Improving MicroBooNE's Sterile Neutrino Sensitivity with PROfit
REU Program at Columbia University - Nevis Labs

June Zey[1]

[1]University of Texas at Austin, Department of Physics, Austin, TX

August 10, 2026

NEVIS LABORATORIES
COLUMBIA UNIVERSITY

Abstract

Anomalous low energy event excesses observed by MiniBooNE have motivated a decades long search for a fourth, "sterile" neutrino mass state that mixes with the three known active flavors but does not interact weakly. MicroBooNE, a 170 ton liquid argon time projection chamber (LArTPC) located on the Booster Neutrino Beamline at Fermilab, was built in part to investigate this excess with substantially improved particle identification over MiniBooNE's Cherenkov detector. This work presents two contributions to MicroBooNE's ongoing 3+1 sterile neutrino search using the PROfit fitting framework: a systematic comparison of BAGELS-based ν_μ disappearance channel configurations against the existing WireCell selection standard, and a novel method for generating an approximate, fully oscillated $\nu_\mu \rightarrow \nu_e$ appearance sample directly from existing ν_μ and ν_e Monte Carlo samples, without the cost of a dedicated simulation.

Contents

1 Background

Neutrinos are fundamental spin-$\frac{1}{2}$ particles that come in three flavors: electron, muon, and tau. They have several unusual properties. First, neutrinos interact only via the weak force (and gravity), giving them an extremely small interaction cross sections. Despite this minuscule coupling, neutrinos are enormously abundant throughout the universe, produced by stars, terrestrial radioactivity, particle accelerators, and as relics of the early universe. Neutrinos were originally hypothesized to be exactly massless, but were later found to have nonzero mass through the discovery of neutrino oscillations by the Super-Kamiokande and SNO collaborations [9, 8], which observed a zenith angle dependent deficit of atmospheric muon neutrinos consistent with flavor oscillation over macroscopic distances a phenomenon impossible for massless neutrinos.

As they propagate through space, neutrinos oscillate between flavor states. This behavior arises because the three flavor eigenstates $(\nu_e, \nu_\mu, \nu_\tau)$, in which neutrinos are produced and detected, are not identical to the three mass eigenstates (ν_1, ν_2, ν_3), in which they propagate. The two bases are related by a unitary transformation known as the **PMNS** (Pontecorvo–Maki–Nakagawa–Sakata) matrix:

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

commonly parameterized in terms of three mixing angles $(\theta_{12}, \theta_{23}, \theta_{13})$ and a CP-violating phase δ_{CP}.

Because the mass eigenstates propagate with slightly different phase velocities, a neutrino produced in one flavor state acquires a nonzero probability of being detected in a different flavor state after traveling a distance L. In the two flavor approximation, this appearance probability is given by

$$P(\nu_\mu \to \nu_e) \approx \sin^2(2\theta_{13}) \sin^2(\theta_{23}) \sin^2\left(\frac{\Delta m^2_{31} L}{4E}\right) \tag{2}$$

where $\Delta m^2_{31} = m^2_3 - m^2_1$ is the relevant mass-squared splitting, E is the neutrino energy, and L is the baseline (distance traveled). The oscillation frequency is set by Δm^2_{31}, while the amplitude is set by the mixing angles.

Figure 1 plots this oscillation probability as a function of L/E for a muon neutrino initial state, showing the characteristic oscillatory exchange of probability between the three flavors.

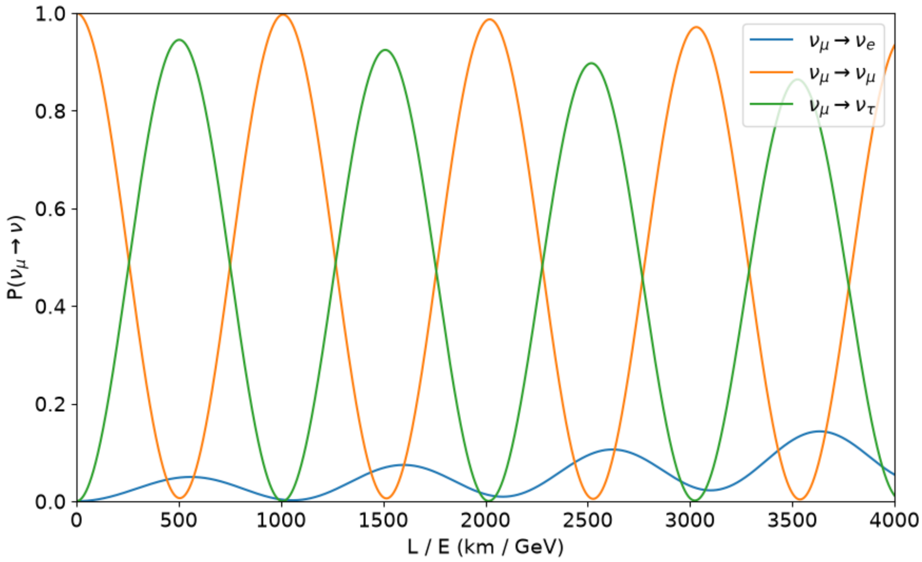

Figure 1: *Neutrino oscillation probability as a function of L/E (baseline over energy), with the initial flavor state taken to be a muon neutrino. The three curves show the probability of detecting a ν_e, ν_μ, or ν_τ after propagation.*

2 The Short-Baseline Anomaly

Neutrino experiments are broadly categorized by the energy scale of their source, which in turn dictates the experimental technique used to detect them. Reactor experiments (Double Chooz, KamLAND, Ricochet, and others) probe MeV scale antineutrinos produced by nuclear fission. Solar experiments (Super-Kamiokande, SNO+, Borexino, and others) detect MeV scale neutrinos produced by fusion processes in the Sun. Accelerator experiments (DUNE, MINOS+, NOvA, T2K, MiniBooNE, and others) instead produce GeV scale neutrino beams using proton accelerators, allowing for controlled studies of oscillation at well defined baselines and energies. This work concerns an accelerator experiment, MicroBooNE, operating on the Booster Neutrino Beamline at Fermilab.

2.1 MiniBooNE

MiniBooNE was a Cherenkov detector experiment located on the Booster Neutrino Beamline (**Boo**ster **N**eutrino **E**xperiment) at Fermilab, designed to search for neutrino oscillations at short baseline. Its mineral-oil Cherenkov detector identified charged-current interactions via their characteristic Cherenkov light rings, using ring sharpness to statistically distinguish muon like from electronlike event topologies.

MiniBooNE observed a persistent, statistically significant excess of low-energy events in its ν_e appearance sensitive sample relative to the predicted background [7]. However, the Cherenkov detector technology did not provide sufficient particle identification or spatial resolution to determine the physical origin of this excess in particular, it could not reliably distinguish a genuine $\nu_e$like signal from misidentified neutral current π^0 background events, which produce a similar two photon Cherenkov signature. To solve this issue a different type of experiment would be necessary, this is what its successor experiment MicroBooNE is intended to do investigate.

2.2 MicroBooNE

MicroBooNE is a 170-ton liquid argon time projection chamber (LArTPC) located on the same Booster Neutrino Beamline, approximately 71 m from the former MiniBooNE detector location. LArTPCs record the ionization left by charged particles as they drift toward finely spaced anode wire planes, providing millimeter scale 3D spatial resolution and calorimetric information far exceeding that of a Cherenkov detector. One of MicroBooNE's central physics goals is to use this improved capability to directly investigate the physical origin of the MiniBooNE low energy excess.

2.3 The Sterile Neutrino Hypothesis

One proposed explanation for the excess is the existence of a fourth neutrino mass state, ν_4, that does not participate in weak interactions a *sterile* neutrino. Because a sterile state cannot be produced or detected via weak processes, it is invisible to direct observation; its presence is instead inferred indirectly through its effect on the interference pattern between the three active mass states, which modifies the oscillation probabilities of Eq. 2. This extended framework is commonly referred to as the **3+1** model (three active mass states plus one sterile state).

Under the 3+1 hypothesis, the appearance oscillation probability takes the short baselines form

$$P(\nu_\mu \to \nu_e) \approx \sin^2(2\theta_{\mu e}) \sin^2\left(\frac{\Delta m^2_{41}\, E}{L}\right) \tag{3}$$

where $\Delta m^2_{41} = m^2_4 - m^2_1$ sets the oscillation *frequency* and $\sin^2(2\theta_{\mu e})$ sets the *amplitude*. Because sterile driven oscillations occur on a much shorter length scale than standard 3 flavor oscillations, short baseline experiments operating at low L/E are uniquely sensitive to this signal. MicroBooNE operates at a characteristic ratio of $L/E \approx 0.5$ km/GeV, placing it directly in the regime where a 3+1 signal would be most visible: any observed ν_e appearance beyond the expected 3 flavor background constitutes evidence for sterile-driven oscillation.

2.4 The Sterile Neutrino Model

The 3+1 model extends the standard 3×3 PMNS matrix to a 4×4 unitary matrix, adding a fourth column that governs the mixing of each active flavor with the mostly sterile mass eigenstate ν_4. The relevant elements are U_{e4}, $U_{\mu 4}$, and $U_{\tau 4}$, which give the amplitude for an electron, muon, or tau neutrino, respectively, to be found in the ν_4 mass eigenstate. Because ν_4 is presumed to be predominantly sterile, these elements are generically small, and it is their pairwise products rather than the individual U_{i4} themselves that determine the strength of any observable short baseline oscillation signal. In the two neutrino approximation commonly used for short baseline fits, the $\nu_\mu \to \nu_e$ appearance probability of Eq. 3 is written in terms of an effective mixing amplitude

$$\sin^2(2\theta_{\mu e}) \equiv 4\,|U_{e4}|^2\,|U_{\mu 4}|^2, \tag{4}$$

so that appearance signals depend on the *product* of the electron and muon flavor couplings to ν_4, and vanish if either coupling is zero. The analogous disappearance channel, used to probe $\nu_\mu \to \nu_\mu$ survival, depends instead on a single flavor's coupling to ν_4:

$$\sin^2(2\theta_{\mu\mu}) \equiv 4\,|U_{\mu 4}|^2 \left(1 - |U_{\mu 4}|^2\right). \tag{5}$$

Because appearance and disappearance measurements constrain different combinations of the underlying U_{i4} elements a product of two mixing amplitudes versus a function of one a null result in one channel does not, by itself, exclude a signal in the other; a value of $\theta_{\mu e}$ consistent with zero can still hide a large $\theta_{\mu\mu}$ if U_{e4} happens to be small. This is one of the key motivations for the multi channel comparison presented in Section, since $\nu_e \to \nu_e$, $\nu_\mu \to \nu_e$, and $\nu_\mu \to \nu_\mu$ each isolate a distinct combination of $|U_{e4}|^2$ and $|U_{\mu4}|^2$, and only by combining all three can the full mixing parameter space be meaningfully constrained.

3 PROfit

3.1 Data analysis

PROfit[11] is a fitting framework for short baseline neutrino oscillation and beyond Standard Model physics searches, developed by Mark Ross Lonergan and collaborators, and adopted as the one of the primary fitting framework for the SBN program [5]. PROfit takes in processed event selection data together with a full accounting of systematic uncertainty on the Monte Carlo prediction, and determines how well a given physics model fits the data within that systematic envelope. Its primary outputs for this analysis are best fit spectra with post fit systematic bands and profiled χ^2 exclusion contours in oscillation parameter space. One of PROfit's main functions is to generate best fit spectra such as the one shown in Figure 2.

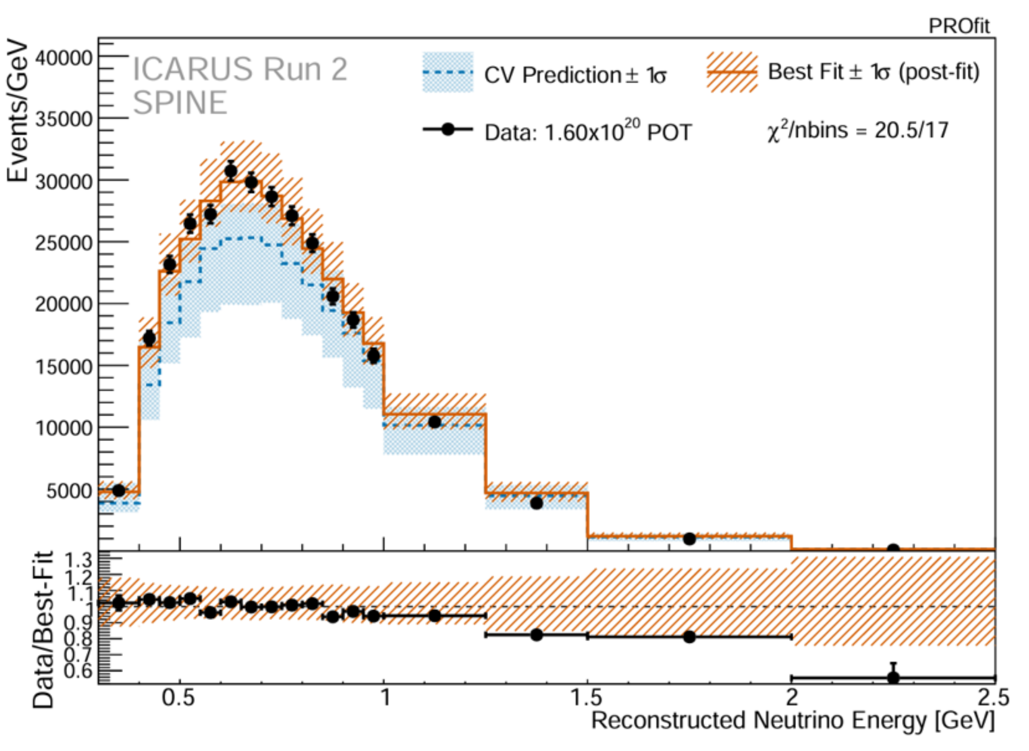

Figure 2: *A reconstructed-energy spectrum with a central value (CV) prediction and a post-fit best-fit overlay, illustrating PROfit's core fitting functionality. Figure reproduced from [5].*

The available parameter space for a sterile neutrino is conventionally summarized using an *exclusion plot* in the $(\sin^2 2\theta, \Delta m^2_{41})$ plane. Each point on such a plot corresponds to an assumed true value of the sterile oscillation parameters; points to the right of a given exclusion curve are ruled out at the stated confidence level by the observed data, while points to the left remain allowed. Another key capability of PROfit is generating sterile neutrino exclusion contours directly from fitted data, as shown in Figure 3.

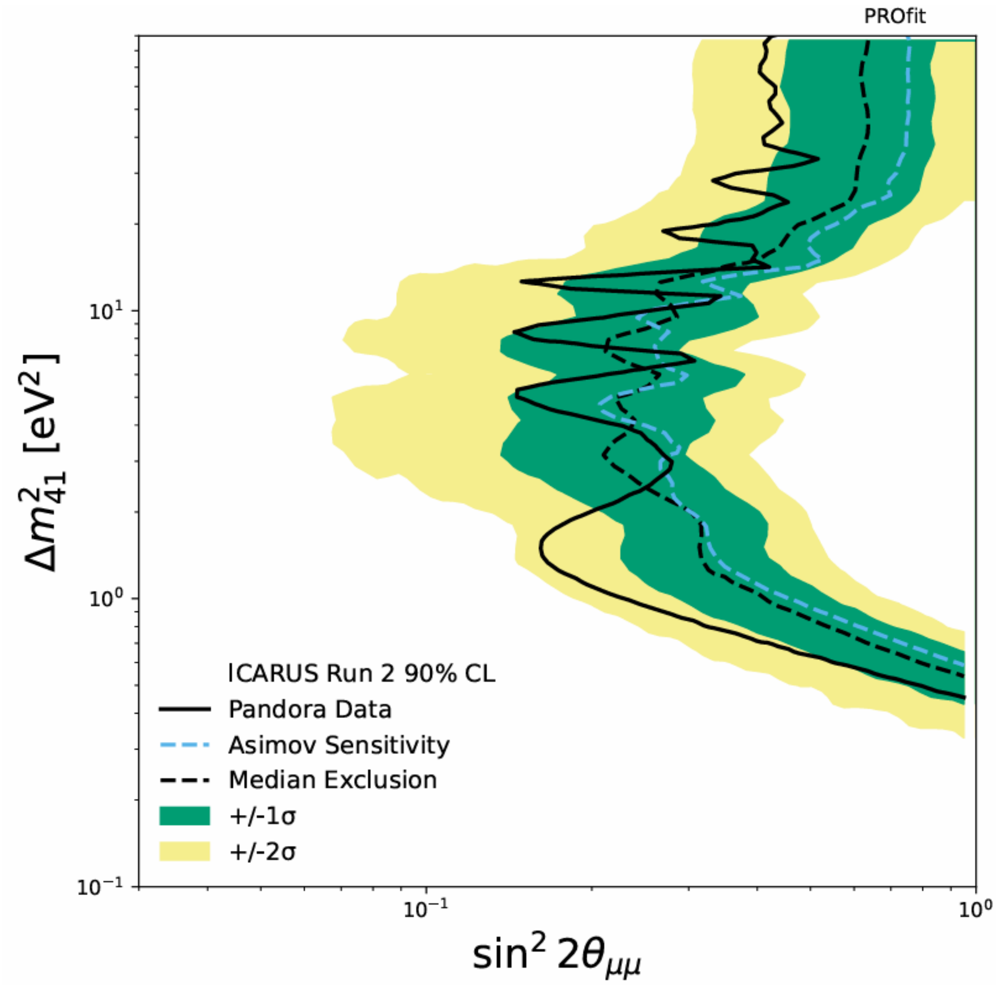

Figure 3: *An example sterile neutrino exclusion contour in the* $(\sin^2 2\theta, \Delta m^2_{41})$ *plane, of the kind PROfit is designed to produce. Figure reproduced from [5].*

3.2 Event selection

The standard for sterile MicroBooNE LArTPC search event selection has been Wire Cell[3]. Wire Cell is a 3D event reconstruction package for liquid argon time projection chambers, converting raw wire plane waveforms into fully reconstructed 3D event topologies. An example of what a signal looks like is shown in Fig. 4.

However, the analysis detailed here uses BAGELS (Broad Analysis of General ELectromagnetic Showers), a boosted decision tree (BDT) based event selection framework. BAGELS is designed to take reconstructed event information from multiple upstream reconstruction and selection frameworks including Wire Cell [3], gLEE[6], Pandora [2], and LANTERN[4] as its input, allowing it to combine information across reconstruction paradigms rather than relying on any single one.

3.3 Channel selection

PROfit takes in certain event typologies which are purely distinguished by weights which are distributed by the event selection used. These come in the form of root files which PROfit takes as the sole input. For a proper analysis many rootfiles are needed to account for uncertainties.

For the analysis carried out ν_μ charged current (ν_μCC) events were selected. These events are organized in a tree structure by final state topology: each event is classified by whether it contains zero or N reconstructed protons (the "0p" / "Np" splitting) and, independently, by whether it contains zero or N reconstructed neutral pions (the "$0\pi_0/N\pi_0$ splitting). An additional

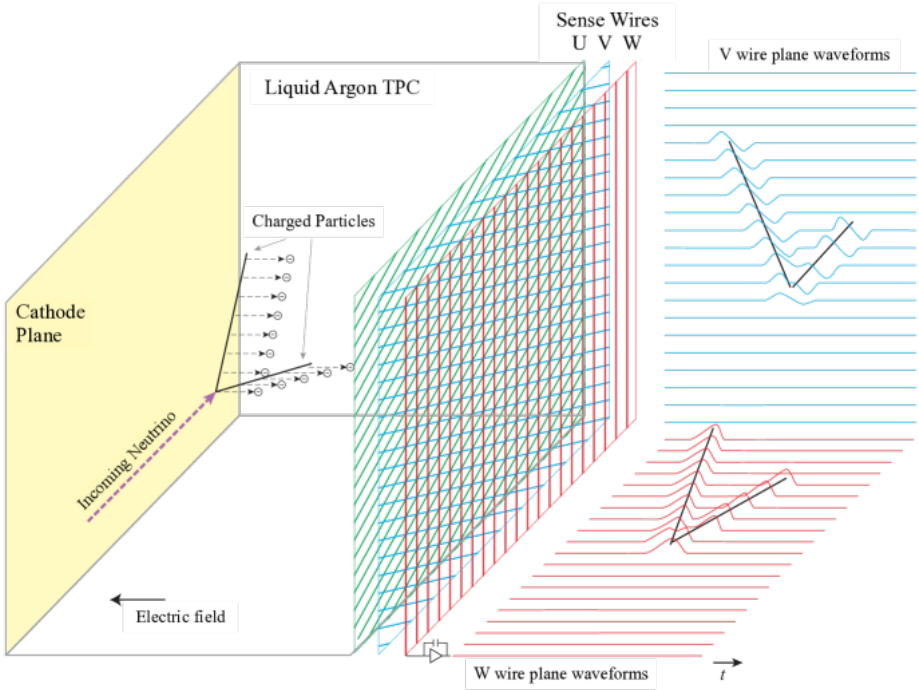

Figure 4: *Example of a neutrino event in a LArTPC with a creation of a charged particle which ionizes the Ar and leaves a "track' behind, which can be read out on the sense wires. Wire Cell takes these raw input voltages and outputs categories such as events to be used by analysis software such as PROfit. illustration from [1].*

geometric constraint separates events by containment: Fully Contained (FC), in which the entire neutrino interaction and its associated tracks are contained within the detector active volume, versus Partially Contained (PC), in which one or more tracks exit the volume. The Wire Cell selection standard uses a 7-channel combination of these topology/containment splits (nueCCFC, nueCCPC, numuCC0pi0PC, numuCC0pi0FC, numuCCpi0PC, numuCCpi0FC, NCpi0FCPC) as the baseline against which alternative BAGELS-based channel configurations are compared.

Exclusion contours are computed via a profiled X^2 statistic, evaluated against Wilks' theorem to obtain approximate confidence levels (typically 90% or 95%). For a given selection configuration, the fraction of oscillation parameter space excluded at a chosen confidence level serves as the primary figure of merit for comparing configurations.

3.4 Fitting and Minimization in PROfit

At its core, PROfit is a fitting and optimization framework built to handle the particular challenges that arise in neutrino oscillation searches. Fits are performed against one of three available goodness-of-fit metrics: the standard Pearson χ^2; the Combined Neyman–Pearson χ^2 [10], which is recommended for most analyses; and a Poisson log-likelihood. All three metrics are supported directly by the framework, and additional metrics can be incorporated without difficulty.

A central complication in short-baseline oscillation fits is that the χ^2 (or likelihood) surface in oscillation-parameter space is not always easily minimized. Because the relevant physics parameter, Δm^2, enters the oscillation probability through a sinusoidal term (Eq. 3), the fit surface is inherently harmonic: a good fit at a given Δm^2 is typically accompanied by comparably good fits at higher harmonics of that frequency, and small perturbations to the data or systematic parameters can shift which harmonic is favored. A purely local, gradient-based minimizer started from a single seed point is therefore prone to becoming trapped in a local minimum corresponding to the wrong harmonic, rather than converging on the true global best fit.

To address this, PROfit combines global and local optimization strategies in a staged approach:

1. **Latin Hypercube Sampling.** The parameter space is first explored broadly using Latin hypercube sampling, which provides uniform coverage of the space without the clustering that can occur under naive random sampling. This step generates a broad set of candidate starting points (3000 samples by default for a global fit).

2. **Particle Swarm Optimization.** A coarse global search is then performed using a particle swarm optimization technique, in which a population ("swarm") of candidate solutions is iteratively updated and converges toward the global minimum (250 iterations across 25 swarms, by default). This step is designed to locate the correct region of parameter space including the correct oscillation harmonic before any local refinement is attempted.

3. **L-BFGS-B Local Minimization.** Once a promising region has been identified, a gradient-based local minimizer (L-BFGS-B, via the LBFGS++ package) refines the solution to find a precise local minimum. L-BFGS-B is a quasi-Newton method that approximates the inverse Hessian using a limited history of past gradient evaluations, making it efficient for the high-dimensional parameter spaces (physics parameters plus systematic nuisance parameters) typical of these fits.

4. **Harmonic Seeding.** Because of the harmonic ambiguity described above, PROfit additionally performs a dedicated scan over Δm^2 once an initial global best fit is found, in order to identify other, comparably good local minima at different harmonics. These are retained as seed points for subsequent fits, helping to guard against a single fit result being an artifact of the minimizer settling into the wrong harmonic.

Because most nuisance parameters (representing systematic uncertainties) behave in an approximately Gaussian manner, this staged strategy concentrates most of its exploratory effort on the harmonic, oscillatory physics parameters, where global exploration is most needed, while relying on the efficient local L-BFGS-B step to refine the systematic-parameter fit once a good region of the physics parameter space has been identified.

4 ν_μ Disappearance Model Studies

The analysis was started out by looking at ν_μ disappearance events. The only events looked at were the previously mentioned $\nu_\mu CC$ events. The goal of this is to investigate and understand the current sterile exclusion space with BAGELS with respect to $sin^2(2\theta_{\mu e})$. The process of this involves trial and error, as I seek to understand what is the most affecting variable in the sensitivity. This is laid out in figure 5. A different Wire Cell channel selection was used due to the current limitation at the time in data available.

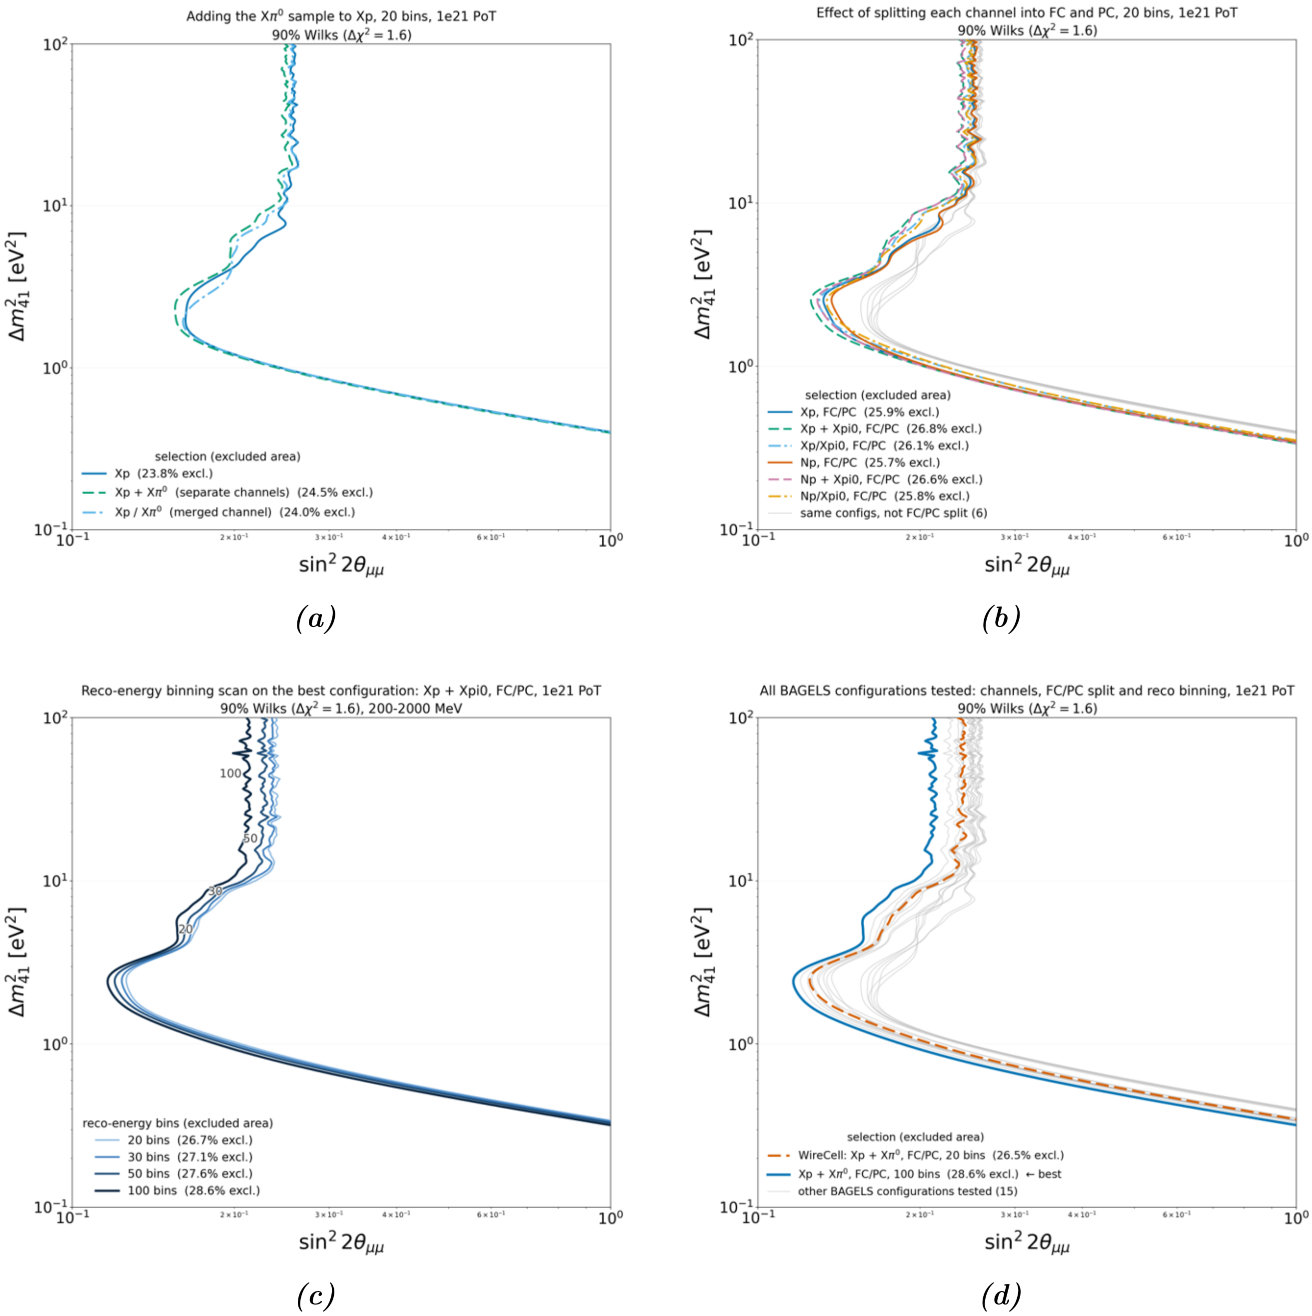

Figure 5: *Exclusion space plots A–D, **(a)** Showing the basic variation that mixing channels can generate in improving exclusion plots. In this case $Xp + X\pi_0$ shows the greatest sensitivity. **(b)** Demonstration of separation of channels into FC and PC selections. This decreases the uncertainty attributed to PC events, as they are now considered separately than FC. **(c)** Varying binning can alter the exclusion, here the general consensus is that more binning means greater exclusion. **(d)** comparison of most sensitive BAGELS selection versus Wire cell.*

5 Generating A Fully Oscillated Sample

To do a comprehensive sterile oscillation, one needs to consider a productive model to investigate. Many models exist for sterile neutrinos, in this work the 3+1 model was used. 3+1, meaning the 3 already existing flavors plus an additional sterile neutrino.

5.1 The Need for a ν_e sample

A complete 3+1 model search requires a simulated sample representing ν_μ neutrinos that have fully oscillated into ν_e by the time they reach the detector. This is refer to as "fully oscillated sample" because of the fact that all the ν_μ have "fully oscillated" into ν_e. Producing such a sample via dedicated simulation is computationally expensive, requiring significant compute time and resources that are not always readily available. This motivated the development of an approximate method for generating a full oscillation sample directly from existing, previously simulated ν_μ and ν_e samples.

5.2 Approximation of a ν_e Appearance

To save computational resources and time, the idea was to use an already existing ν_μ and ν_e data to conjure the fully oscillated sample. The core idea is to generate a ν_e appearance sample from the preexisting ν_μ sample while correctly retaining the systematic uncertainty structure appropriate to each stage of the neutrino's history. Because the neutrino beam originates as a ν_μ beam, the flux related systematic uncertainties (beam production, focusing, hadron production, etc.) must be inherited from the original ν_μ sample. However, once the neutrino has oscillated into a ν_e and interacts in the detector, the relevant detector and cross section related systematics are properly those of a ν_e interaction, not a ν_μ interaction.

To achieve this, events in the ν_e and ν_μ samples are matched by the nearest corresponding energy.(e.g., a ν_μ event at 1.100 GeV is matched to the ν_e event nearest it in true energy, such as 1.101 GeV). This matching is one-to-one, and functions on picking ν_μ-ν_e pairs with a allowed energy difference of $\Delta E_{\mu e} = 1MeV$, which is generous enough for energies of the neutrinos produced for MicroBooNE of $0 - 2GeV$. This process continues until there are no more pairs founded within the $1MeV$ limit. Because of the nature of picking ν_μ events, the systematics which determines how uncertainties are handled, can easily be picked from which one of the two samples are required. This allows each generated full oscillation event to carry the flux systematics of its matched ν_μ progenitor while retaining the detector and cross section systematics native to the ν_e sample it is built from. This can be demonstrated as PROfit outputs fractional uncertainties as shown in figure 8. The final fully oscillated (fullosc) sample is thus a hybrid construction: ν_μ flux systematics layered onto a ν_e like detector/interaction response.

This process took a ν_μ with a sample size of $\approx 700,000$ events and $\nu_e$size of $\approx 119,000$ events, and generated ν_{osc} with $\approx 76,000$ events. The code used to generate this full osc sample can be found at [12]. The final output of this is an exact copy of the input ν_e file with an additional weighting called "fullosc". This weighting determines whether the event is a identified ν_μ-ν_e pair, with 1 being yes and 0 being no.

5.3 Additional ν_e weighting

All I've done so far is recreate the appearing ν_e with replaced systematics and information from ν_μ, but to generate a fully oscillated sample, I have to match the size of the the ν_μ events. Which recalling were about $\approx 700,000$, while the fully oscillated pairs came out to be $\approx 76,000$. Meaning I have to add an additional weighting to normalize ν_{osc} to ν_μ. The process for this is much simpler as the pairs are already found, it is manner of applying a appropriate weighting to each pair to correctly match the ν_μ distribution. This came out to be:

$$W_{weight} = \frac{\rho_{\nu_{mu}}(E)}{\rho_{\nu_{osc}}(E)} \tag{6}$$

Where $\rho_{\nu_{mu}}$, is the ν_μ channel density and $\rho_{\nu_{osc}}$ is the ν_{osc} channel density both functions of the binned energy. A graphical demonstration of why this is necessary is shown 6

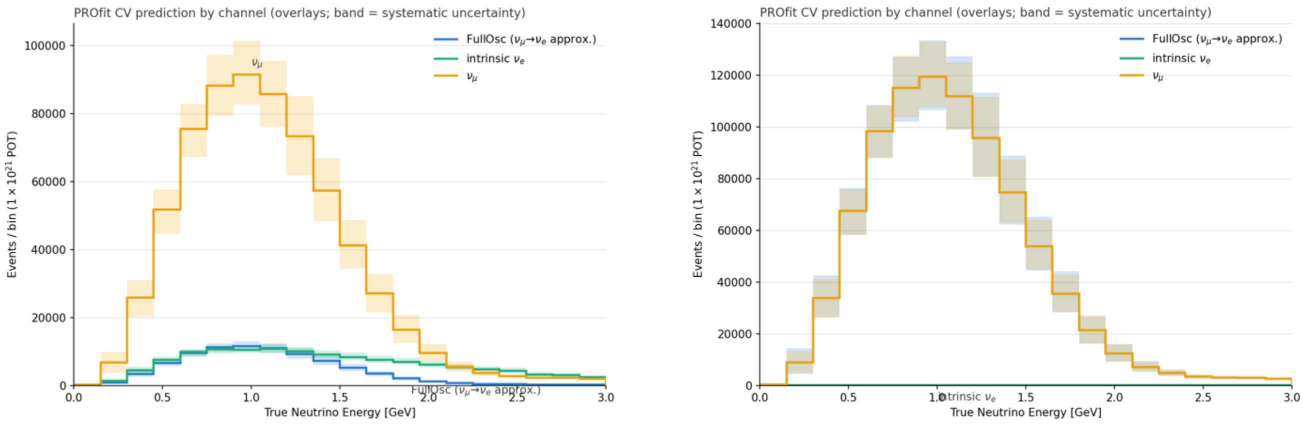

(a) *(b)*

Figure 6: ***(a)*** *The initial output of* ν_{os} *compared to* ν_μ *and* ν_e*, notice the immense discrepancy between* ν_μ *and* ν_{osc}*.* ***(b)*** *the application of the additional weighting to* ν_{osc}*, the normalizing of Protons on Target (PoT) causes additional events to be generated, as PROfit fits inputs to the chosen PoT. The invisible* ν_e *matches as it makes up* 0.05% *of the incoming neutrino flux.*

5.4 ν_e & ν_μ Cross Section Reweighting

An another concern is how cross section are handled for ν_{osc}. As ν_μ and ν_e couple equally to the Z and W bosons, but in kinematically they differ. A muon is fundamentally differing from a electron in the sole fact of its rest mass being $105MeV$ while a electron is $.511MeV$. This restricts the phase space because of the mass energy of the muon being vastly greater than that of an electron.

The weight W_{weight}, defined in Eq. 6, is constructed as a ratio of ν_μ to paired-ν_e *event rates* in bins of true neutrino energy. Because an event rate is proportional to the product $\Phi(E)\,\sigma(E)$, the flux and cross section are degenerate in that observable: matching the ν_μ event rate necessarily transfers the ν_μ flux and the ν_μ cross section together. The intrinsic ν_e sample already carried the correct σ_{ν_e}, and the reweighting overwrites it. Since a full oscillation sample requires $\Phi_{\nu_\mu}\sigma_{\nu_e}$ the flux of the ν_μ that oscillated, but the cross section of the ν_e that interacts the resulting template is in error by a factor $\sigma_{\nu_\mu}/\sigma_{\nu_e}$. This weight is than added to the ν_{osc} sample as "numu nue xs ratio weight".

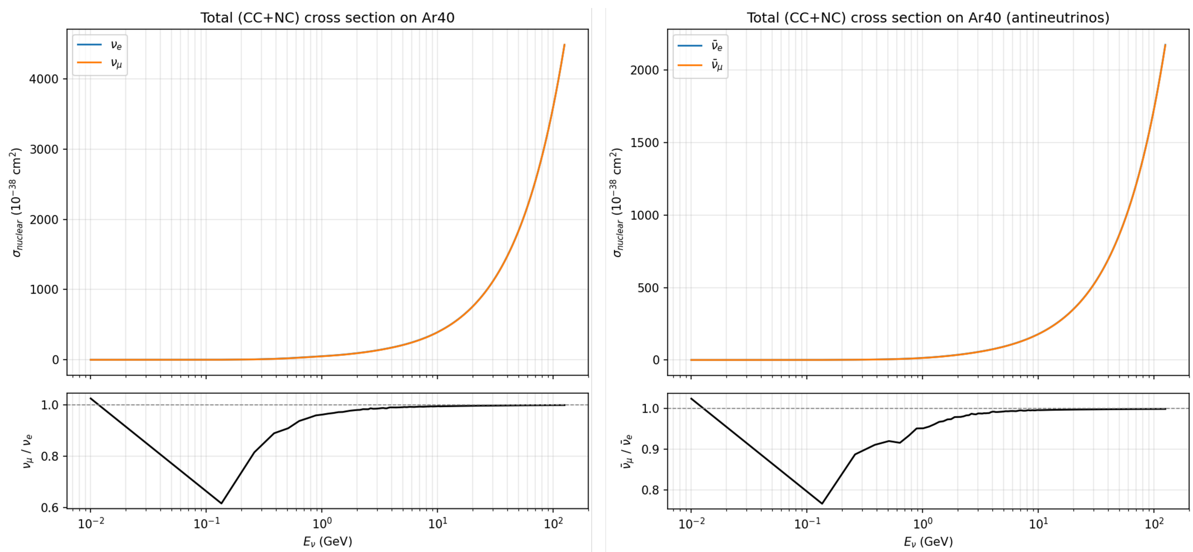

Figure 7: *Plotting Cross section for* ν_μ *&* ν_e *and its anti variants for Ar40.*

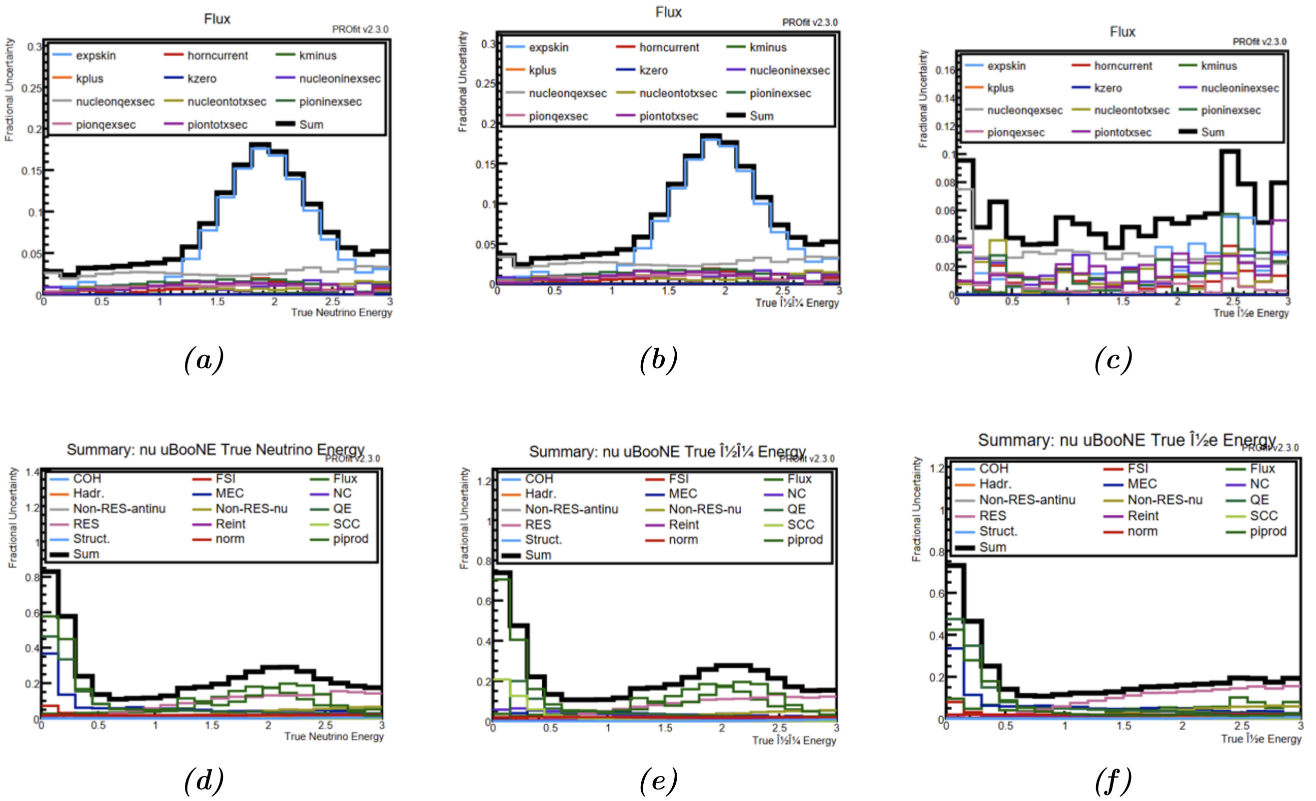

Figure 8: *fractionial uncertainties.* ***top row****, flux for* ν_μ,ν_e, ν_{osc} *respectively.* ***bottom row****, total uncertainties.*

The code for generating these weightings is found in [12].

6 3+1 Model Studies

Having access to fully oscillated sample allows ample comparison of Wire Cell and BAGELS selection. Recalling the Wire Cell selection, a 7-channel combination of these topology/containment splits (nueCCFC, nueCCPC, numuCC0pi0PC, numuCC0pi0FC, numuCCpi0PC, numuCCpi0FC, NCpi0FCPC). However there are many ways PROfit can represent a 3+1 model oscillation.

6.1 PROfit 3+1 Models

Although 3+1 refers to the 3 flavor plus an additional sterile, in PROfit there are various differing models for this behavior. Each are constructed as a way to investigate a parameter of sterile oscillation. As Rather than fitting the extended-PMNS mixing elements or angles directly (Section 2.3), PROfit also supports fitting directly in the physical, observable oscillation amplitudes themselves $\sin^2(2\theta_{ee})$, $\sin^2(2\theta_{\mu e})$, and $\sin^2(2\theta_{\mu\mu})$ with the remaining, unconstrained degree of freedom profiled over internally by the fit. This is convenient because it lets the analysis be presented directly in the variable of physical interest without a post-fit remapping step. Because each two-dimensional parameterization must resolve a degeneracy in mapping back to the underlying mixing elements, each of the three variants below additionally restricts to the small-mixing-angle regime, $\theta_{14}, \theta_{24} < 45°$.

6.1.1 3+1 Model 3A: $\sin^2(2\theta_{ee})$-Focused

This parameterization fits directly on the effective electron disappearance amplitude. The two fitting parameters are

$$\sin^2(2\theta_{ee}) \in [0,1] \quad \text{(effective electron disappearance amplitude)}, \tag{7}$$

$$r^A \in [0,1] \quad \text{(ratio parameter)}, \tag{8}$$

where r^A is defined as the ratio of the appearance to the electron disappearance amplitude,

$$r^A \equiv \frac{\sin^2(2\theta_{\mu e})}{\sin^2(2\theta_{ee})} = \frac{|U_{\mu 4}|^2}{1-|U_{e4}|^2} = \sin^2(\theta_{24}). \tag{9}$$

The remaining physical amplitudes follow as

$$\sin^2(2\theta_{\mu e}) = r^A \sin^2(2\theta_{ee}), \tag{10}$$

$$\sin^2(2\theta_{\mu\mu}) = 4\left(1-|U_{\mu 4}|^2\right)|U_{\mu 4}|^2, \quad \text{with} \quad |U_{\mu 4}|^2 = \frac{r^A}{2}\left(1+\sqrt{1-\sin^2(2\theta_{ee})}\right). \tag{11}$$

6.1.2 3+1 Model 3B: $\sin^2(2\theta_{\mu\mu})$-Focused

This parameterization is the symmetric counterpart to 3A, obtained by swapping the roles of U_{e4} and $U_{\mu 4}$. The two fitting parameters are

$$\sin^2(2\theta_{\mu\mu}) \in [0,1] \quad \text{(effective muon disappearance amplitude)}, \tag{12}$$

$$s^B \in [0,1] \quad \text{(ratio parameter)}, \tag{13}$$

where

$$s^B \equiv \frac{\sin^2(2\theta_{\mu e})}{\sin^2(2\theta_{\mu\mu})} = \frac{|U_{e4}|^2}{1-|U_{\mu 4}|^2}. \tag{14}$$

The remaining physical amplitudes follow as

$$\sin^2(2\theta_{ee}) = 4\left(1 - |U_{e4}|^2\right)|U_{e4}|^2, \quad \text{with} \quad |U_{e4}|^2 = \frac{s^B}{2}\left(1 + \sqrt{1 - \sin^2(2\theta_{\mu\mu})}\right), \tag{15}$$

$$\sin^2(2\theta_{\mu e}) = s^B \sin^2(2\theta_{\mu\mu}). \tag{16}$$

Because ν_μ disappearance is typically the best-measured of the three channels at MicroBooNE, Parameterization 3B is a natural choice when the disappearance result is the primary quantity of interest and ν_e disappearance can be largely neglected.

6.1.3 3+1 Model 3C: $\sin^2(2\theta_{\mu e})$-Focused

This parameterization instead fits directly on the effective appearance amplitude, making it the most natural choice when comparing directly to appearance-channel exclusion contours. The two fitting parameters are

$$\sin^2(2\theta_{\mu e}) \in [0, 1] \quad \text{(effective appearance amplitude)}, \tag{17}$$

$$w^C \in (0, \infty) \quad \text{(geometric mixing ratio)}, \tag{18}$$

where w^C is the geometric-mean ratio of electron to muon mixing,

$$w^C \equiv \sqrt{\frac{|U_{e4}|^2}{|U_{\mu 4}|^2}}. \tag{19}$$

$w^C = 1$ corresponds to equal electron and muon mixing strength, $w^C > 1$ to electron-dominated mixing, and $w^C < 1$ to muon-dominated mixing. The remaining physical amplitudes follow as

$$\sin^2(2\theta_{ee}) = 4|U_{e4}|^2\left(1 - |U_{e4}|^2\right), \quad \text{with} \quad |U_{e4}|^2 = \frac{w^C\sqrt{\sin^2(2\theta_{\mu e})}}{2}, \tag{20}$$

$$\sin^2(2\theta_{\mu\mu}) = 4|U_{\mu 4}|^2\left(1 - |U_{\mu 4}|^2\right), \quad \text{with} \quad |U_{\mu 4}|^2 = \frac{\sqrt{\sin^2(2\theta_{\mu e})}}{2w^C}. \tag{21}$$

The unitarity constraint $|U_{e4}|^2 + |U_{\mu 4}|^2 < 1$ translates, under this parameterization, to the bound

$$\frac{\sqrt{\sin^2(2\theta_{\mu e})}\left(w^C + 1/w^C\right)}{2} < 1. \tag{22}$$

Of the three amplitude-focused parameterizations, 3C carries the highest per-amplitude computational cost, owing to the additional division required to recover $|U_{\mu 4}|^2$ from.

6.2 First glimpse at 3+1 with Full Osc

The base is represented in the three 3+1 models shown in 9. This being the standard upon which the fruits of our efforts will be compared against. The idea being to improve the sensitivity with BAGELS. The ways one can manipulate BAGELS is similar to the case shown for ν_μ disappearance in terms of binning/channel selection, and splitting into FC/PC channels. All of this is present in the analysis carried out in 10.

Results being that, a maximal channel split, finer binning, and lower ν_e BDT score cut value all contribute to sensitivity. Efforts with affecting ν_μ score yielded no visible alteration, thus are not shown. Finally displaying this in a manner highlighting the improvement can be seen in 11. This clearly demonstrates a marginal improvement has occurred.

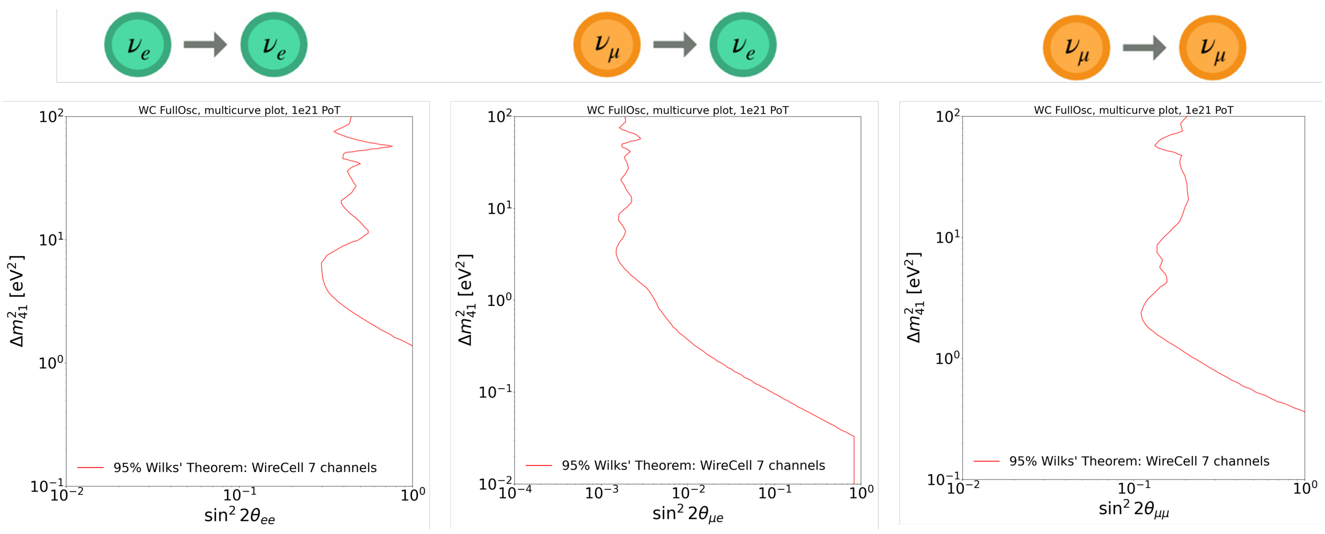

Figure 9: ***Left***, *Using PROfit model "3+1 3A", **Middle**: Using PROfit model "3+1 3C", **Right**: using PROfit model "3+1 3B". [5].*

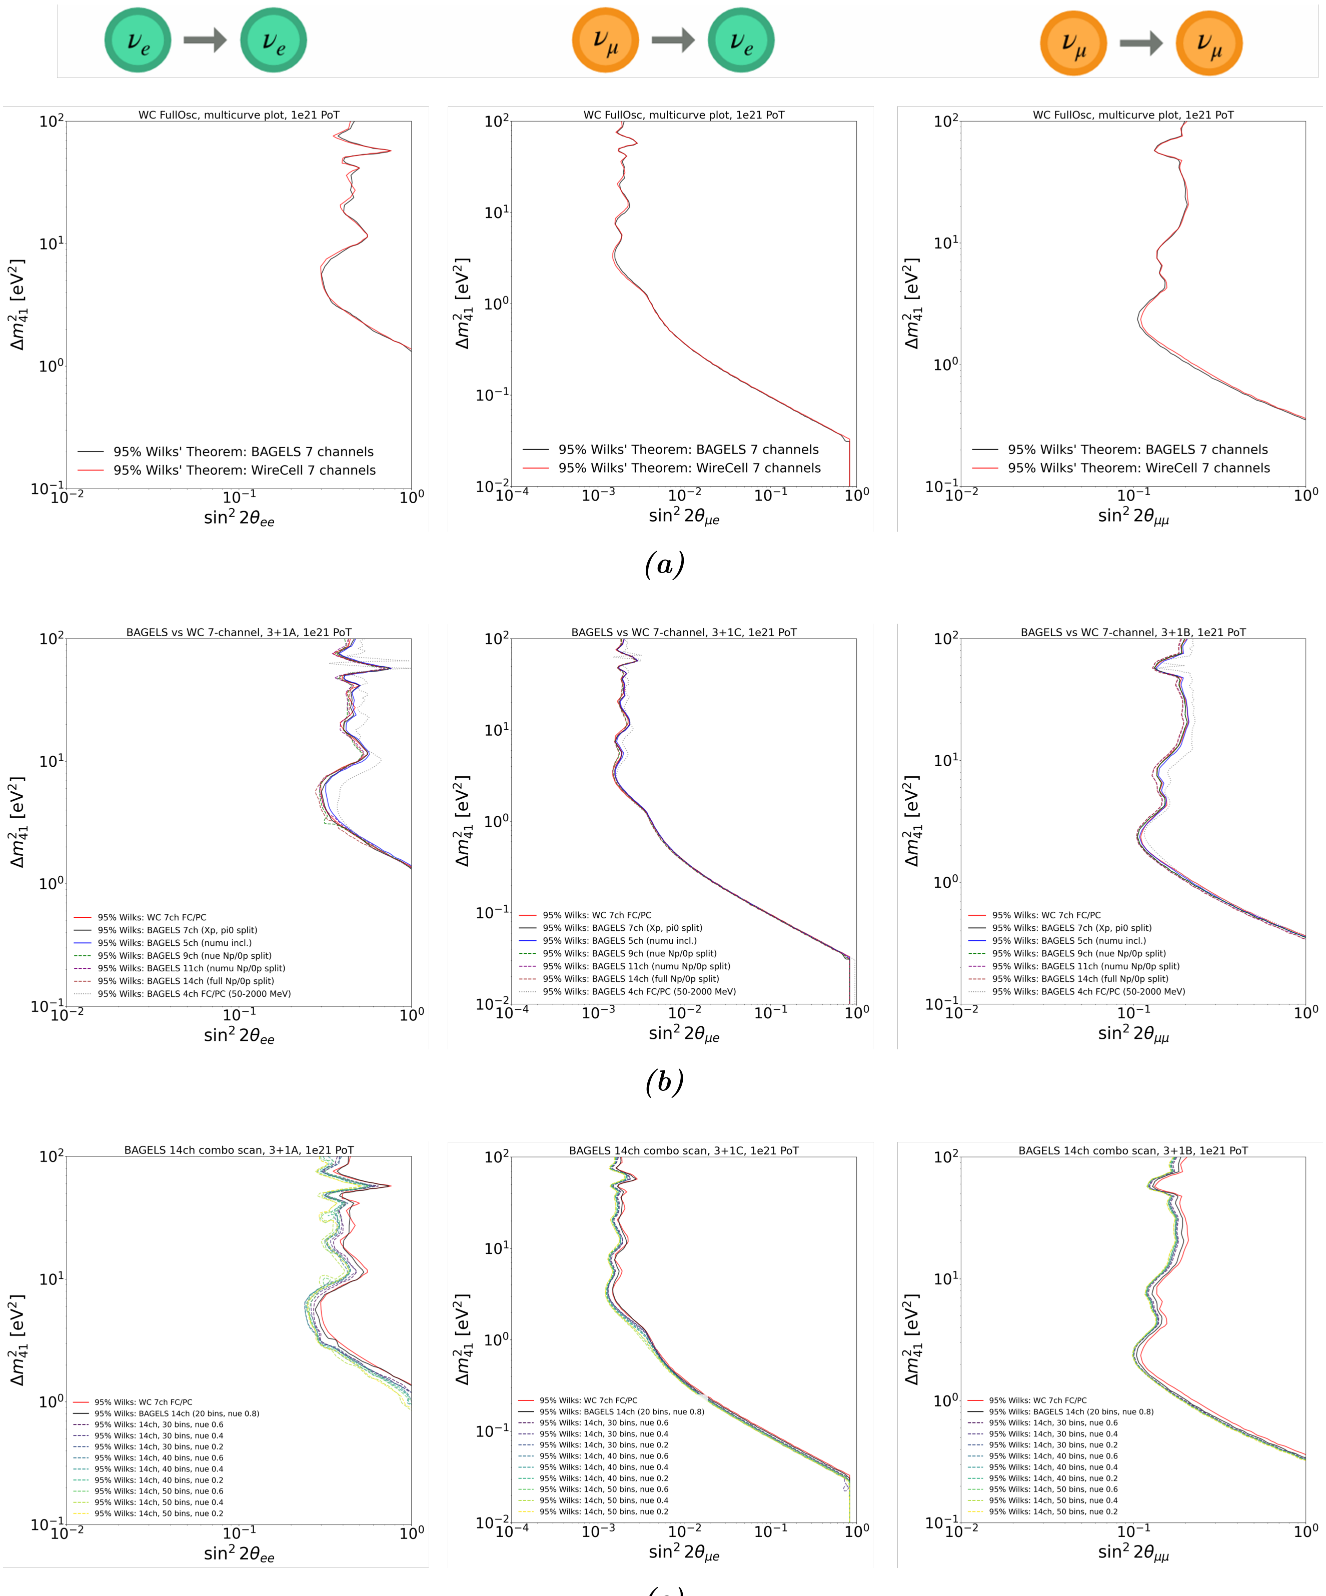

Figure 10: ***(a)*** *Comparison of Wire Cell with BAGELS of the same channel selection.* ***(b)*** *Splitting channels shows that maximal 14 channel improves the sensitivity the most.* ***(c)*** *Changing the nue BDT cut value, and increasing binning size with 14 channels shows improvement over wire cell*

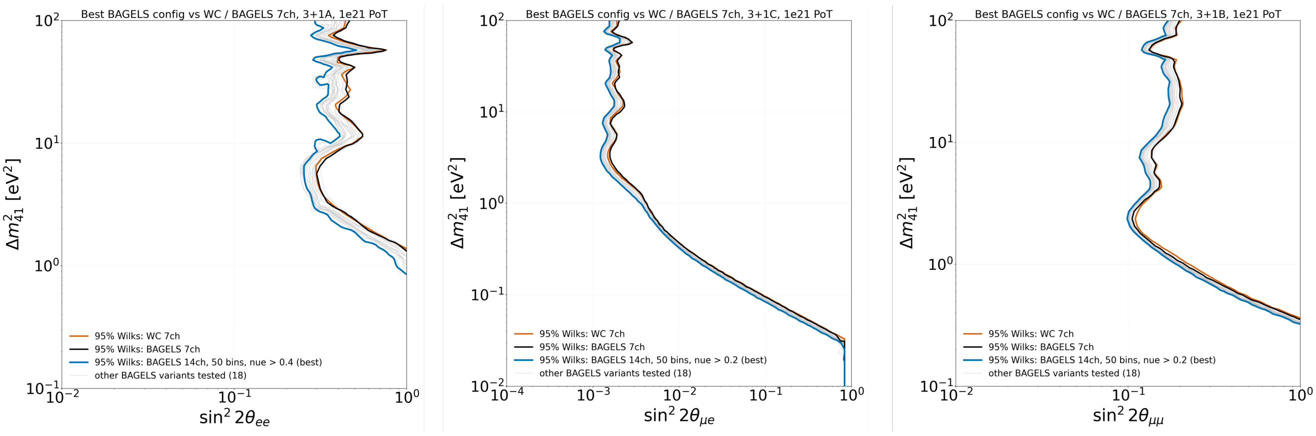

Figure 11: *The most sensitive channel selection of BAGELS compared to Wire-Cell. [5].*

7 Summary and Conclusions

This work showcases an explanation of the functionality of PROfit, and its ability to streamline the process of analysis for neutrino oscillation analysis. By manipulating channels, binning, and reconstruction categories. PROfit has immense capability in all things related to analysis. This was highlighted in the ν_μ disappearance studies.

Furthermore, a novel method of generating a fully oscillated ν_e appearance simulation sample was developed and studied. Leading to a marginal improvement over the Wire Cell selection, further work will have to be done in rigorously validating this approximation in the future, and developing a similar NuMI fully oscillated numu to nue appearance simulation using the same method, for use in BNB+NuMI oscillation analyses.

8 Acknowledgements

I would like to thank Lee Hagaman and Professor Mark Ross-Lonergan personally as without their efforts I wouldn't have been able to conduct the work done here. I also would like to thank Georgia Karagiorgi and Reshmi Mukherjee for their work on leading such a wonderful experience for me and my cohort. This material is based upon work supported by the National Science Foundation under Grant No. PHY-2349438.

References

[1] Design and construction of the MicroBooNE detector. *JINST*, 12:P02017, 2017.

[2] The Pandora multi-algorithm approach to automated pattern recognition of cosmic-ray muon and neutrino events in the MicroBooNE detector. *Eur. Phys. J. C*, 78:82, 2018.

[3] Wire-cell 3d pattern recognition techniques for neutrino event reconstruction in large LArTPCs: Algorithm description and quantitative evaluation with MicroBooNE simulation. *JINST*, 17:P01037, 2022.

[4] Reconstruction and selection of neutrino interactions in MicroBooNE using deep convolutional neural networks. Technical Report MICROBOONE-NOTE-1123-PUB, Fermi National Accelerator Laboratory, 2024.

[5] First search for sterile neutrino oscillation leading to ν_μ disappearance in the booster neutrino beam at ICARUS. 3 2026.

[6] P. Abratenko et al. Search for neutrino-induced neutral-current δ radiative decay in MicroBooNE and a first test of the MiniBooNE low energy excess under a single-photon hypothesis. *Phys. Rev. Lett.*, 128:111801, 2022.

[7] A. A. Aguilar-Arevalo et al. Significant excess of electronlike events in the MiniBooNE short-baseline neutrino experiment. *Phys. Rev. Lett.*, 121:221801, 2018.

[8] Q.R. Ahmad et al. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, 2002.

[9] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[10] Xiangpan Ji et al. Combined neyman–pearson chi-square: An improved approximation to the poisson likelihood ratio. 2019.

[11] Mark Ross-Lonergan and collaborators. PROfit: A framework for short-baseline neutrino oscillation fits. `https://github.com/markrosslonergan/Elephant_Vanishes.git`, 2026. Accessed: 2026-07-31.

[12] June Zey. Undimensional-Regularization: A ROOT macro for building a full-oscillation ($\nu_\mu \to \nu_e$) approximation sample for MicroBooNE. `https://github.com/juneszeyz/Undimensional-Regularization`, 2026. Accessed: 2026-07-31.