systematics, we effectively change the nuisance parameter space from an S- to a T-dimensional space. We need to insist that $T > S$ so that we can be sure that the Truth Bin Systematics are capable of capturing the same number of degrees of freedom as the original suite of systematics. In practice, we can set $T \gg S$, in order to improve the chances that the T Truth Bin Systematics can capture the potentially complex behavior of the original systematic model.

The fact that we would be encoding S degrees of freedom into the T Truth Bin Systematics suggests that there should be some degenerate degrees of freedom. This would lead to the correlation matrix being singular. In order to guarantee that an inverse is calculable, we will boost the diagonal elements of the covariance matrix a small amount. In essence, we are slightly (but negligibly) increasing the variances of individual truth bins.

5.5 Calculating Covariance Matrices

In order to calculate a covariance matrix, we have tried two methods: Monte Carlo Estimation and Linear Approximation, each with pros and cons.

Monte Carlo Estimation

The conceptually simpler way to estimate the covariance matrix between Truth Bins is to calculate it via Monte Carlo sampling. We throw the underlying systematic model parameters randomly millions of times, and use the resultant shifts in truth bins to estimate the covariance matrix U:

$$U_{ij} = \mathrm{E}[t_i t_j] - \mathrm{E}[t_i]\mathrm{E}[t_j]$$

where t_i is the fractional shift on the ith Truth Bin and the expectation values $\mathrm{E}[\cdots]$ are calculated over the throws of the original systematic parameters.

One final subtlety: due to nonlinearities in the systematic model, it is possible that the average shift to a Truth Bin is nonzero. As such, when calculating the covariance matrix, care must be taken to also save the average shift for each bin $E[t_i]$. This matters when it comes time to calculate the penalty term, as the "central value" for Truth Bin Systematic i should correspond to $E[t_i]$.

Linear Approximation

This approach is based on the propagation of errors [29]. To first order, we can approximate the covariance matrix between truth bins as follows:

$$U_{ij} = \sum_{k,l} \left.\frac{\partial t_i}{\partial s_k}\frac{\partial t_j}{\partial s_l}\right|_{\mathrm{CV}} V_{kl}$$

where t_i represents the relative shift on the ith truth bin, s_k represents the amount of change applied to the kth systematic, the partial derivatives are evaluated at the central value of the systematic parameters, and $V_{kl} = \mathrm{cov}[s_k, s_l]$ is the covariance matrix of the original systematic model parameters.

In matrix notation we can rewrite the above as:

$$U = A^T V A$$

where $A_{ij} = \frac{\partial t_j}{\partial s_i}\Big|_{\mathrm{CV}}$.

This means that if we can calculate $\frac{\partial t_i}{\partial s_k}$, the rate of change of truth bin i under changes of the underlying systematic k, we can calculate the covariance matrix on Truth Bin Systematics.

In practice, we take a linear approximation by shifting the underlying systematic to two points, and calculating the difference between the resultant shifts to the truth bin. Unlike in the Monte Carlo Estimation, we assume the Truth Bin Systematic shifts are centered at zero. One caveat is that by construction, this method cannot accurately handle systematics that are severely nonlinear.

One final note on this method is that it is in principle reversible: starting from a Truth Bin Covariance matrix U_{ij}, it should be possible to find a covariance matrix V_{kl} such that $U = A^T V A$. The upshot is that the original Truth Bin Covariance Matrix could instead have been created from a different model:

$$V_1 \xrightarrow{U=A_1^T V_1 A_1} U \xrightarrow{U=A_2^T V_2 A_2} V_2$$

The power here is that it is in principle possible to use a truth binning to translate a tuning from one interaction model to another, like NEUT into GENIE and vice versa.

5.6 Applications of this Technique

One potential end goal for this method is to make cross-section tunes portable from one neutrino event generator to another. This could also provide an apples-to-apples comparison between two otherwise unrelated parametrizations of the neutrino cross-section model.

Future joint fits between NOvA and T2K, for example, could use Truth Bin Systematics to examine and implement correlations between the two cross-section parametrizations.

While the joint fit represents a promising future application, we show in this thesis that Truth Bin Systematics can be successfully applied for the existing cross-section parametrization in a single experiment. As such, we will implement Truth Bin Systematics for NOvA alone.

5.7 Choosing a Truth Binning for NOvA

To replace NOvA's systematics with a set of Truth Bin Systematics we need to ensure that we are capturing as much physics as is relevant to the NOvA cross-section model.

As such, we cannot get away with one kinematic variable like in the toy example. Instead we need a binning in several variables, both kinematic and topological. Ideally, this choice of variables should provide the suite of Truth Bin Systematics with enough degrees of freedom to encapsulate the underlying cross-section model variations.

Kinematic Variables

Consider the generic neutrino scattering process (Figure 5.2):

$$\nu_\ell + n \rightarrow \ell^- + \text{(hadronic system)}$$

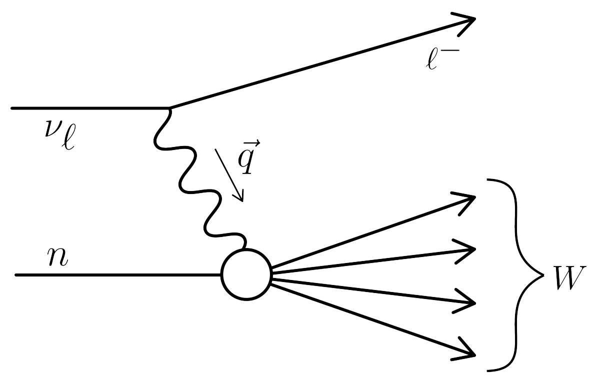

Figure 5.2: Simple neutrino scattering process. The primary kinematic variables beyond the incident neutrino energy E_ν, are the four-momentum transfer $\vec{q}$ from the neutrino into the nucleon, and the invariant mass of the hadronic system W. Note that the components of q^μ contain the information needed to calculate W. Other calculable variables of interest include E_{had}, which is the energy of the hadronic system (accessible via q^0) and the lepton transverse momentum p_T (accessible via the spatial components of $\vec{q}$).

In NOvA, our primary observables are the travel direction of and energy deposited by the charged lepton, and the energy deposited by the hadronic system. Because we can assume that the incident neutrino is coming from the beam source, and assuming that the nucleon is at rest before the interaction, we can use the observables to derive many others: the incident neutrino energy E_ν, the four-momentum transfer into the hadronic system $\vec{q}$, the lepton transverse momentum p_T, etc.

In fact, the kinematics of this process is uniquely determined by E_ν and $\vec{q}$. Going further, it can be shown that only two degrees of freedom of $\vec{q}$ are sufficient to completely describe the kinematics of the interaction [16]. Therefore, there is no obvious benefit to going beyond three kinematic dimensions when setting up Truth Bin Systematics.

Topological Categories

We want to encapsulate many different types of interactions via truth bins.

First, we consider the flavor of the incoming neutrino as well as the charge of the weak current in the scattering process. This yields five categories: ν_μ CC, $\overline{\nu}_\mu$ CC, ν_e CC, $\overline{\nu}_e$ CC, and NC.

Next, we need a proxy for interaction types. Because different cross-section model parametrizations demarcate interaction types differently, we find a different topological marker. The number of pions produced from the interaction provides such a marker. Thus, we can divide the sample into three buckets:

- 0 π: Mostly Quasielastic (QE), some 2-particle-2-hole (2p2h) events
- 1 π: Mostly Resonant (RES), some Deep Inelastic Scattering (DIS) events
- 2+ π: Mostly DIS

Final-state interactions within the nucleus can cause events of various interaction types to appear in other pion-count categories (e.g. some RES events are in the 0 π sample, while some QE events are in the 1 π sample).

Final Bin Totals

In summary, we will be considering the following space of truth bins:

$$\left\{E_\nu, Q^2, q^0\right\} \times \{\nu \text{ Flavor}\} \times \{\pi \text{ Count}\}\,.$$

If we were to use standard rectangular bins the kinematic space with $O(n)$ divisions per axis, then we would have as many as $O(n^3)$ bins. Since each truth bin corresponds to a nuisance parameter we need to fit over, computationally, we cannot afford to have n be much greater than 10 ($>$ 1000 systematics per topological category). However, low values of n mean that each truth bin covers a very large region of phase space, reducing the flexibility of the final Truth Bin Systematics.

The problem of finding a suitable kinematic binning required the development of a new analysis tool: the Voronoi histogram. Chapter 6 will go into greater detail on the method. For now, suffice it to say that Voronoi histograms make it possible to create usable binnings over data of any dimensionality with any desired number of bins.

If we choose Voronoi histograms with 100 bins for each topological category, a full fit would involve 1,500 systematic parameters.

5.8 Closure Test: ν_μ CC Events

We want to verify that Truth Bin Systematics are capable of capturing the behavior of the original cross-section systematic parameters. To start, let us consider a smaller set of events: ν_μ CC only.

We can use the observed distribution of ν_μ CC events in each of the three pion-multiplicity samples to construct 100-bin Voronoi histograms. This gives us 300 truth bins. To encode the impact of the cross-section parametrization on these truth bins, we calculate the covariance matrix (Figure 5.3).

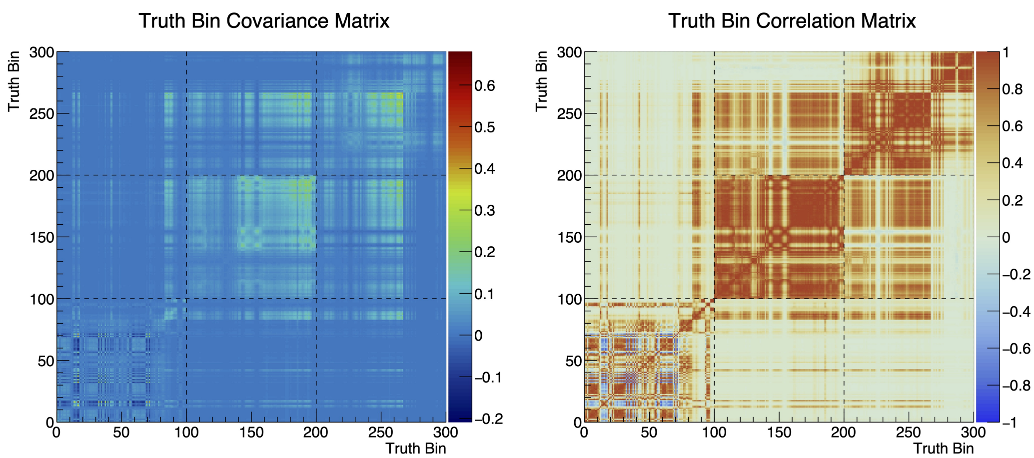

Figure 5.3: The bin-to-bin covariance (left) and correlation (right) matrices for ν_μ CC events in 300 truth bins. The dashed black lines separate the 300 total bins into 100 each with 0 π, 1 π, and 2+ π respectively. Note that each bin represents a region of 3D phase space; the relative ordering of bins within each pion-count sample is not important. The correlation matrix is calculated by throwing NOvA's cross-section parameters millions of times, and seeing how the spectra changes in truth bins.

With this 300×300 covariance matrix, we can construct our Truth Bin Systematics. If the Truth Bin Systematics accurately mimic the original cross-section systematic parameters, then we should expect that random throws of Truth Bin Systematics should look similar to random throws of cross-section systematic parameters.

We can assess if this is true by seeing how random throws of Truth Bin Systematics shift spectra in reconstructed phase space. We can construct a bin-to-bin covariance matrix on these spectra.

While NOvA's cross-section model assumes no correlations between systematics, Truth Bin Systematics are by definition highly correlated. Thus, we need to ensure that "random" throws of Truth Bin Systematics obey the correlation matrix from Figure 5.3. This can be done by using the Cholesky decomposition of the inverse of the correlation matrix to weight a random vector [15].

Also, since the goal is to observe the action of systematic parameters on MC spectra, we will skip the normal oscillation and extrapolation steps, and will instead be acting on raw Far Detector MC.

The results of this closure test can be seen in Figure 5.4.

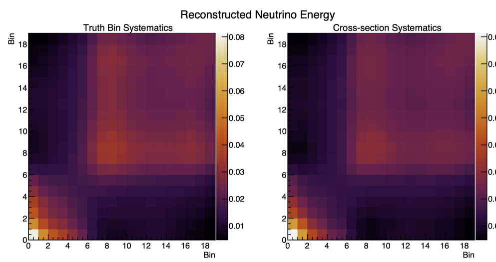

(a) Reconstructed Neutrino Energy

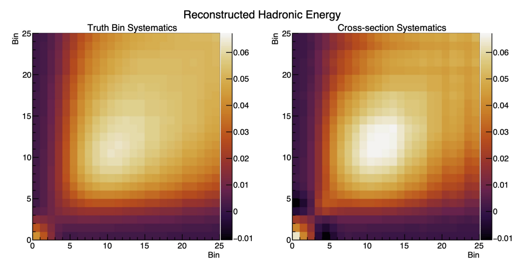

(b) Reconstructed Hadronic Energy

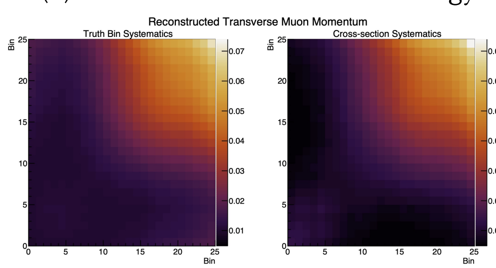

(c) Reconstructed Muon p_T

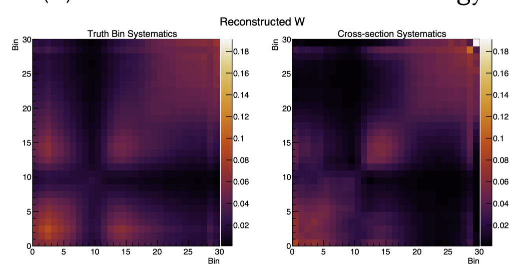

(d) Reconstructed W

Figure 5.4: Bin-to-bin covariance matrices on unextrapolated reconstructed spectra under the action of Truth Bin Systematics (left in each subfigure) and the original GENIE cross-section parametrization (right). These covariance matrices are calculated by randomly throwing the suite of systematic parameters, and calculating the covariance matrix of the shifts experienced by each bin. In all tested variables, there is remarkable agreement between both suites of systematic parameters.

There is remarkably good agreement between the covariance matrices produced by throwing the Truth Bin Systematics and the original cross-section systematics. We can look closer at this agreement by focusing on the diagonal elements of the covariance matrix. Figure 5.5 shows the standard deviation of shifts on each bin induced by both systematic suites.

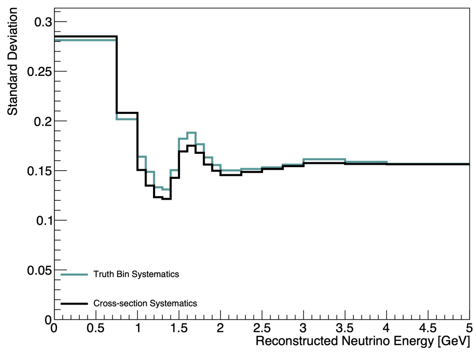

(a) Reconstructed Neutrino Energy

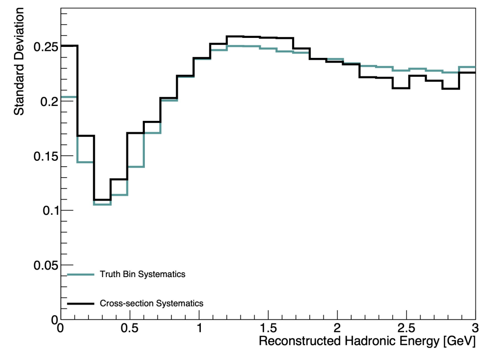

(b) Reconstructed Hadronic Energy

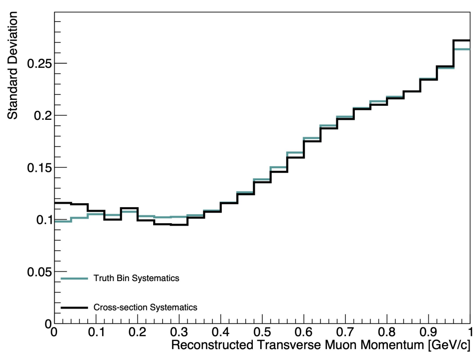

(c) Reconstructed Muon p_T

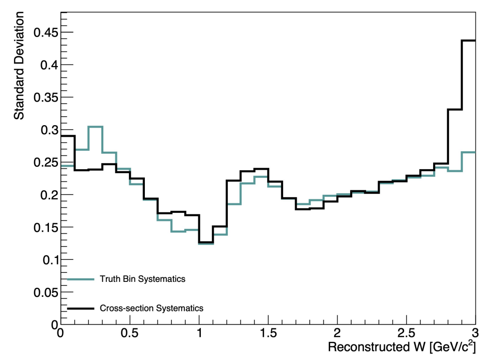

(d) Reconstructed W

Figure 5.5: For each bin of the reconstructed spectrum, we calculate the standard deviation of shifts caused by either Truth Bin Systematics (teal) or the original GENIE cross-section parametrization (black). We see good agreement between both suites of systematic parameters. Note that in these plots, the bins are mapped onto the actual reconstructed variable x-axis, instead of the bin number as shown in previous figures. These curves correspond to the square roots of the diagonal elements of the covariance matrices depicted in Figure 5.4

Again, we see excellent agreement between the actions of Truth Bin Systematics and those of the original cross-section parameters. We have thus demonstrated that Truth Bin Systematics are capable of encapsulating the NOvA cross-section model uncertainties in a model-agnostic parameter space.

Chapter 6

VORONOI HISTOGRAMS

A common problem in particle physics applications relates to dealing with binning histograms in multiple dimensions. In applications where limiting the number of bins is important, naïve approaches of using rectangular bins can lead to major inefficiencies. Indeed, in Section 5.7 we saw that we needed bins in three dimensions, but since each bin corresponded to a new systematic parameter, we desperately needed to keep the total number of bins low. In this chapter, I will introduce the concept of using Voronoi diagrams to partition space in an efficient manner for the use of histograms.

6.1 Voronoi Diagrams

Consider a space $\mathscr{P}$, in which are distributed a set of n points s_i, known as "sites". Define the "Voronoi diagram" as a partition of the space into n "Voronoi cells", one for each site. This partition is defined as follows.

Let $\mathrm{dist}(a, b)$ be a distance function between two points a, b in $\mathscr{P}$. Then we define the Voronoi cell for site s_i as the set of points $p \in \mathscr{P}$ such that $\mathrm{dist}(p, s_i) \leq \mathrm{dist}(p, s_j)$ for all $j \neq i$.

A simple example of such a diagram can be seen in Figure 6.1.

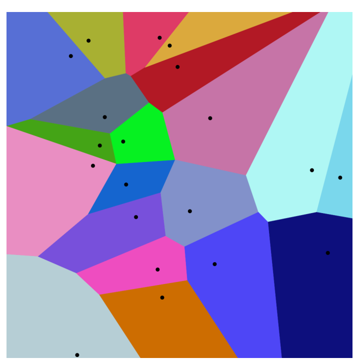

Figure 6.1: A simple example of a Voronoi diagram. Each black dot is a site, and the surrounding colored region denotes its Voronoi cell. For example, consider the blue region at the top left of the image, and the site contained therein. Each point in this blue region is closer to this site than it is to any other site. File from Wikimedia Commons [111].

Put another way, the Voronoi diagram partitions space into cells based on which of the n sites is closest.

History

Given that the idea behind Voronoi Diagrams is naturally motivated, the concept has been discovered and rediscovered numerous times. Figure 6.2 highlights a few such early uses before mathematical formalism as described above was described by Georgy Voronoy.

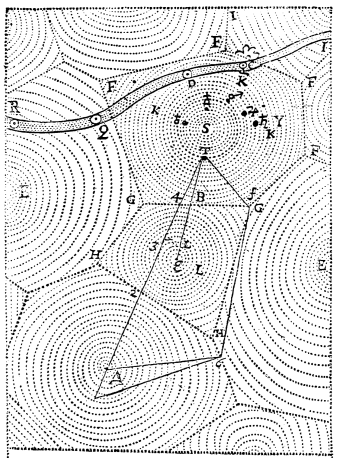

(a) From René Descartes's *Le Monde* (1664)

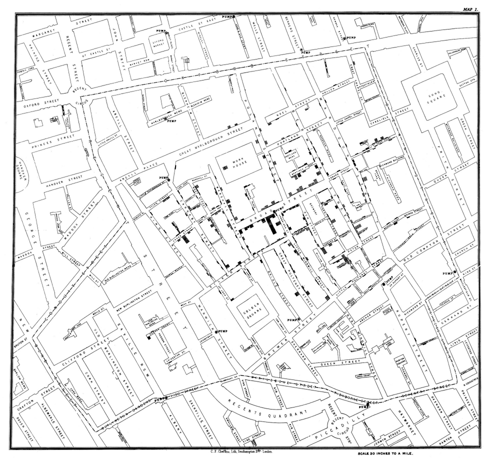

(b) From John Snow's *On the mode of communication of cholera* (1854)

Figure 6.2: (a) An early version of a Voronoi diagram found in René Descartes's 1664 treatise on science *Le Monde*. The "Suns" are denoted by labels S, ε, E and A in the above figure [112]. (b) John Snow's 1854 map of cholera cases relative to the locations of various water pumps in the Soho neighborhood of London. Black bars indicate fatal cholera cases surrounding the water pump on Broad Street [113].

Descartes's 1664 treatise *Le Monde* describes an attempt to divide the night sky into several regions, each of which consists of a "heaven turning about the Sun" [112]. This effectively partitions space into a weighted Voronoi Diagram where stars are sites. It is unclear whether this visualization technique originated with Descartes or was simply a contemporary way of describing the night sky [114].

One of the first practical applications came in 1854 at the dawn of modern epidemiology. The physician John Snow charted the locations of fatal cholera cases in the Soho neighborhood of London, and found that: "nearly all the deaths had taken a short distance from the pump [on Broad Street]. There were only ten deaths decidedly nearer to another street pump." [113]. Here, Snow had effectively used a Voronoi diagram where local street water pumps acted as sites, and found that almost all cholera cases in the outbreak were within the Voronoi Cell corresponding to the Broad Street pump. Snow's hand drawn map can be seen in Figure 6.2b.

Other early uses and applications of Voronoi Diagrams came in such varied fields as meteorology, geology, crystallography, ecology, demography, and information theory [114].

6.2 Capacity-Constrained Voronoi Diagrams

Consider a continuous distribution $f(x)$ over a variable x which we intend to split into N bins over x. Our goal is for each bin to contain the same integrated density (that is, the integral of the distribution over each bin contains $\frac{1}{N}$ of the total integral of the distribution). This can be done by scanning across the domain of f, creating bin boundaries as needed:

1. Create a bin boundary x_0 at the start of the domain of f.
2. Find x_1 such that the integral of $f(x)$ from x_0 to x_1 equals $\frac{1}{N}$ of the total integral. Create a new bin boundary at x_1.
3. For i from 1 to $N-1$, find x_{i+1} such that the integral of $f(x)$ from x_i to x_{i+1} equals $\frac{1}{N}$ of the total integral. Create a new bin boundary at x_{i+1}.

By construction, we see that there is exactly one set of bin-edges for 1-dimensional histograms such that each bin contains the same integrated density, or “capacity”.

In contrast, in two or more dimensions, there are infinitely many ways to partition space such that each bin of a histogram would have the same integrated capacity. A trivial example of such a “capacity-constrained” partition would be to run the above algorithm over just one of the dimensions of the underlying distribution.

To reduce the openness of this problem, I propose the following practical criteria for finding good “capacity-constrained” binnings:

- The bins should be “compact” or “round”.
- Figuring out which bin a given point lies within should run in $O(Nd)$ time with N bins in d dimensions.
- The binning should be storable with $O(Nd)$ floating-point numbers.

In general, binnings which require axis-aligned bin boundaries may fail any of the above criteria. Rather than using a grid-based partition of space as the binning, we can use a Voronoi diagram. Thus, our goal will be to find a set of Voronoi sites which partition space into a capacity-constrained binning.

In most use cases, we want bins to group together data points which are physically close to each other in Cartesian space. This naturally leads one to consider algorithms for k-means clustering as a starting point. k-means clustering is a technique often used in unsupervised machine learning applications. Lloyd's Algorithm for k-means clustering involves iteratively updating a set of Voronoi sites until a certain stopping criterion is met [115]. This means that Lloyd's algorithm provides a natural jumping-off point for creating an algorithm for choosing Voronoi sites.

Balzer (2009) lays out such an algorithm based on Lloyd's algorithm which can create a Voronoi diagram that can balance the contents of each Voronoi cell [116]. This so-called capacity-constrained Voronoi Diagram iteratively jitters sites towards a better global configuration, where the goodness-of-fit is determined by the integrated density per bin. Examples from Balzer (2009) can be seen in Figure 6.3

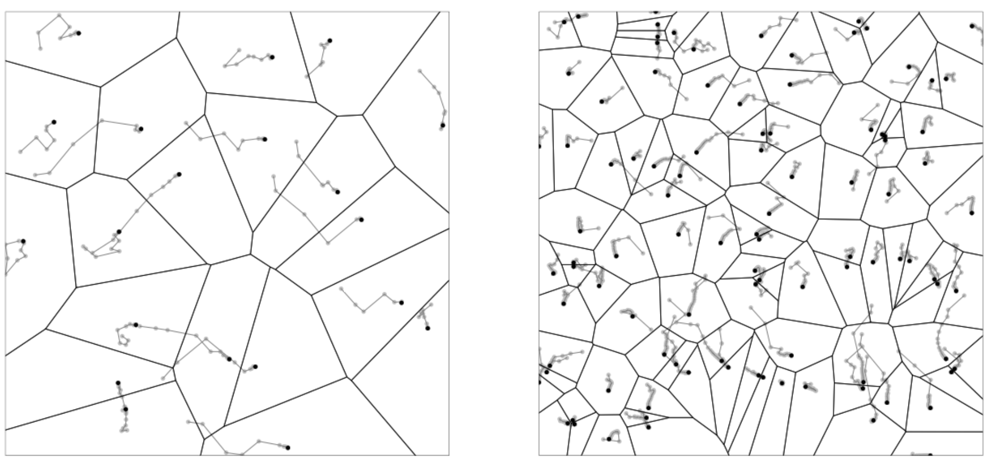

(a) 20 Sites, equal capacities

(b) 100 Sites, unequal capacities

Figure 6.3: Example from Balzer (2009) [116] showing the path individual Voronoi sites take to find a capacity-constrained configuration. In the above, the area of each cell needs to be the same for each site (left) or each site needs a different, specified capacity (right).

In practice, the underlying distributions we run on are discrete. Specifically, consider a cloud of points we intend to bin into a histogram. The Voronoi diagram will partition space such that each point lands in the cell corresponding to the closest site. As such, the capacity of a given cell is calculated by summing the weights from each point which is assigned to this cell. An important consequence of this is that it may not always be possible to create a perfectly capacity-constrained histogram, as the above constructive proof in 1D relies on the partitioned distribution being continuous.

Regardless, the algorithm can still providing useful results. Instead of planning on a perfectly capacity-constrained histogram, if we allow a stopping criterion, we can get sufficiently capacity-constrained distribution of sites that satisfy the original need behind this technique: avoiding excessive empty bins. The figure of merit used as the stopping criterion is related to the variance of capacities across all bins. A low figure of merit corresponds to all bins having similar capacities, while a high figure of merit corresponds to wide variance in bin capacities.

Finally, an important consideration is the choice of distance metric. In principle a Voronoi binning can be made with the raw Euclidean distance. In practice, however, the histogram's axes could have numerical values could be orders of magnitude apart. As a result, we rescale both axes before calculating the Voronoi binning via the Euclidean metric.

6.3 Algorithm Performance

When implemented, The Capacity-constrained Voronoi site-finding algorithm shows some sensitivity to the initial conditions. This can be seen in Figure 6.4.

In the course of preparing the NOvA analysis, we adjust 2p2h models using a histogram in $q^0 - q^3$ space. This is a good test space to demonstrate Voronoi histograms, as depicted in Figure 6.5. The standard rectangular binning is compared to a Voronoi binning with an equal number of bins, and as intended, the histograms maintain high resolutions in areas of high data density. Unrolling the Voronoi histogram into a 1D distribution shows just how much more uniform the Voronoi binning is than a standard rectangular binning.

Figure 6.6 shows that the number of bins can be greatly reduced while maintaining resolution in areas of interest. Excitingly, these gains over rectangular binnings only get more pronounced as the dimensionality increases.

6.4 Area Normalization

One subtlety heretofore ignored is that Voronoi diagrams will create cells which are unbounded in space. This is because every point in space always have at least one "closest site". This means that every point in space belongs to a Voronoi cell. Specifically, these unbounded cells will correspond to sites which lie on the convex hull of the set of sites.

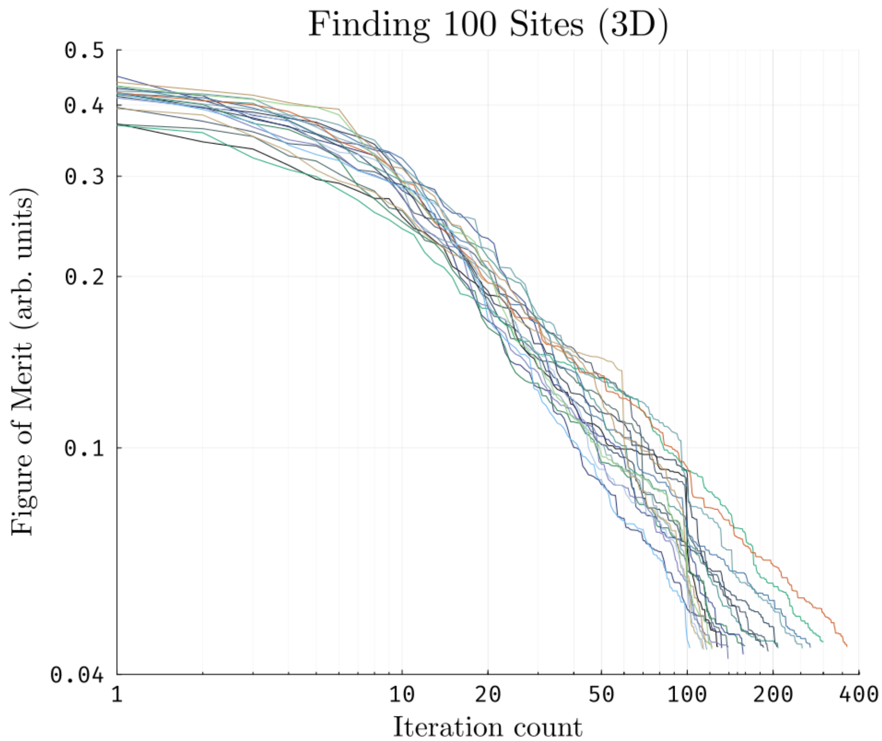

Figure 6.4: Evolution of the Voronoi site-finding algorithm on the same three-dimensional data using 20 different random seeds. The above plot is log-log in iteration count vs a figure-of-merit the algorithm optimizes. Note that despite some mitigating measures, the algorithm performance is quite sensitive to the initial distribution of sites.

The upshot is that certain functionality we expect from rectangular-binned histograms will not work the same way as Voronoi counterparts.

Most obviously is area normalization, as is used to provide a color axis in Figs. 6.5d and 6.6. The cells with infinite area would not be area-normalizable. Thus, these sites need to be clipped to finite area. The most natural way to bound a Voronoi histogram is to take a page from traditional rectangular histograms and impose a rectangular bounding box on the space described by the Voronoi histogram.

In practice, this requires effectively ray-tracing along individual edges of the Voronoi diagram to see if they intersect the edges of the bounding box. The 2D edge cases in the algorithm that needed to be handled by my implementation of the algorithm are depicted in Figure 6.7.

Once all Voronoi cells have finite area and the topology of each cell has been fixed by clipping to the bounding box, calculating cell areas is simple via the shoelace formula [117]. This allows for area normalizing all bin contents, and allows for color-mapped

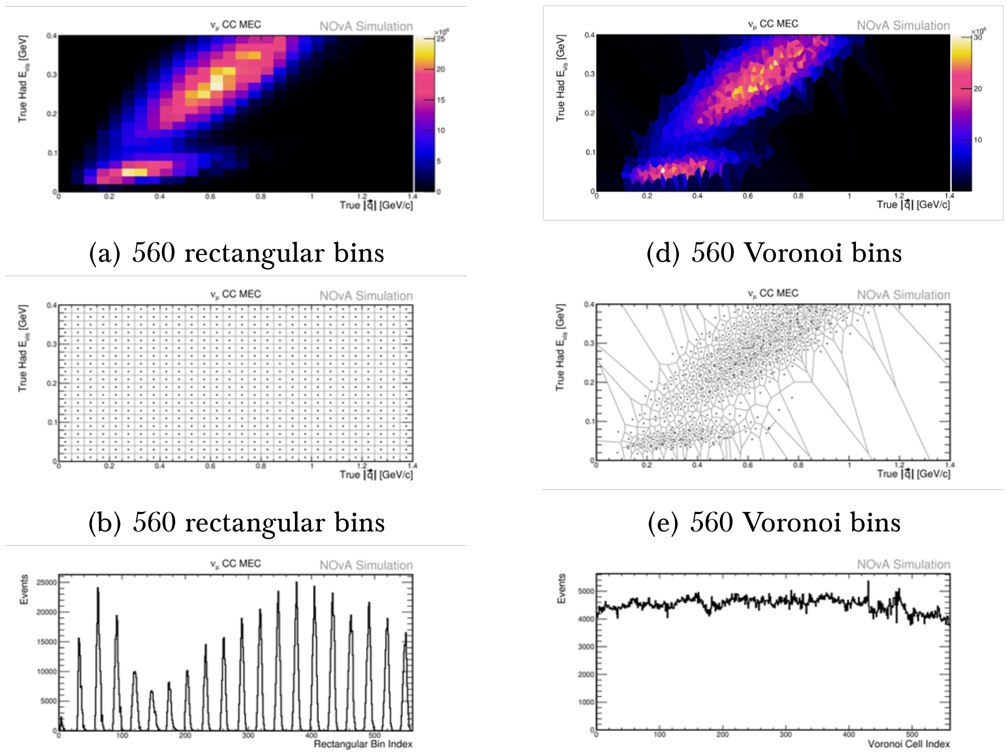

(a) 560 rectangular bins

(d) 560 Voronoi bins

(b) 560 rectangular bins

(e) 560 Voronoi bins

(c) Unwrapped rect. hist.

(f) Unwrapped Voronoi hist.

Figure 6.5: An example usage of Voronoi Histograms. The first row depicts the 2D histograms, each with 560 bins. The underlying binnings are depicted in the second row, which makes it clear that the Voronoi binning has high resolution in the region of interest, but low resolution in areas with few events. Finally, the bottom row shows an unwrapping of the top histograms into a 1D spectrum. The uniformity of the rectangular bins means that some bins will have very high capacities, while others have almost no events. In contrast, the Voronoi binning chosen for this distribution leads to a fairly uniform spectrum.

2D Voronoi histograms as above.

While the functionality of unbounded Voronoi histograms will work without modifications at higher dimensions, volume normalization becomes substantially more complicated beyond two dimensions. First of all, there are substantially more edge cases like those shown in Figure 6.7 in higher dimensions.

More pressingly, there is no equivalent of the shoelace formula to provide a closed-form expression for the volume of a polytope [118]. Instead, the state-of-the-art techniques for measuring volumes instead rely on random sampling techniques [119].

Despite these difficulties, it is important to note that higher dimensional Voronoi histograms still have uses, even if they are not area normalizable. See Section 5.7 for

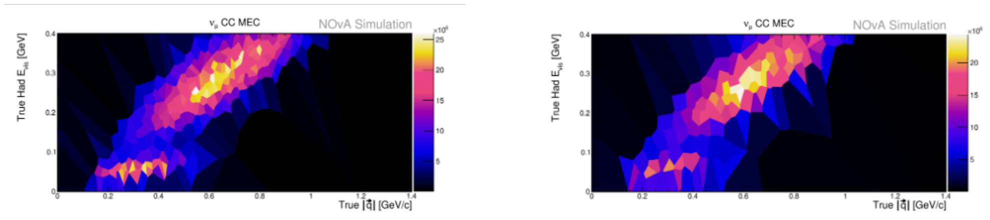

(a) 280 bins (b) 140 bins

Figure 6.6: $q^0 - q^3$ Voronoi histograms with fewer bins. With half or even one quarter the bins as the original rectangular binning (Fig. 6.5a), Voronoi histograms maintain the resolution in the regions of interest.

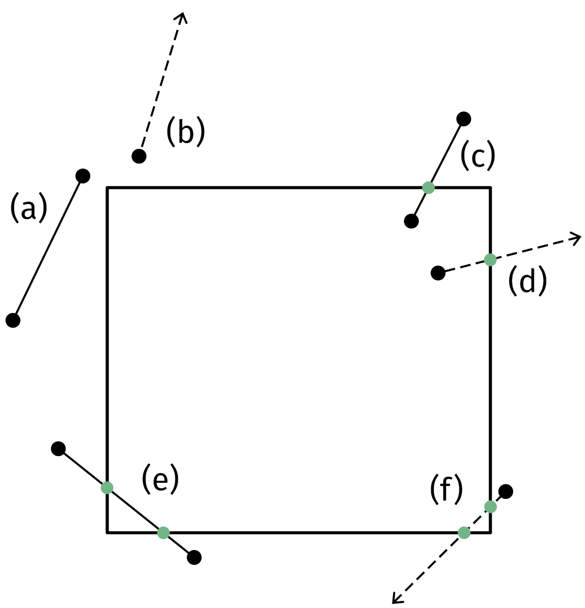

Figure 6.7: Two-dimensional edge cases for clipping a Voronoi diagram to a bounding box. Each of these edge cases will modify the number of vertices and edges for cells on the boundary of the histogram. Cases (a) and (b) are finite and infinite edges respectively which lie entirely outside the bounding box and so can be removed from the appropriate cells. Cases (c) and (d) correspond to finite/infinite edges which escape the bounding box and need to be clipped to the point of intersection with the bounding box. Finally, cases (e) and (f) involve finite/infinite edges whose vertices are outside the bounding box, but the edges partially fall within the box. These edges need to be clipped to the points where they enter/exit the box.

a situation where 3D Voronoi histograms are still useful.

6.5 Possible Extensions

Function Fitting

While the above work is useful for creating bounded histograms in two dimensions, as well as unbounded histograms in two or more dimensions, there are further uses for histograms. One example common in particle physics is function fitting, in which the integral of a function across a bin is compared against the measured content of the bin.

In traditional rectangular histograms, these integrals are often approximated by evaluating the function at the center of the bin, and taking that as the average value across the bin. This simple approximation does not work as easily for Voronoi cells as the nonrectangular shape means that it is not clear what to use as a representative point for evaluating the function. Candidates include the centroid of the cell, or the location of the site generating the cell. Either way, if the function varies significantly along the size of the cell, a simple one-point approximation will be inadequate. Instead, sampling methods will yield more accurate approximations of the integral of the function across the cell. This is true in the case of rectangular bins too, although sampling methods are simpler to implement in the case of rectangular bins.

The volume measurement algorithms in the above section depend on being able to sample points from within the polytopes. As such, if such sampling methods are implemented to area normalize these higher-dimensional Voronoi histograms, it will be simple to apply these in function fitting applications as well.

Alternative site-finding heuristics

One assumption made early in this process is that the capacity of each bin needs to be constant from bin to bin. However, this is a choice made to solve a specific problem, namely that of finding efficient binnings. It is equally easy to imagine different heuristics for optimization. One simple example is alluded to in the above section: choosing bins that don't contain large gradients of an underlying function. This would make it easier to approximate the value of the function's integral across the bin.

Chapter 7

CONCLUSION

Neutrino oscillation measurements are entering a transitional period. NOvA and T2K are both mature experiments with many years of beam data and hundreds of observed far-detector neutrinos between them. However, definitive statements on the open questions of neutrino oscillation still elude us. What is the mass ordering? Do neutrinos violate CP symmetry?

We can expect the next generation of neutrino oscillation experiments to start turning on by the end of this decade. Between DUNE, Hyper-Kamiokande, and JUNO, we will soon see major progress on these outstanding questions of neutrino oscillations.

For now, though, NOvA and T2K have successfully completed the first ever joint fit between long baseline accelerator neutrino oscillation experiments. We presented measurements of Δm^2_{32} at world-leading precision: $2.429^{+0.039}_{-0.035} \times 10^{-3}$ eV^2 in the normal mass ordering or $-2.477 \pm 0.035 \times 10^{-3}$ eV^2 in the inverted. We exclude $\delta_{CP} = +\frac{\pi}{2}$ from 3σ credible intervals. If we assume the inverted ordering (which our results give a very slight preference for), our results exclude CP conservation from 3σ credible intervals.

These excellent results were built on the hard work of both collaborations and the joint working group which bridged them. Each experiment will continue to take data, and will continue to discuss options for supporting future joint analyses [120]. Both experiments have updated data sets available, and plan to release joint analyses up until the release of the final datasets around 2027. There is always much more to learn from each other.

"GENIE in a Bottle" is a promising framework for handling cross-sections in a model-agnostic way. This technique can find use in a future joint data analysis, whether between NOvA and T2K or between future experiments.

Finally, Voronoi histograms are a very extensible idea. There are many potential applications where precise binning is needed in high-dimensional spaces.

BIBLIOGRAPHY

[1] Kai Zuber. *Neutrino Physics*. CRC Press, 2020. ISBN: 9781315195612. DOI: 10.1201/9781315195612. URL: http://dx.doi.org/10.1201/9781315195612.

[2] Enrico Fermi. "Tentativo di una Teoria Dei Raggi β". In: *Il Nuovo Cimento* 11.1 (1934), pp. 1–19. ISSN: 1827-6121. DOI: 10.1007/bf02959820. URL: http://dx.doi.org/10.1007/BF02959820.

[3] C. L. Cowan et al. "Detection of the Free Neutrino: a Confirmation". In: *Science* 124.3212 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103. eprint: https://www.science.org/doi/pdf/10.1126/science.124.3212.103. URL: https://www.science.org/doi/abs/10.1126/science.124.3212.103.

[4] Frank Close. *Half-life*. Basic Books, 2015. ISBN: 978-0465069989.

[5] B. Pontecorvo. "Mesonium and anti-mesonium". In: *Sov. Phys. JETP* 6 (1957), p. 429.

[6] G. Danby et al. "Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos". In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36. URL: https://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[7] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the Unified Model of Elementary Particles". In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. ISSN: 0033-068X. DOI: 10.1143/PTP.28.870. eprint: https://academic.oup.com/ptp/article-pdf/28/5/870/5258750/28-5-870.pdf. URL: https://doi.org/10.1143/PTP.28.870.

[8] C.V. Achar et al. "Detection of muons produced by cosmic ray neutrinos deep underground". In: *Physics Letters* 18.2 (1965), pp. 196–199. ISSN: 0031-9163. DOI: https://doi.org/10.1016/0031-9163(65)90712-2. URL: https://www.sciencedirect.com/science/article/pii/0031916365907122.

[9] F. Reines et al. "Evidence for High-Energy Cosmic-Ray Neutrino Interactions". In: *Phys. Rev. Lett.* 15 (9 Aug. 1965), pp. 429–433. DOI: 10.1103/PhysRevLett.15.429. URL: https://link.aps.org/doi/10.1103/PhysRevLett.15.429.

[10] John N. Bahcall and Raymond Davis. "Solar Neutrinos: A Scientific Puzzle". In: *Science* 191.4224 (1976), pp. 264–267. DOI: 10.1126/science.191.4224.264. eprint: https://www.science.org/doi/pdf/10.1126/

science.191.4224.264. URL: https://www.science.org/doi/abs/10.1126/science.191.4224.264.

[11] Alain Blondel. "The Number of Neutrinos and the Z Line Shape". In: *The Standard Theory of Particle Physics*. 2016. Chap. Chapter 8, pp. 145–160. DOI: 10.1142/9789814733519_0008. eprint: https://worldscientific.com/doi/pdf/10.1142/9789814733519_0008. URL: https://worldscientific.com/doi/abs/10.1142/9789814733519_0008.

[12] K. Kodama et al. "Observation of tau neutrino interactions". In: *Physics Letters B* 504.3 (2001), pp. 218–224. ISSN: 0370-2693. DOI: https://doi.org/10.1016/S0370-2693(01)00307-0. URL: https://www.sciencedirect.com/science/article/pii/S0370269301003070.

[13] B. Aharmim et al. "Electron energy spectra, fluxes, and day-night asymmetries of ^{8}B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory". In: *Physical Review C* 72.5 (Nov. 2005). ISSN: 1089-490X. DOI: 10.1103/physrevc.72.055502. URL: http://dx.doi.org/10.1103/PhysRevC.72.055502.

[14] J. Hosaka et al. "Solar neutrino measurements in Super-Kamiokande-I". In: *Phys. Rev. D* 73 (11 2006), p. 112001. DOI: 10.1103/PhysRevD.73.112001. URL: https://link.aps.org/doi/10.1103/PhysRevD.73.112001.

[15] S. Navas et al. "Review of particle physics". In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[16] Mark Thomson. *Modern particle physics*. Cambridge University Press, 2013.

[17] B. Pontecorvo. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge". In: *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

[18] C. Jarlskog. "Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation". In: *Phys. Rev. Lett.* 55 (10 1985), pp. 1039–1042. DOI: 10.1103/PhysRevLett.55.1039. URL: https://link.aps.org/doi/10.1103/PhysRevLett.55.1039.

[19] S P Mikheev and A Y Smirnov. "Resonance enhancement of oscillations in matter and solar neutrino spectroscopy". In: *Sov. J. Nucl. Phys. (Engl. Transl.); (United States)* 42:6 (Dec. 1985).

[20] L. Wolfenstein. "Neutrino oscillations in matter". In: *Phys. Rev. D* 17 (9 1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369. URL: https://link.aps.org/doi/10.1103/PhysRevD.17.2369.

[21] G. C. Branco et al. "Minimal scenarios for leptogenesis and CP violation". In: *Phys. Rev. D* 67 (7 2003), p. 073025. DOI: 10.1103/PhysRevD.67.073025. URL: https://link.aps.org/doi/10.1103/PhysRevD.67.073025.

[22] S. F. King. "Leptogenesis–Maki-Nakagawa-Sakata link in unified models with natural neutrino mass hierarchy". In: *Phys. Rev. D* 67 (11 2003), p. 113010. DOI: 10.1103/PhysRevD.67.113010. URL: https://link.aps.org/doi/10.1103/PhysRevD.67.113010.

[23] R.B. Patterson. "Prospects for Measurement of the Neutrino Mass Hierarchy". In: *Annual Review of Nuclear and Particle Science* 65.Volume 65, 2015 (2015), pp. 177–192. ISSN: 1545-4134. DOI: https://doi.org/10.1146/annurev-nucl-102014-021916. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-nucl-102014-021916.

[24] Raul Jimenez et al. "Can we measure the neutrino mass hierarchy in the sky?" In: *Journal of Cosmology and Astroparticle Physics* 2010.05 (2010), p. 035. DOI: 10.1088/1475-7516/2010/05/035. URL: https://dx.doi.org/10.1088/1475-7516/2010/05/035.

[25] Hitoshi Murayama and Carlos Peña Garay. "Neutrinoless double beta decay in light of SNO salt data". In: *Phys. Rev. D* 69 (3 2004), p. 031301. DOI: 10.1103/PhysRevD.69.031301. URL: https://link.aps.org/doi/10.1103/PhysRevD.69.031301.

[26] Arthur B. McDonald. "Nobel Lecture: The Sudbury Neutrino Observatory: Observation of flavor change for solar neutrinos". In: *Rev. Mod. Phys.* 88 (3 July 2016), p. 030502. DOI: 10.1103/RevModPhys.88.030502. URL: https://link.aps.org/doi/10.1103/RevModPhys.88.030502.

[27] F. P. An et al. "Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay". In: *Phys. Rev. Lett.* 130 (16 Apr. 2023), p. 161802. DOI: 10.1103/PhysRevLett.130.161802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.161802.

[28] H. Minakata et al. "Determining neutrino mass hierarchy by precision measurements in electron and muon neutrino disappearance experiments". In: *Phys. Rev. D* 74 (5 Sept. 2006), p. 053008. DOI: 10.1103/PhysRevD.74.053008. URL: https://link.aps.org/doi/10.1103/PhysRevD.74.053008.

[29] R. L. Workman and Others. "Review of Particle Physics". In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

[30] C. A. Argüelles et al. "Measuring Oscillations with a Million Atmospheric Neutrinos". In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041055. DOI: 10.1103/PhysRevX.13.041055. URL: https://link.aps.org/doi/10.1103/PhysRevX.13.041055.

[31] K. Abe et al. "Atmospheric neutrino oscillation analysis with external constraints in Super-Kamiokande I-IV". In: *Phys. Rev. D* 97 (7 Apr. 2018), p. 072001. DOI: 10.1103/PhysRevD.97.072001. URL: https://link.aps.org/doi/10.1103/PhysRevD.97.072001.

[32] R. Abbasi et al. "Measurement of atmospheric neutrino mixing with improved IceCube DeepCore calibration and data processing". In: *Phys. Rev. D* 108 (1 July 2023), p. 012014. DOI: 10.1103/PhysRevD.108.012014. URL: https://link.aps.org/doi/10.1103/PhysRevD.108.012014.

[33] P. Adamson et al. "Combined Analysis of ν_μ Disappearance and $\nu_\mu \rightarrow \nu_e$ Appearance in MINOS Using Accelerator and Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 112 (19 May 2014), p. 191801. DOI: 10.1103/PhysRevLett.112.191801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.112.191801.

[34] P. Adamson et al. "Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset". In: *Phys. Rev. Lett.* 125 (13 Sept. 2020), p. 131802. DOI: 10.1103/PhysRevLett.125.131802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.125.131802.

[35] L. Aliaga et al. "Design, calibration, and performance of the MINERvA detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 743 (2014), pp. 130–159. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2013.12.053. URL: https://www.sciencedirect.com/science/article/pii/S0168900214000035.

[36] P Stowell et al. "Tuning the GENIE pion production model with MINERvA data". In: *Physical Review D* 100.7 (2019), p. 072005.

[37] J. Mousseau et al. "Measurement of partonic nuclear effects in deep-inelastic neutrino scattering using MINERvA". In: *Phys. Rev. D* 93 (7 Apr. 2016), p. 071101. DOI: 10.1103/PhysRevD.93.071101. URL: https://link.aps.org/doi/10.1103/PhysRevD.93.071101.

[38] A. Filkins et al. "Double-differential inclusive charged-current ν_μ cross sections on hydrocarbon in MINERvA at $<E_\nu> \sim 3.5$ GeV". In: *Phys. Rev. D* 101 (11 June 2020), p. 112007. DOI: 10.1103/PhysRevD.101.112007. URL: https://link.aps.org/doi/10.1103/PhysRevD.101.112007.

[39] J. Kleykamp et al. "Simultaneous Measurement of ν_μ Quasielasticlike Cross Sections on CH, C, H_2O, Fe, and Pb as a Function of Muon Kinematics at MINERvA". In: *Phys. Rev. Lett.* 130 (16 Apr. 2023), p. 161801. DOI: 10.1103/PhysRevLett.130.161801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.161801.

[40] L. Aliaga et al. "Neutrino flux predictions for the NuMI beam". In: *Phys. Rev. D* 94 (9 2016), p. 092005. DOI: 10.1103/PhysRevD.94.092005. URL: https://link.aps.org/doi/10.1103/PhysRevD.94.092005.

[41] A. Bercellie et al. "Simultaneous Measurement of Muon Neutrino ν_μ Charged-Current Single π^+ Production in CH, C, H_2O, Fe, and Pb Targets in MINERvA". In: *Phys. Rev. Lett.* 131 (1 July 2023), p. 011801. DOI: 10.1103/

PhysRevLett.131.011801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.131.011801.

[42] C. M. Marshall et al. "Measurement of Neutral-Current K^+ Production by Neutrinos using MINERvA". In: *Phys. Rev. Lett.* 119 (1 July 2017), p. 011802. DOI: 10.1103/PhysRevLett.119.011802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.119.011802.

[43] T. Cai et al. "Measurement of the axial vector form factor from antineutrino-proton scattering". In: *Nature* 614.7946 (Feb. 2023), pp. 48–53. ISSN: 1476-4687. DOI: 10.1038/s41586-022-05478-3. URL: http://dx.doi.org/10.1038/s41586-022-05478-3.

[44] Sowjanya Gollapinni et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report (Volume II: DUNE Physics)". In: (Jan. 2020). DOI: 10.2172/1699419. URL: https://www.osti.gov/biblio/1699419.

[45] A. Abed Abud et al. "Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report". In: *Instruments* 5.4 (2021). ISSN: 2410-390X. DOI: 10.3390/instruments5040031. URL: https://www.mdpi.com/2410-390X/5/4/31.

[46] DUNE Collaboration et al. *The DUNE Science Program.* 2025. DOI: 10.48550/ARXIV.2503.23291. URL: https://arxiv.org/abs/2503.23291.

[47] Chris Marshall. *DUNE: Overview & Science.* P5 Town Hall at Fermilab and Argonne. Mar. 21, 2023. URL: https://indico.fnal.gov/event/58272/contributions/262187/.

[48] Hyper-Kamiokande Proto-Collaboration et al. *Hyper-Kamiokande Design Report.* 2018. arXiv: 1805.04163 [physics.ins-det]. URL: https://arxiv.org/abs/1805.04163.

[49] Ali Ajmi. "Status of the Hyper-Kamiokande Experiment". In: *Proceedings of 16th International Conference on Heavy Quarks and Leptons — PoS(HQL2023).* HQL2023. Sissa Medialab, Apr. 2024, p. 098. DOI: 10.22323/1.462.0098. URL: http://dx.doi.org/10.22323/1.462.0098.

[50] JUNO Collaboration. "JUNO physics and detector". In: *Progress in Particle and Nuclear Physics* 123 (2022), p. 103927. ISSN: 0146-6410. DOI: https://doi.org/10.1016/j.ppnp.2021.103927. URL: https://www.sciencedirect.com/science/article/pii/S0146641021000880.

[51] Fengpeng An et al. "Neutrino physics with JUNO". In: *Journal of Physics G: Nuclear and Particle Physics* 43.3 (Feb. 2016), p. 030401. ISSN: 1361-6471. DOI: 10.1088/0954-3899/43/3/030401. URL: http://dx.doi.org/10.1088/0954-3899/43/3/030401.

[52] JUNO Collaboration. *JUNO Detector Begins Liquid Filling Process*. Dec. 23, 2024. URL: http://juno.ihep.cas.cn/PPjuno/202412/t20241223_895765.html.

[53] P. Adamson et al. "The NuMI neutrino beam". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (2016), pp. 279–306. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2015.08.063. URL: https://www.sciencedirect.com/science/article/pii/S016890021501027X.

[54] D. A. A. Wickremasinghe and K. Yonehara. *Updates and Lessons Learned from NuMI Beamline at Fermilab*. 2024. arXiv: 2412.12368 [physics.acc-ph]. URL: https://arxiv.org/abs/2412.12368.

[55] *First NuMI Beam from Main Injector*. Fermilab Today. Dec. 10, 2004. URL: https://history.fnal.gov/historical/accelerator/numi_first_beam_from_main_injector.html.

[56] Robert Ainsworth et al. "High intensity operation using proton stacking in the Fermilab Recycler to deliver 700 kW of 120 GeV proton beam". In: *Phys. Rev. Accel. Beams* 23 (12 Dec. 2020), p. 121002. DOI: 10.1103/PhysRevAccelBeams.23.121002. URL: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.121002.

[57] Alex Himmel. email communication. June 24, 2024.

[58] Anna B Cooleybeck. "Blessing package for ND All POT plots". NOvA Internal Note 62927. 2024.

[59] Anna B Cooleybeck. "Blessing package for FD All POT plots". NOvA Internal Note 62786. 2024.

[60] Marco Del Tutto and Gareth Kafka. "Blessing Package - Beam Simulation Plots". NOvA Internal Note 13524. 2015.

[61] E.C. Dukes et al. "The NOvA power distribution system". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 902 (2018), pp. 123–137. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2018.06.021. URL: https://www.sciencedirect.com/science/article/pii/S0168900218307423.

[62] D S Ayres et al. *The NOvA Technical Design Report*. Tech. rep. Oct. 2007. DOI: 10.2172/935497. URL: https://www.osti.gov/biblio/935497.

[63] S. Mufson et al. "Liquid scintillator production for the NOvA experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 799 (2015), pp. 1–9. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2015.07.026. URL: https://www.sciencedirect.com/science/article/pii/S0168900215008554.

[64] A Norman et al. “The NOvA Timing System: A system for synchronizing a long baseline neutrino experiment”. In: *Journal of Physics: Conference Series* 396.1 (Dec. 2012), p. 012034. ISSN: 1742-6596. DOI: 10.1088/1742-6596/396/1/012034. URL: http://dx.doi.org/10.1088/1742-6596/396/1/012034.

[65] M Fischler et al. “NOvA Event Building, Buffering and Data-Driven Triggering From Within the DAQ System”. In: *Journal of Physics: Conference Series* 396.1 (Dec. 2012), p. 012020. ISSN: 1742-6596. DOI: 10.1088/1742-6596/396/1/012020. URL: http://dx.doi.org/10.1088/1742-6596/396/1/012020.

[66] B. P. Abbott et al. “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living Reviews in Relativity* 23.1 (July 2020), p. 3. ISSN: 1433-8351. DOI: 10.1007/s41114-020-00026-9. URL: https://doi.org/10.1007/s41114-020-00026-9.

[67] S Al Kharusi et al. “SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy”. In: *New Journal of Physics* 23.3 (Mar. 2021), p. 031201. ISSN: 1367-2630. DOI: 10.1088/1367-2630/abde33. URL: http://dx.doi.org/10.1088/1367-2630/abde33.

[68] Yibing Zhang. “Ana2020 Numu FHC Event Displays”. NOvA Internal Note 45826. 2020.

[69] Yibing Zhang. “Ana2020 Nue FHC Event Displays”. NOvA Internal Note 45850. 2020.

[70] Daniel Stephen Pershey. “A Measurement of ν_e Appearance and ν_μ Disappearance Neutrino Oscillations with the NOvA Experiment”. PhD thesis. California Institute of Technology, 2018. DOI: 10.7907/RC0W-ZG32. URL: https://resolver.caltech.edu/CaltechTHESIS:06012018-215047231.

[71] Micah Groh. “Blessed Plots: Event CVN 2020”. NOvA Internal Note 43911. 2020.

[72] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(03)01368-8. URL: https://www.sciencedirect.com/science/article/pii/S0168900203013688.

[73] C. Andreopoulos et al. “The GENIE neutrino Monte Carlo generator”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 614.1 (2010), pp. 87–104. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2009.12.009. URL: https://www.sciencedirect.com/science/article/pii/S0168900209023043.

[74] Chris Hagmann, David Lange, and Douglas Wright. "Cosmic-ray shower generator (CRY) for Monte Carlo transport codes". In: *2007 IEEE Nuclear Science Symposium Conference Record*. Vol. 2. 2007, pp. 1143–1146. DOI: 10.1109/NSSMIC.2007.4437209.

[75] A Aurisano et al. "The NOvA simulation chain". In: *Journal of Physics: Conference Series* 664.7 (2015), p. 072002. ISSN: 1742-6596. DOI: 10.1088/1742-6596/664/7/072002. URL: http://dx.doi.org/10.1088/1742-6596/664/7/072002.

[76] Adam Lister. "Calibration: A One Page Summary". NOvA Internal Note 65198. 2025.

[77] M. A. Acero et al. "Improved measurement of neutrino oscillation parameters by the NOvA experiment". In: *Phys. Rev. D* 106 (3 2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004. URL: https://link.aps.org/doi/10.1103/PhysRevD.106.032004.

[78] Shiqi Yu. "Blessing package: Extrapolation Cartoon for Ana2018". NOvA Internal Note 29612. 2018.

[79] Brajesh Choudhary. "Ana2024 3F Summary of Systematic Uncertainties". NOvA Internal Note 63147. 2024.

[80] Luke Vinton. "Measurement of Muon Neutrino Disappearance with the NOvA Experiment". University of Sussex, Jan. 2018. DOI: 10.2172/1423216. URL: http://dx.doi.org/10.2172/1423216.

[81] M. A. Acero et al. "New constraints on oscillation parameters from ν_e appearance and ν_μ disappearance in the NOvA experiment". In: *Physical Review D* 98.3 (Aug. 2018). ISSN: 2470-0029. DOI: 10.1103/physrevd.98.032012. URL: http://dx.doi.org/10.1103/PhysRevD.98.032012.

[82] Thomas Warburton. "Ana 20 NuMu Quantile Plots". NOvA Internal Note 47242. 2020.

[83] Aaron Mislivec. "Pt Extrapolation Blessed Plots for the 2020 Analysis". NOvA Internal Note 45980. 2020.

[84] M. A. Acero et al. "Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data". In: *The European Physical Journal C* 80.12 (2020). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-020-08577-5. URL: http://dx.doi.org/10.1140/epjc/s10052-020-08577-5.

[85] M. A. Acero et al. "Expanding neutrino oscillation parameter measurements in NOvA using a Bayesian approach". In: *Phys. Rev. D* 110 (1 2024), p. 012005. DOI: 10.1103/PhysRevD.110.012005. URL: https://link.aps.org/doi/10.1103/PhysRevD.110.012005.

[86] Thomas J. Faulkenberry. "A tutorial on generalizing the default Bayesian *t*-test via posterior sampling and encompassing priors". eng. In: *Communications for Statistical Applications and Methods* 26.2 (Mar. 2019), pp. 217–238. ISSN: 2287-7843. URL: http://http://www.csam.or.kr/journal/view.html?doi=10.29220/CSAM.2019.26.2.217.

[87] K Abe et al. "The T2K experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (2011), pp. 106–135.

[88] K. Abe et al. "Proposal for an Extended Run of T2K to 20×10^{21} POT". In: (2016). DOI: 10.48550/ARXIV.1609.04111. URL: https://arxiv.org/abs/1609.04111.

[89] Joe Walsh. "Constraining the T2K neutrino oscillation parameter results using data from the off-axis near detector, ND280: Implementation of a nucleon removal energy systematic uncertainty treatment in the BANFF fit". PhD thesis. Lancaster University, 2022. DOI: 10.17635/lancaster/thesis/1604.

[90] W. K. Hastings. "Monte Carlo sampling methods using Markov chains and their applications". In: *Biometrika* 57.1 (1970), pp. 97–109. ISSN: 0006-3444. DOI: 10.1093/biomet/57.1.97. URL: http://dx.doi.org/10.1093/BIOMET/57.1.97.

[91] K. Abe et al. "Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target". In: *The European Physical Journal C* 83.9 (2023). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-11819-x. URL: http://dx.doi.org/10.1140/epjc/s10052-023-11819-x.

[92] The T2K and NOvA Collaborations. "Joint neutrino oscillation analysis of data from the T2K and NOvA experiments". 2025. Submitted to Journal.

[93] Zoya Vallari. *Results from a joint analysis of data from NOvA and T2K*. Fermilab Joint Experimental–Theoretical Physics Seminar. Feb. 16, 2024. URL: https://indico.fnal.gov/event/62062/.

[94] Sabya Sachi Chatterjee and Antonio Palazzo. "Neutrino Non-standard Interactions: A Solution to the NOνA and T2K Tension". In: *High Energy Physics Symposium*. Springer. 2022, pp. 67–72.

[95] Ushak Rahaman and Sushant K Raut. "On the tension between the latest NOνA and T2K data". In: *The European Physical Journal C* 82.10 (2022), p. 910.

[96] Ivan Esteban et al. "NuFit-6.0: updated global analysis of three-flavor neutrino oscillations". In: *Journal of High Energy Physics* 2024.12 (Dec. 2024). ISSN: 1029-8479. DOI: 10.1007/jhep12(2024)216. URL: http://dx.doi.org/10.1007/JHEP12(2024)216.

[97] Ryan Patterson et al. "Development of a joint oscillation analysis by the NOvA and T2K collaborations – Snowmass LOI". NOvA Internal Note 47117. 2020.

[98] The T2K and NOvA Collaborations. *Bifrost.* Accessed 2025. URL: https://github.com/nova-t2k/bifrost.

[99] The T2K and NOvA Collaborations. *DummyLLH.* Accessed 2025. URL: https://github.com/nova-t2k/DummyLLH.

[100] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *The European Physical Journal C* 71.2 (2011). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0. URL: http://dx.doi.org/10.1140/epjc/s10052-011-1554-0.

[101] L. Aliaga and Others. "Use of the NOvA and T2K Flux Predictions in a Joint Analysis". NOvA Internal Note 54582. 2023.

[102] Jeremy Wolcott. "Measurement of the charged-current quasi-elastic cross-section for electron neutrinos on a hydrocarbon target". PhD thesis. University of Rochester, 2015. URL: http://hdl.handle.net/1802/30416.

[103] E. Cataño Mur and Others. "Assessment of possible correlations between NOvA and T2K detector simulation and systematic uncertainties". NOvA Internal Note 54582. 2023.

[104] E. Atkin and Others. "Impact of Interaction Model Parameters on the NOvA-T2K Joint Oscillation Fit". NOvA Internal Note 54582. 2023.

[105] Melanie Day and Kevin S. McFarland. "Differences in quasielastic cross sections of muon and electron neutrinos". In: *Phys. Rev. D* 86 (5 2012), p. 053003. DOI: 10.1103/PhysRevD.86.053003. URL: https://link.aps.org/doi/10.1103/PhysRevD.86.053003.

[106] E. Atkin and Others. "NOvA-T2K Joint Oscillation Fit: Fake Data Studies". NOvA Internal Note 54582. 2023.

[107] Edward Atkin and Others. "Data fit results from a Joint T2K-NOvA analysis". NOvA Internal Note 54582. 2023.

[108] Andrew Gelman, Xiao-Li Meng, and Hal Stern. "POSTERIOR PREDICTIVE ASSESSMENT OF MODEL FITNESS VIA REALIZED DISCREPANCIES". In: *Statistica Sinica* 6.4 (1996), pp. 733–760. ISSN: 10170405, 19968507. URL: http://www.jstor.org/stable/24306036 (visited on 05/17/2025).

[109] Particle Data Group et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020), p. 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104. eprint: https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673722/ptaa104.pdf. URL: https://doi.org/10.1093/ptep/ptaa104.

[110] Hiroshi Nunokawa, Stephen Parke, and Renata Zukanovich Funchal. "Another possible way to determine the neutrino mass hierarchy". In: *Physical Review D* 72.1 (2005), p. 013009.

[111] Balu Ertl. *File:Euclidean Voronoi diagram.svg*. [Online; accessed 2-August-2024]. 2020. URL: https://commons.wikimedia.org/w/index.php?title=File:Euclidean_Voronoi_diagram.svg&oldid=486001345.

[112] René Descartes. *The World and Other Writings*. Trans. by Stephen Gaukroger. Cambridge texts in the history of philosophy. Cambridge University Press, 1998. ISBN: 0521631580.

[113] John Snow. "On the mode of communication of cholera". In: (1855). URL: https://www.ph.ucla.edu/epi/snow/snowmap1.pdf.

[114] Atsuyuki Okabe et al. *Spatial Tessellations*. en. 2nd ed. Wiley Series in Probability and Statistics. Chichester, England: John Wiley & Sons, May 2000.

[115] Stuart Lloyd. "Least squares quantization in PCM". In: *IEEE Transactions on Information Theory* 28.2 (1982), pp. 129–137.

[116] Michael Balzer. "Capacity-constrained Voronoi diagrams in continuous spaces". In: *2009 sixth international symposium on Voronoi diagrams*. IEEE. 2009, pp. 79–88.

[117] Wikipedia contributors. *Shoelace formula*. [Online; accessed 16-August-2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Shoelace_formula.

[118] Z Furedi and I Barany. "Computing the volume is difficult". In: *Proceedings of the eighteenth annual ACM symposium on Theory of computing*. 1986, pp. 442–447.

[119] László Lovász and Santosh Vempala. "Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm". In: *Journal of Computer and System Sciences* 72.2 (2006), pp. 392–417.

[120] NOvA-T2K Leadership. *Report from leadership meeting*. T2K+NOvA workshop Spring 2025. Apr. 2, 2025.

[121] Fermilab Test Beam Facility. "Beam Structure and Delivery Path". Accessed: April 4 2025. URL: https://ftbf.fnal.gov/beam-delivery-path/.

[122] "NOvA Test Beam Internal Wiki". Accessed: April 2, 2025.

[123] Teresa Lackey. "Proton Scattering in NOvA Test Beam". PhD thesis. Indiana University, July 2022. URL: https://www.osti.gov/biblio/1882555.

[124] Varun Ranga Raj. "Test Beam Beamline Timing Studies". NOvA Internal Note 36414. 2019.

[125] Alex Sousa. "Test Beam Analysis Update". NOvA Internal Note 63459. 2024.

[126] David F. Dueñas Tonguino. "Blessing package: NOvA Test Beam Pions". NOvA Internal Note 61866. 2024.

[127] Teresa Lackey. "What questions can we answer with TestBeam data?" NOvA Internal Note 65201. 2025.

[128] Nicholas Metropolis et al. *Equation of state calculations by fast computing machines.* 1953. DOI: 10.2172/4390578. URL: https://www.osti.gov/biblio/4390578.

[129] Andrew Gelman et al. *Bayesian data analysis, second edition.* 2nd ed. Chapman & Hall/CRC Texts in Statistical Science. Chapman & Hall/CRC, 2003.

Appendix A

THE NOvA TEST BEAM EXPERIMENT

As shown in Figure 2.21, some of the largest individual systematics in NOvA's oscillation measurements come from detector calibration, lepton reconstruction, and detector response modeling. As mentioned in Section 2.6, the extrapolation procedure significantly helps with reducing systematic errors. However, errors could be reduced even further if we could observe particles of known species, energy, and momentum pass through a NOvA detector. This motivated the construction of the NOvA Test Beam experiment.

As part of my work on NOvA, I spent significant time at Fermilab helping prepare the detector and beamline components for data taking.

A.1 Experiment Overview

The NOvA Test Beam detector was built in the MC7 enclosure along the MCenter secondary beamline at the Fermilab Test Beam Facility (FTBF). Figure A.1 depicts the beam path at FTBF.

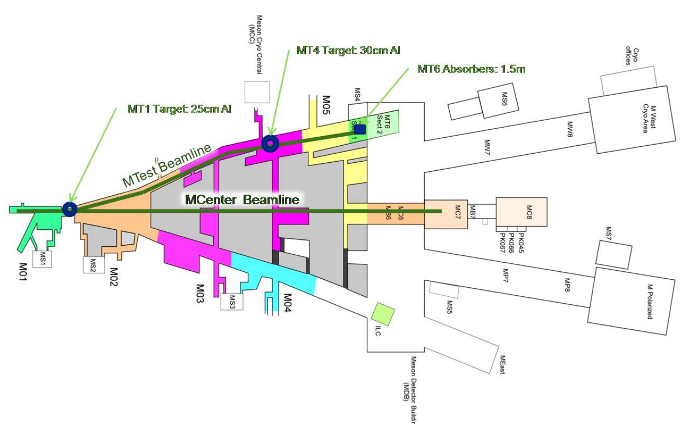

Figure A.1: Beam Paths at the Fermilab Test Beam Facility. Primary beam consists of 120 GeV protons coming from the switchyard (see Figure 2.1 for a diagram of the full Fermilab accelerator complex). Secondary beam is sent down the MCenter beamline towards the MC7 enclosure, where the NOvA Test Beam Detector was located [121].

Protons from the main injector are sent to the Switchyard. From there, some of these 120 GeV protons are diverted towards the MCenter beamline. At MC6, the MCenter beam is sent through secondary beamline components, converting the beam to lower-energy protons and pions. These particles are then sent to the tertiary beam source in the MC7 enclosure. By instrumenting this tertiary beamline, NOvA can identify particles before they reach the NOvA Test Beam detector.

Figure A.2: Photograph of the MC7 enclosure in Fermilab taken with a fish-eye lens. The beam comes from the right side. The analyzer magnet is located near the center-right of the image, above the orange safety netting. The NOvA Test Beam Detector is the large black rectangular prism located at the left of the image [122].

The detector consisted of 63 planes in two blocks, with a total mass of 30 tonnes. This is about 10% the volume of the Near Detector. Most of the detector electronics are leftovers from the construction of the Far Detector.

Because the Test Beam detector was designed to reuse NOvA detector parts, we could mostly reuse the Data Acquisition (DAQ) system from the main detectors. Aside from accounting for the much smaller detector volume, the biggest change made to the Test Beam DAQ was accounting for the much higher event rates. As a way to mitigate noisy channels, FEBs automatically shut-off when event rates go too high. Unfortunately for Test Beam, the event rates from beam events could be enough to shut off FEBs, especially for cells near the front of the detector. As a result, the Test Beam DAQ was modified to repeatedly send reset signals to all FEBs.

A.2 Beamline components

The Tertiary beamline was instrumented with three main sets of detectors: a Time-of-Flight (ToF) system for measuring particle energies, a set of Multi-Wire Proportional Chambers (MWPCs) for measuring momenta, and a Cherenkov counter for tagging electrons. Figure A.3 shows the tertiary beamline schematically.

The ToF system consisted of three scintillator paddles. Two of these paddles were optically mated to photomultiplier tubes (PMTs), and the third was instead mated to a Silicon photomultiplier (SiPM) for readout. At 1 GeV particle momenta, this system provided the ability to separate protons, kaons, and faster species. However,

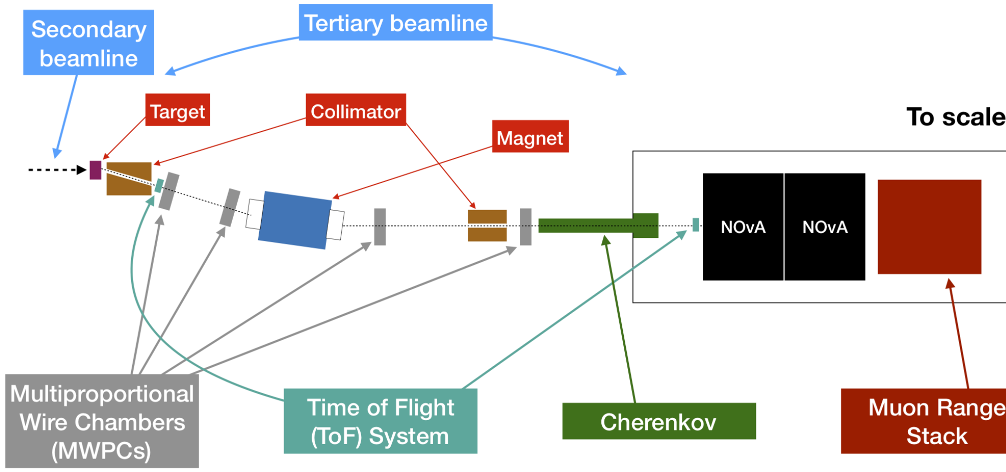

Figure A.3: Schematic of the Test Beam Tertiary beamline. 64 GeV protons from the secondary beam strike a copper target. The daughter particles from the collision pass through several beamline detectors for particle tagging. The analyzer magnet is used to select for specific particles. An additional Time-of-Flight scintillator paddle was inserted between the Cherenkov counter and the fourth MWPC later during data-taking [122].

the other beamline components were needed to properly separate muons, pions, and electrons [123].

The MWPCs and analyzer magnet were responsible for momentum reconstruction. Each MWPC consists of a set of horizontal and vertical anode wires set at a large voltage relative to a cathode. A passing tertiary beam particle passing through can ionize the argon-based gaseous medium inside the chamber. These ions will drift toward the anode wires. By seeing which wires produce signals coincident with a beam trigger, we can locate the original beam particle's trajectory in three-dimensional space. Because MWPCs are located both before and after the analyzer magnet, it becomes possible to measure how much the particle bends in the magnetic field, enabling momentum measurements.

Finally, the Cherenkov counter consists of a tube filled with 1 atm CO_2. Carbon Dioxide's low index of refraction means that only electrons will be traveling fast enough to produce Cherenkov light. This allows for differentiating electrons from muons and pions.

A new Data Acquisition System (DAQ) had to be built specifically for the beamline components, and specifically. Each beamline component had to be triggered to record all events with each 4.2 ms beam spill and recorded all events, and care had to be taken to match specific events from each detector [124].

A.3 Current Status

Data was taken from 2019 to 2022. In this time, the NOvA Test Beam detector accumulated events associated with 33,975 particles that pass analysis cuts [125]. This sample includes over 9,000 protons, 2,000 electrons, and 350 kaons. Figure A.4 shows spectra of the particles seen in the Test Beam detector.

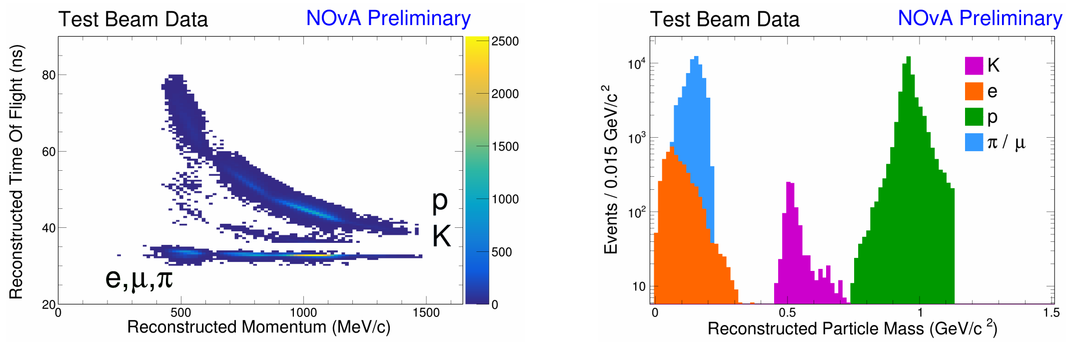

(a) Momentum vs Time-of-Flight (b) Reconstructed masses

Figure A.4: Spectra of reconstructed analysis-quality events at the Test beam. These events passed beamline and mass cuts [126].

Because their masses are so similar, it is not possible to separate muons and charged pions with exclusively beamline components. As such, work is ongoing to use data from the NOvA Test Beam Detector to distinguish these particles.

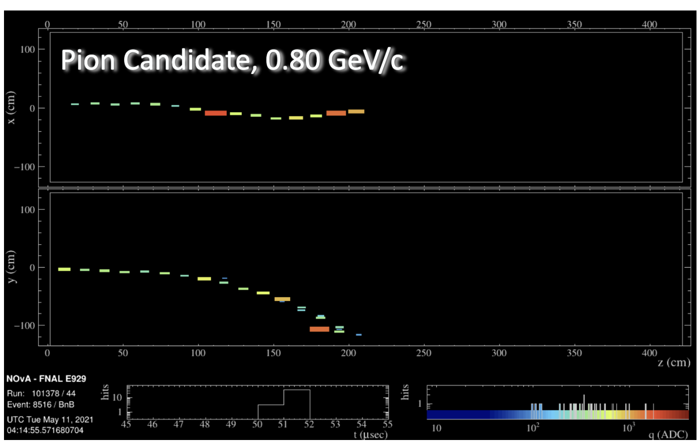

Figure A.5: Event Display for a 0.8 GeV/c pion candidate in the NOvA Test Beam detector [125].

NOvA is currently preparing several analyses to use the collected Test Beam data to improve detector response models for protons, pions, and electrons. This data can also be used to improve our calibration by building energy scale systematics which are dependent on particle species and energy. Finally, we can use Test Beam data as a way to test reconstruction algorithms [127].

Appendix B

MR^2T^2 MCMC

NOvA has built the Aria fitter which implements the Metropolis-Rosenbluth-Rosenbluth-Teller-Teller (MR^2T^2) MCMC algorithm (also known in literature as the Metropolis algorithm) [128]. This algorithm is able to estimate a posterior distribution and find Bayesian credible intervals.

Consider Bayes' Rule:

$$P(\theta|x) = \frac{P(x|\theta)P(\theta)}{P(x)} \tag{B.1}$$

where x is data, θ is the model parameters, $P(\theta|x)$ is the posterior distribution, $P(x|\theta) := \mathscr{L}(x, \theta)$ is the likelihood of the data given the model parameters, and $P(\theta)$ is the prior distribution for θ.

The goal is to approximate the posterior $P(\theta|x)$. As data is fixed, let us make the definition $\pi(\theta) := P(\theta|x)$. We will use a Markov chain, which is a sequence of random variables θ_i such that θ_{i+1} only depends on θ_i. If the Markov chain holds the following property:

$$\left(\lim_{n\to\infty} \theta_n\right) \sim \pi(\theta)$$

Then, new iterations from the Markov chain will be distributed according to the posterior distribution, as desired. The evolution of the Markov chain is governed by its transition probability $P(\theta'|\theta)$, namely the probability of moving to state θ' given that the chain is at state θ

This can happen if π is the unique stationary distribution of the Markov chain. Most Markov chains have a unique stationary distribution as long as they can reach all possible states of θ with nonzero probability [129]. In order for that stationary distribution to be equal to our desired posterior π, we can apply a stricter condition on the transition probability [90]:

$$\pi(\theta)P(\theta'|\theta) = \pi(\theta')P(\theta|\theta') \tag{B.2}$$

We can break the transition probability up into two pieces:

$$P(\theta'|\theta) = g(\theta'|\theta)A(\theta', \theta) \tag{B.3}$$

where $g(\theta'|\theta)$ is the probability of proposing θ' as the next step in the chain, given that we started at θ, and $A(\theta',\theta)$ is the probability of accepting the proposal and actually shifting the chain to θ'. Using Equation B.2 we can see:

$$\begin{aligned}
\frac{\pi(\theta)}{\pi(\theta')} &= \frac{P(\theta|\theta')}{P(\theta'|\theta)} \\
&= \frac{g(\theta|\theta')A(\theta,\theta')}{g(\theta'|\theta)A(\theta',\theta)} \\
\Rightarrow \frac{A(\theta',\theta)}{A(\theta,\theta')} &= \frac{\pi(\theta')g(\theta|\theta')}{\pi(\theta)g(\theta'|\theta)}
\end{aligned}$$

Now, we can take advantage of Bayes' Rule (Eq. B.1), remembering that we defined $\pi(\theta) = P(\theta|x)$ and $\mathscr{L}(x,\theta) = P(x|\theta)$:

$$\begin{aligned}
\frac{A(\theta',\theta)}{A(\theta,\theta')} &= \frac{\frac{P(x|\theta')P(\theta')}{P(x)}g(\theta|\theta')}{\frac{P(x|\theta)P(\theta)}{P(x)}g(\theta'|\theta)} \\
&= \frac{\mathscr{L}(x,\theta')P(\theta')g(\theta|\theta')}{\mathscr{L}(x,\theta)P(\theta)g(\theta'|\theta)}
\end{aligned} \tag{B.4}$$

Thus, we must choose $A(\theta,\theta')$ to satisfy this property. The MR^{2}T^2 choice is:

$$A(\theta',\theta) = \min\left(1, \frac{\mathscr{L}(x,\theta')P(\theta')g(\theta|\theta')}{\mathscr{L}(x,\theta)P(\theta)g(\theta'|\theta)}\right)$$

If the proposal density $g(\theta'|\theta)$ is chosen to be a function of $|\theta - \theta'|$, and therefore symmetric in θ and θ', then $g(\theta|\theta') = g(\theta'|\theta)$. The acceptance probability simplifies:

$$A(\theta',\theta) = \min\left(1, \frac{\mathscr{L}(x,\theta')P(\theta')}{\mathscr{L}(x,\theta)P(\theta)}\right) \tag{B.5}$$

Note that if $\mathscr{L}(x,\theta')P(\theta') > \mathscr{L}(x,\theta)P(\theta)$, then $A(\theta',\theta) = 1$ and the proposal to jump to θ' is always accepted. In this case, the acceptance probability of the reverse will be $A(\theta,\theta') = \frac{\mathscr{L}(x,\theta)P(\theta)}{\mathscr{L}(x,\theta')P(\theta')}$, and Equation B.4 holds.

Thus in MR^{2}T^2 MCMC, the acceptance ratio is simply:

$$a(\theta',\theta) = \frac{\mathscr{L}(x,\theta')P(\theta')}{\mathscr{L}(x,\theta)P(\theta)}$$

To go from one step to the next in the chain, generate a proposal θ' and a number α uniformly between 0 and 1. If $\alpha < a(\theta',\theta)$, then accept the proposal θ' and add it to the chain. If $\alpha > a(\theta',\theta)$, then reject the proposal and add θ to the chain again. After enough iterations, θ_n will be distributed according to the desired posterior distribution.