Details of the NOvA Oscillation Analyses

Erica Smith, for the NOvA Collaboration[*]

Indiana University

E-mail: ess3@iu.edu

NOvA is a long-baseline neutrino oscillation experiment that uses the NuMI beam from Fermilab. Two functionally identical detectors are placed off-axis, separated by 810 km. We study electron neutrino appearance and muon neutrino disappearance to probe unknown physics parameters such as the octant of θ_{23}, mass hierarchy, and CP violation. This document will discuss the details of these analyses, focusing on the latest results which include the full neutrino dataset and the first antineutrino dataset taken by the experiment.

The 20th International Workshop on Neutrinos (NuFact2018)
12-18 August 2018
Blacksburg, Virginia

---

[*]Speaker.

1. Introduction

NOvA is a long-baseline neutrino oscillation experiment that uses the NuMI beam from Fermi-lab. The NOvA detectors are functionally identical tracking calorimeters; the Near Detector (ND) measures the flavor content and energy spectrum of the beam before neutrino oscillations and the Far Detector (FD) measures the oscillated neutrino energy spectrum. The ND is 290 tons, located 1 km from the source, and has a 100 m overburden. The 14 kton FD is located on the surface 810 km from the source in Ash River, MN. The detectors are made of extruded PVC cells filled with liquid scintillator; the cells are assembled into planes and stacked in an alternating fashion to allow for 3D reconstruction of events. The detectors are placed 14.6 mrad off-axis of the beam, resulting in a narrow band beam that peaks at 2 GeV, putting NOvA just above the first oscillation maximum for electron neutrino appearance and above the visible energy of most neutral current (NC) backgrounds.

By measuring muon neutrino disappearance at the FD, we can make precision measurements of the mixing angle θ_{23} and Δm^2_{32}; observation of electron neutrino appearance allows us to probe the mass hierarchy and CP violation, and determine the octant of θ_{23}. These proceedings describe the event selection, energy reconstruction, and background estimation for the most recent charged current (CC) analyses which use a dataset of 8.85E20 protons on target in neutrino mode and 6.91E20 protons on target in antineutrino mode.

2. Event Selection

NOvA's ν_eCC and ν_μCC event selection is performed with a convolutional visual neural network (CVN) [1]. CVN is trained on the top and side 2D views of the event's calibrated hits. The network trains on each view separately for a given event and the last network layer combines the neural output from the previous stages in a conventional neural network to form an event classification. This method yielded an effective increase of 30% exposure from previous methods. To account for differences between neutrino and antineutrino events, we use separate trainings for each mode, increasing efficiency in antineutrino mode by up to 14%.

Two data-driven techniques are used to test the robustness of this method. The first, muon-removed electron-added (MRE), relies on the large statistics sample of ν_μCC events in the ND. We select ν_μCC candidates in data and Monte Carlo (MC), remove the muon, and add a simulated electron, leaving the hadronic portion of the event intact. We compare the ν_eCC event selection for the MRE procedure performed on data and MC. This test allows us to ensure the selection of ν_eCC events does not change with the hadronic portion of the event, and also reassures us that if there are differences in hadron reconstruction between data and MC, we will still select ν_eCC events properly. Selection efficiency in neutrino (antineutrino) mode agrees within 3% (2%).

The second data-driven technique used to test CVN is muon-removed bremsstrahlung, or MR-Brem. This method uses the large statistics sample of cosmic-ray muons at the FD. For cosmic-ray muons that produce bremsstrahlung, the muon is removed, leaving the electromagnetic shower. We expect the ν_eCC selector to respond to these electromagnetic showers in the same way in data and MC, and the selection efficiency for data and MC agrees within 2% for both neutrino and antineutrino modes.

3. Energy Reconstruction

3.1 ν_μ CC

In ν_μ CC events, neutrino energy is the sum of the muon energy, calculated by converting track length to energy, and the hadronic energy, calculated by converting visible calorimetric energy to true energy. Muon energy resolution is ~3% and hadronic energy resolution is ~30%. Oscillation sensitivity for ν_μ disappearance depends on the ability to see the deficit of events in the ν_μ energy spectrum at the FD; poor energy resolution causes events to feed into this region. To avoid degrading the spectrum measured using high resolution events by events with low energy resolution, we separate our high and low energy resolution events; because the hadronic energy resolution is much larger than muon energy resolution, we can estimate neutrino energy resolution with the fraction of hadronic energy in an event. We split the neutrino energy spectrum into quartiles based on the hadronic energy fraction and extrapolate each quartile to the FD separately. Although the primary motivation was to improve the energy spectrum reconstruction, this procedure also isolates the majority of the background in quartile 4, which improves the NC systematic impact by a factor of 2-4. Treating each quartile separately also allows for better extrapolation of cross sections [2].
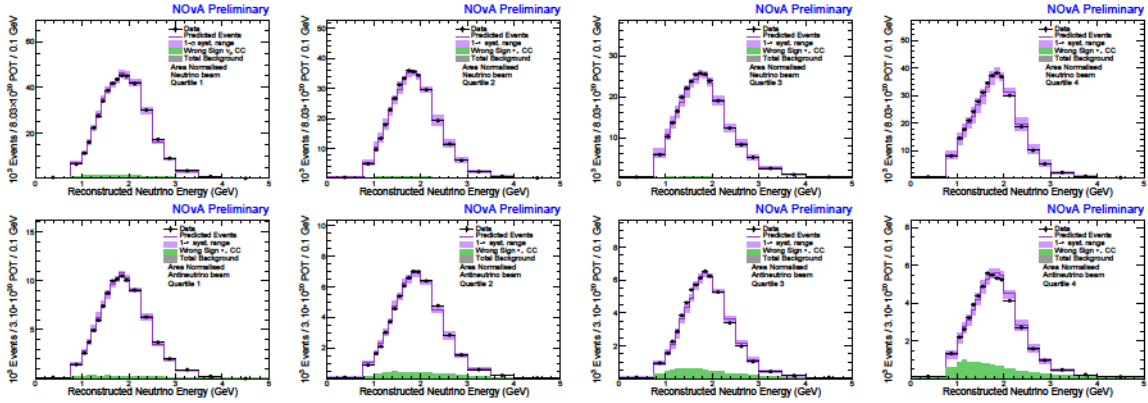

Figure 1: Neutrino energy spectrum for ν_μ-selected events in the ND, separated by quartile. (Far left) Quartile 1 has the lowest hadronic energy fraction and thus the highest energy resolution and (far right) quartile 4 has the highest hadronic energy fraction and lowest energy resolution, for (top) neutrino mode and (bottom) antineutrino mode. Data are shown in black, the purple lines show the area-normalized MC, the light purple regions show the shape-only systematic, and the green is wrong-sign contamination ($\overline{\nu}_\mu$ for neutrino beam and ν_μ for antineutrino beam).

3.2 ν_e CC

In ν_e CC events, neutrino energy is a combination of the electromagnetic energy and the hadronic energy. The calorimetric energy of each population is considered separately because the detector response is different for electromagnetic and hadronic energy. The separate components are identified with a version of CVN that functions as a particle identifier. This CVN, completely independent of the event selector, is trained on four images as input: the particle track and the full event for context, for the top and side detector views. Electromagnetic energy is taken to be

the calorimetric energy of all CVN-identified electrons and photons, and the hadronic energy is the difference between the total calorimetric energy of the event and the electromagnetic energy. The neutrino energy calculation is a quadratic fit to the electromagnetic and hadronic energies, $E = aE_{EM} + bE_{had} + cE_{EM}^2 + dE_{had}^2$. During training, we weight the events by a function that flattens the true energy spectrum that is implicit in the simulation, which minimizes the bias in energy reconstruction between 1-4 GeV. The resulting v_e energy resolution is 10.7% (8.8%) in neutrino (antineutrino) mode.

4. Background Estimation

Oscillation sensitivity for v_e appearance depends on separating signal from background. To achieve this separation, we bin the v_e energy spectrum in three particle ID (PID) bins, two core and one peripheral, and nine energy bins; PID binning separates the sample by purity, and energy binning separates appeared v_e from beam v_e. Core events are those that pass preselection and cosmic rejection cuts at the FD. Peripheral events fail either pre-selection or cosmic rejection but pass a tight selection cut and a secondary cosmic rejection; no energy reconstruction is performed for these events.
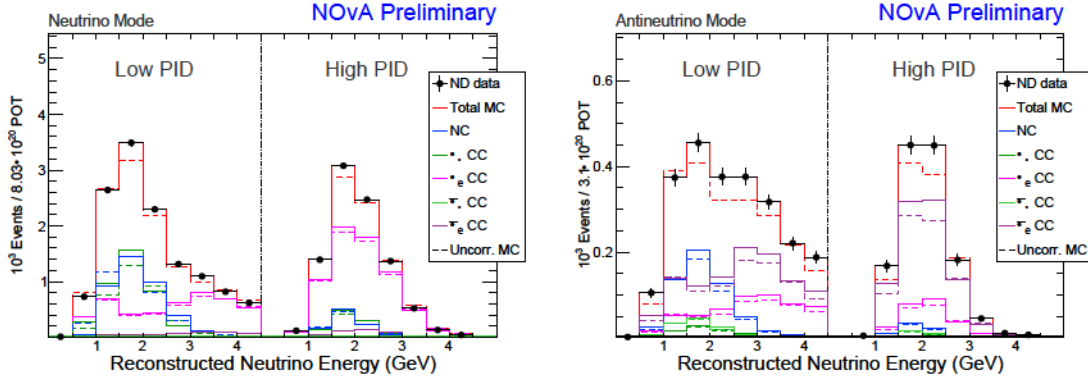

Figure 2: Neutrino energy for v_eCC-like events in the ND in (left) neutrino mode and (right) antineutrino mode for high and low PID bins, broken down by contribution. Data are in black; uncorrected MC is indicated by dashed lines and corrected MC by solid lines. Red is total MC, blue is NC, dark green is v_μCC, light purple is v_eCC, light green is $\overline{v}_\mu$CC, and dark purple is $\overline{v}_e$CC. MC is scaled to match data; in neutrino mode, background contributions are scaled by two data-driven techniques, while in antineutrino mode, background contributions are scaled proportionally.

At the ND, all v_e-selected events are background as no oscillations have occurred yet. This v_e-like population is used to constrain backgrounds for v_e appearance at the FD. Because each background has different origins and will extrapolate differently to the FD, we normalize the MC to match data in a way that takes individual background contributions into account. A first order solution is to scale up each background by the same factor until data and MC agree, which is the method used for antineutrino mode. For the neutrino mode, we constrain backgrounds using two data-driven techniques using ND data. The first technique decomposes intrinsic beam v_e flux into contributions from its hadronic parents. Pion decays mostly result in lower energy neutrinos while

kaon decays mostly result in higher energy neutrinos. Because ν_es and ν_μs have the same parents, we can use contained ν_μ data to constrain the pion flux, and higher energy uncontained ν_μ events to constrain the kaon flux. We see a deficit in the MC for the ν_μ energy spectrum for ν_μs from a pion parent and an excess for ν_μs from a kaon parent, so we scale the contributions from pion parents up and from kaon parents down; this scaling is then applied to the pion and kaon contributions to the beam ν_e flux. The second technique constrains NC backgrounds by looking at Michel electron multiplicity. Due to the pion decay chain we expect to have an enhancement of ν_eCC events when there are no Michel electrons, and a higher number of NC events as the multiplicity increases. For a given number of Michel electrons, we find the ratio between NC and CC events, then scale the MC to the data based on this ratio. Both methods are done separately for every PID and energy bin.

5. Summary

NOvA uses a convolutional visual neural network (CVN) for event selection, which gives an effective 30% increase in exposure over traditional methods. The robustness of this event selection is tested with two data-driven techniques, which show a data/MC agreement to within a few percent. Muon energy is calculated from the track length, and the hadronic energy is calculated from the visible calorimetric energy; neutrino energy is taken to be the sum of these quantities. For ν_eCC events the electromagnetic and hadronic populations are selected using a version of CVN that serves as a PID, and the total energy is a quadratic sum of the calorimetric energies of these populations. Backgrounds for the ν_eCC appearance at the FD, such as intrinsic beam ν_es and NC events, are estimated using data from the ND and extrapolated to the FD separately.

References

[1] A. Aurisano *et al.*, JINST **11**, P09001 (2016).

[2] J. Wolcott [NOvA Collaboration], PoS NUFACT **2018**, 098 (2018).