

An Altarelli Cocktail for the MiniBooNE Anomaly?

Vedran Brdar^{1,2, a} and Joachim Kopp^{3,4, b}

¹*Fermi National Accelerator Laboratory, Batavia, IL 60510, USA*

²*Northwestern University, Dept. of Physics & Astronomy, Evanston, IL 60208, USA*

³*Theoretical Physics Department, CERN, 1211 Geneva 23, Switzerland*

⁴*Johannes Gutenberg University Mainz, 55099 Mainz, Germany*

We critically examine a number of theoretical uncertainties affecting the MiniBooNE short-baseline neutrino oscillation experiment in an attempt to better understand the observed excess of electron-like events. We re-examine the impact of fake charged current quasi-elastic (CCQE) events, the background due to neutral current π^0 production, and the single-photon background. For all processes, we compare the predictions of different event generators (GENIE, GiBUU, NUANCE and NuWro) and, for GENIE, of different tunes. Where MiniBooNE uses data-driven background predictions, we discuss the uncertainties affecting the relation between the signal sample and the control sample. In the case of the single-photon background, we emphasize the large uncertainties in the radiative branching ratios of heavy hadronic resonances. We find that not even a combination of uncertainties in different channels adding up unfavorably (an ‘‘Altarelli cocktail’’) appears to be sufficient to resolve the MiniBooNE anomaly. Varying the radiative branching ratios of the $\Delta(1232)$ and $N(1440)$ resonances by $\pm 2\sigma$, however, reduces its significance from 4σ to less than 3σ . We finally investigate how modified background predictions affect the fit of a $3 + 1$ sterile neutrino scenario. We carefully account for full four-flavor oscillations not only in the signal, but also in the background and control samples. We emphasize that because of the strong correlation between MiniBooNE’s ν_e and ν_μ samples, a sterile neutrino mixing only with ν_μ is sufficient to explain the anomaly, even though the well-known tension with external constraints on ν_μ disappearance persists.

CONTENTS

1	Introduction	2
2	Background estimates from Monte Carlo simulations	3
2.1	Charged Current Events	5
2.2	Neutral Current π^0 Production	6
2.3	Neutral Current Single γ Production	8
3	Data-Driven Background Estimates	8
3.1	Neutral Current π^0 Production	9
3.2	Neutral Current Single γ Production	10
4	Uncertainties in the Radiative Branching Ratios of Heavy Baryonic Resonances	11
5	Impact on Sterile Neutrino Fits	12
5.1	2-Flavor vs. 4-Flavor Fits to MiniBooNE Data	12
5.2	Dependence on Background Predictions	14
6	Summary and Conclusions	18
A	Parameter Scans in the $3 + 1$ Sterile Neutrino Model	19
	References	19

^a vedran.brdar@northwestern.edu

^b jkopp@cern.ch

1 INTRODUCTION

The decision to downgrade the BooNE proposal to the MiniBooNE experiment [1] has been, in retrospect, both a curse and a blessing for neutrino physics. On the one hand, MiniBooNE has given us one of the most intriguing anomalies particle physics has seen in recent years: a 4.8σ [2] excess of electron neutrinos (ν_e) in a beam consisting mostly of muon neutrinos (ν_μ). This observation has led to significant progress in our understanding of neutrino–nucleus interactions [3], progress that will be invaluable to future neutrino experiments. The anomaly has also given rise to a tremendous amount of theoretical and phenomenological work interpreting the excess as a hint for new physics, for instance in the form of sterile neutrinos [4–14]. On the other hand, if MiniBooNE hadn’t been stripped of its second detector, we might have known right away whether the anomaly is due to “new physics” or due to imperfect modeling of Standard Model effects.

In any case, the situation is being rectified now, with Fermilab’s new short-baseline program consisting of not one but three additional detectors: MicroBooNE [15], SBND [16], and ICARUS [17, 18]. These detectors are located at different baselines, L , from the primary target and should therefore be able to unambiguously determine whether MiniBooNE’s ν_e excess oscillates with L or not. Moreover, they are liquid argon time projection chambers which, compared to MiniBooNE’s mineral oil-based Čerenkov detector, offer much better event reconstruction capabilities and will therefore be much better at distinguishing a possible neutrino oscillation signal from various backgrounds.

Our goal in this paper is to add several novel aspects to the discussion of background processes and theoretical uncertainties in MiniBooNE. Ultimately, we would like to determine whether an accumulation of small deviations in different background channels adding up in an inauspicious way – often dubbed an “Altarelli cocktail” [19] – could be sufficient to explain the MiniBooNE anomaly.

We begin in section 2 by re-calculating several of MiniBooNE’s most important backgrounds. First, in section 2.1, we address MiniBooNE’s event reconstruction: the signal process charged current quasi-elastic (CCQE) neutrino–nucleus scattering is identified by the exclusive presence of a single e^\pm or μ^\pm , and the energy of the incoming neutrino is calculated from the energy of this charged lepton and its direction with respect to the beam axis. However, events may be incorrectly classified as CCQE if additional final state particles such as pions are either reabsorbed before they leave the target nucleus, or are missed by the detector. The resulting misreconstruction of neutrino energies has been discussed previously in refs. [20–34], and while it leads to distortions of neutrino energy spectra, the effect has been found to be too small to explain the MiniBooNE anomaly. Our novel contribution compared to previous works will be threefold: (i) we compare predictions of different event generators, namely GENIE, GiBUU, NUANCE, and NuWro, to better estimate how CCQE energy reconstruction depends on theory errors; (ii) we include the impact of fake CCQE events in the ν_μ sample, which MiniBooNE analyze together with the ν_e sample to better constrain the neutrino flux; (iii) we work with more up-to-date data than previous studies, in particular the data from ref. [34].¹

Second, in section 2.2 we will study MiniBooNE’s π^0 background, comparing again the predictions of different event generators (see ref. [35] for previous related work in this direction). The π^0 background arises from neutral current (NC) interactions in which a single π^0 is produced. To the MiniBooNE detector, the photons from π^0 decay look the same as e^\pm from a charged current (CC) ν_e or $\bar{\nu}_e$ interaction. Therefore, if one of the two photons is missed, or if the two are so close to each other that they merge into a single reconstructed photon, NC π^0 production can mimic CCQE ν_e interactions, the signal MiniBooNE is looking for.

The third background we address in this paper is the single-photon background (section 2.3).

¹ We do not consider the even more recent data from the 2020 update of the MiniBooNE anomaly [2], which corresponds to a roughly 50% further increase in statistics and elevates the significance of the anomaly from 4.7σ to 4.8σ .

Single photons can originate from radiative decays of hadronic resonances like the $\Delta(1232)$, or from coherent production off the target nucleus. As for the π^0 -induced background, a single photon can mimic the CCQE ν_e signal. We compare predictions of different event generators and tunes to estimate the theoretical uncertainties affecting the single-photon background.

In the second part of the paper, section 3, we shift our focus towards data-driven estimates for the π^0 s and single photons in an attempt to more closely follow the approach the MiniBooNE collaboration is taking in predicting backgrounds. For both the π^0 background and the single-photon background, a suitable control sample are single π^0 events in which the two photons from the decay are separately reconstructed. On the one hand, this control sample constrains the rate of π^0 production. But since at MiniBooNE energies most π^0 s stem from the decay of hadronic resonances, it also constrains the production rate of such resonances and thus rate of radiative resonance decay events. Even with data-driven background estimates, theoretical uncertainties enter when translating the event rate in the control sample into a number of expected background events in the signal region. To estimate these theoretical uncertainties, we develop a mock-up of MiniBooNE's data-driven π^0 and single-photon analyses, anchoring these background rates to the measured spectrum of π^0 events and comparing the impact of different theoretical models and different event generators/tunes on the translation between the control and signal samples.²

In section 4, we emphasize the impact of large uncertainties in the radiative branching ratios of hadronic resonances. We will find that these uncertainties can significantly affect the predictions of the single-photon background. Importantly, this remains true even when data-driven methods are used.

In the final part of the paper, section 5, we fit a 3 + 1 sterile neutrino scenario (3 standard active neutrinos and 1 additional eV-scale sterile neutrino, ν_s) to MiniBooNE data. We first emphasize that in a full four-flavor fit, the ν_e background can be affected by significant $\nu_e \rightarrow \nu_s$ disappearance, and the ν_μ control sample that is used for flux normalization can suffer from sizeable $\nu_\mu \rightarrow \nu_s$ disappearance. In a two-flavor fit, on the other hand, disappearance effects are negligible. We will show how this disparity affects the preferred parameter regions of the 3 + 1 scenario. We then investigate how the fit changes depending on which event-generator is used for the background predictions, and on whether the background prediction is taken directly from the Monte Carlo (as in section 2) or whether data-driven methods are used (as in section 3). We summarize and conclude in section 6.

2 BACKGROUND ESTIMATES FROM MONTE CARLO SIMULATIONS

We begin by individually considering various background processes relevant to MiniBooNE's sample of CCQE ν_e -like events. We focus in particular on the CC ν_e background due to the ν_e contamination in the beam (section 2.1), NC π^0 production (section 2.2), and NC single-photon production (section 2.3). Even though the MiniBooNE collaboration is not relying on Monte Carlo simulations alone, but rather on data-driven background estimates wherever possible, a comparison of Monte Carlo-only predictions will give a first indication of where large theoretical uncertainties are lurking. We employ in particular the following event generators: GENIE v3.00.04 [39], NUANCE v3.000 [40], NuWro v19.02.2-35-g03c3382 [41], and GiBUU (2019 release) [42].

While GENIE, NuWro, and GiBUU are actively used state-of-the-art tools, NUANCE is, to the best of our knowledge, not under active development any more. Nevertheless, NUANCE will be crucial for our analysis because it is the main generator used by the MiniBooNE collaboration [43].

² MiniBooNE's data-driven estimates of the π^0 background have also been scrutinized recently in refs. [36, 37], focusing in particular on the effect of π^0 re-absorption (an effect that has also been included in MiniBooNE's analyses [38]).

Generator	Tune	Ref.	Comments
NUANCE	–	[40]	the generator used by MiniBooNE
GiBUU	–	[42]	theory-driven generator
NuWro	–	[41]	
GENIE	G18_01a_02_11a	[39, 44]	GENIE baseline tune; see [44] for naming conventions
	G18_01b_02_11a		different FSI implementation compared to G18_01a_02_11a
	G18_02a_02_11a		updated res./coh. scattering models compared to G18_01a_02_11a
	G18_02b_02_11a		updated res./coh. scattering models and different FSI
	G18_10a_02_11a		theory-driven configuration; similar to G18_02a
	G18_10b_02_11a		theory-driven configuration; similar to G18_02b

TABLE I. Event generators and tunes used in this work.

Indeed, for NUANCE, we work with flux and configuration files that were kindly provided to us by the MiniBooNE collaboration and are dated April/May 2007.

GiBUU differs from the other three generators in that it employs a more holistic approach: rather than piecing together largely independent theoretical models for different kinematic regimes (quasi-elastic scattering, resonance production, deep-inelastic scattering, etc.) and subprocesses (primary interaction, final state interactions, etc.), it uses the same inputs such as nuclear ground state, nuclear potentials, and production/absorption amplitudes for all kinematic regimes. By solving a set of quantum transport equations, one of its strengths is the accurate simulation of final state interactions (FSI).

In the case of GENIE, we will consider six different tunes [44]. The naming convention for these tunes is **G18_XXy_02_11a**, where tunes with **XX=01** can be considered baseline tunes, those with **XX=02** feature updated implementations of resonant and coherent scattering, and those with **XX=10** also employ updated models for CCQE and two-particle/two-hole (2p2h) interactions [21] as well as an improved description of the nuclear initial state in terms of a local Fermi gas (with radius-dependent Fermi momentum, as opposed to a relativistic Fermi gas with a Fermi momentum that is the same everywhere in the nucleus). The lower-case letter **y** indicates how FSI are treated, with **y=a** corresponding to a simple implementation of hadron–nucleus cross sections, while **y=b** stands for a more sophisticated hadronic cascade in which interactions of hadrons with individual nucleons are recursively simulated. The code **02_11a**, finally, describes the data sets that the models have been tuned to, which are the same for all tunes considered here.

Note that GiBUU, NUANCE, and NuWro do not implement radiative decays of heavy baryonic resonances (e.g. $\Delta(1232) \rightarrow N + \gamma$) by default. As these decays are an important source of single-photon events in MiniBooNE and thus an important background to the ν_e appearance search, we have implemented them manually by randomly replacing the pion in $\Delta(1232) \rightarrow N + \pi^0$ events by a photon with the same energy. We do this for 0.6% of all $\Delta(1232) \rightarrow N + \pi^0$ events, corresponding to the branching ratio of $\Delta(1232) \rightarrow N + \gamma$ according to ref. [45].

The different generators and tunes used in this work are also summarized in table I.

Our strategy is the same for each of the three considered background channels (CC neutrino scattering, NC π^0 production, and NC single-photon production), and can be described as follows:

- (i) From a Monte Carlo simulation using the NUANCE generator, we predict the event sample under consideration. In doing so, we make our best effort to reproduce the cuts and implement the efficiency factors of the real MiniBooNE analysis (which also employed the NUANCE generator).
- (ii) The predicted event spectrum from (i) is then compared with the corresponding prediction obtained by the MiniBooNE collaboration [34]; the differences, which are expected to be

mild, are compensated by bin-by-bin tuning.

- (iii) We then predict the same event sample using GiBUU, NuWro, as well as six different GENIE tunes, using the same cuts and efficiency factors as for NUANCE. We then apply the tuning factors determined in step (ii) as the ratio between our NUANCE prediction and MiniBooNE’s. This final tuning greatly alleviates any residual differences between our simplified analysis and the one employed by the MiniBooNE collaboration, yielding background predictions that are accurate enough to compare to data in a meaningful way.

In our analysis we only consider positive horn polarity (neutrino mode) data which mostly drives the statistical significance of the reported excess..

2.1 Charged Current Events

We start by considering CC neutrino interactions. To MiniBooNE’s $\nu_\mu \rightarrow \nu_e$ oscillation search, such interactions are relevant not only for the signal, but also for part of the background. This is because the beam, though consisting mostly of muon neutrinos, unavoidably contains a small admixture of electron neutrinos, mostly from the decays of kaons and muons.³ This intrinsic ν_e background accounts for $\mathcal{O}(10\%)$ of the total background at the lowest measurable neutrino energies $E_\nu \sim 200$ MeV, and for almost all background events at $E_\nu > 1$ GeV. On top of this, the sterile neutrino fit includes also CC ν_μ events, which are used as a control sample to normalize the flux. A change in the CC ν_μ rate will thus indirectly affect predictions for the intrinsic ν_e background and for the $\nu_\mu \rightarrow \nu_e$ signal.

Following the strategy introduced in the beginning of this section, we first compute the expected rate of CC ν_e and CC ν_μ events using NUANCE [46]. Out of all simulated events, we keep those that contain exactly one charged lepton (electron or muon) and no detectable mesons. We define a “detectable” meson as a neutral pion, a charged pion above the Čerenkov threshold, or a meson heavier than a pion. For CC ν_e -like events, we apply a 20% detection efficiency [34], while for CC ν_μ events, the efficiency is assumed to be 35% [47].

Like the MiniBooNE collaboration, we reconstruct the neutrino energy E_ν in each event based on the assumption that the event topology is indeed $\nu_{e,\mu} + n \rightarrow e^-/\mu^- + p$ (or the corresponding processes for anti-neutrinos). In this case, E_ν can be calculated as [43]

$$E_\nu = \frac{2m'_n E_\ell - (m_n'^2 + m_\ell^2 - m_p^2)}{2[m'_n - E_\ell + \sqrt{E_\ell^2 - m_\ell^2} \cos \theta_\ell]}, \quad (1)$$

where E_ℓ is the charged lepton’s energy ($\ell = e, \mu$), m_ℓ is its mass, and θ_ℓ is the direction of its momentum vector relative to the beam axis. The proton and neutron masses are denoted as m_n and m_p , respectively, while $m'_n \equiv m_n - E_B$, with E_B the binding energy in the nucleus. We set $E_B = 34$ MeV, corresponding to neutrons bound in a ^{12}C nucleus [48]. It is important to keep in mind that eq. (1) will yield an incorrect value for E_ν in fake CCQE events, that is events which contain extra final state particles (for instance pions), but in which these extra final state particles are missed, either because they are reabsorbed by the nucleus in which they are produced, or because they fall below the experimental thresholds [28, 33, 49].

We compare in fig. 1 our predicted E_ν spectra (colored histograms) to the ones used by MiniBooNE [34] (gray histograms). For the case of ν_μ interactions (which are observed essentially without backgrounds), we also compare to data (black points with error bars in fig. 1). Focusing

³ Here, and in the following, “neutrino” refers to both neutrinos and anti-neutrinos, unless stated otherwise.

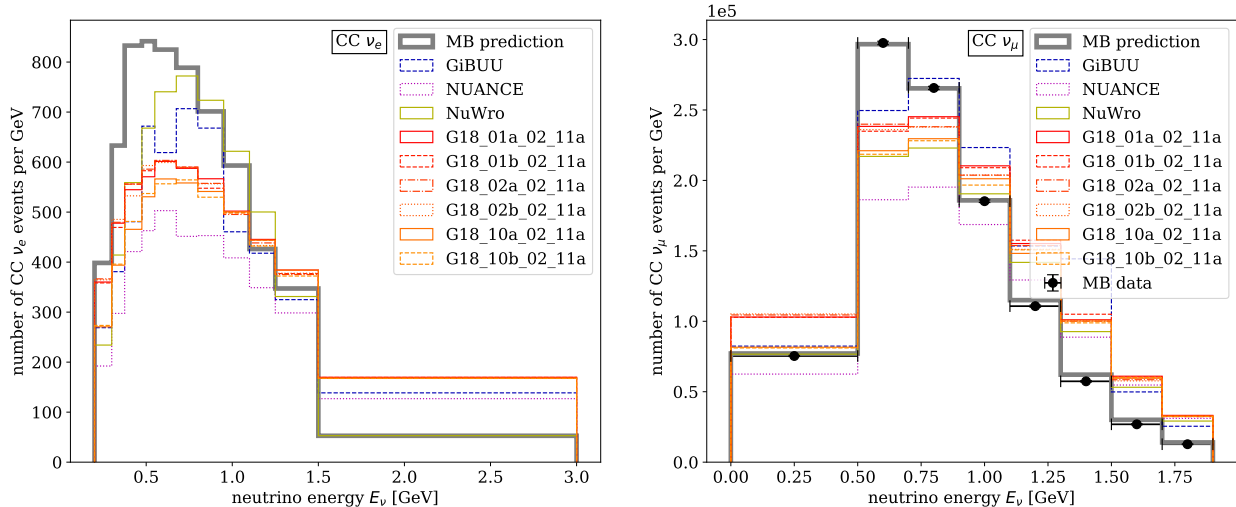


FIG. 1. (a) Monte Carlo-only predictions for the CC ν_e background to MiniBooNE’s ν_e appearance search from different Monte Carlo event generators, in particular GiBUU (blue dashed), NUANCE (purple dotted), NuWro (yellow solid), and GENIE (orange/red) with different tunes as explained in table I. The solid gray histogram corresponds to the official background prediction by the MiniBooNE collaboration. (b) Monte Carlo-only predictions for the CC ν_μ events that are used for flux normalization. We show results for the same generators and tunes as in (a), but we also compare to MiniBooNE’s data (black points with error bars).

first on the differences between different event generators and tunes, we observe that predictions vary by $\mathcal{O}(10\%)$. One striking observation is that NuWro predicts relatively large CC ν_e rates compared to the other generators, while for CC ν_μ interactions, its predictions are among the lowest. GiBUU’s predictions are overall relatively large, which is a reflection of the well-known fact that GiBUU predicts lower pion production rates than observed in MiniBooNE (while being consistent with the pion production rates in MINER ν A and T2K) [50–52]. Here, this deficit means that a larger number of CC interactions will be identified as CCQE, and fewer will be vetoed because of the presence of extra pions. Regarding the comparison between our predictions and MiniBooNE’s there are certain discrepancies; namely, MiniBooNE predicts somewhat higher event rates compared to us both in the ν_e channel, and their predicted spectrum is more peaked in the ν_μ channel. This indicates that our simplified cuts do not fully capture MiniBooNE’s true acceptance and efficiency. As discussed above, for the purpose of the sterile neutrino fits which we will present in section 5, we will eliminate this bias by applying additional energy-dependent tuning factors which are obtained as the ratio of MiniBooNE’s prediction and our prediction using NUANCE. That way, we ensure that using the same generator as MiniBooNE, our predictions exactly match theirs. After this tuning, the differences between our predictions expose the differences between event generators while being fairly robust with respect to the simplifications of our analysis.

2.2 Neutral Current π^0 Production

Neutral pions are frequently produced in neutrino interactions. Of particular concern to MiniBooNE’s ν_e appearance search are neutral current interactions of the form $\nu + N \rightarrow \nu + N + \pi^0$. In this case, the π^0 , or rather the two photons into which it promptly decays, are the only visible interaction products. If one of the photons leaves the fiducial volume before showering, or if the laboratory frame opening angle between the two photons is small, the event will contain a single

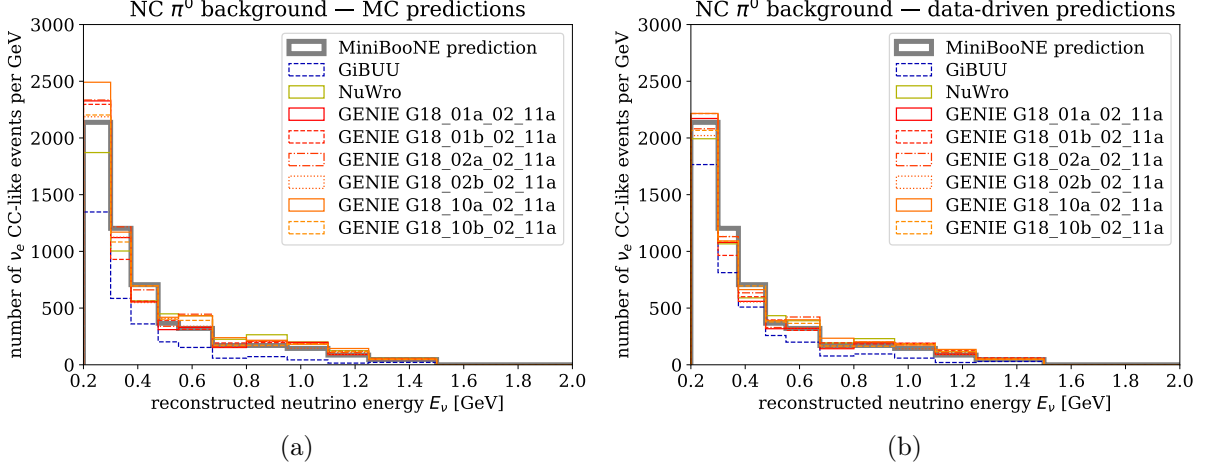


FIG. 2. (a) Monte Carlo-only predictions and (b) data-driven predictions for the NC π^0 background to MiniBooNE’s ν_e appearance search from different Monte Carlo event generators, in particular GiBUU (blue dashed), NuWro (yellow solid), and GENIE (orange/red) with different tunes as explained in table I. We do not show NUANCE results here because they are used for tuning our analysis and thus would not be independent, The solid gray histogram corresponds to the official background prediction by the MiniBooNE collaboration.

electromagnetic shower that can be mistaken for an e^\pm from a CC ν_e or $\bar{\nu}_e$ interaction.

To predict the contribution of NC π^0 events to the ν_e background, we proceed as follows. First, out of all simulated neutrino events, we select those which have one or several π^0 ’s in the final states, no e^\pm or μ^\pm , and no other charged particles above the Čerenkov threshold. We then generate the photons from π^0 decay, and we apply Gaussian energy and angular smearing to their 4-momentum vectors. We use a 10° angular resolution, and an energy resolution given by $(\Delta_E/E)^2 = [0.08\sqrt{\text{GeV}/E}]^2 + [0.024/(E/\text{GeV})]^2$ [53].

On average, photons propagate around 50 cm before converting to an e^+e^- pair and starting an electromagnetic shower [54]. Therefore, we pick the point at which this happens by randomly drawing from an exponential distribution with a mean of 50 cm. Photons converting outside the detector volume are discarded. If a photon converts in the veto region outside the fiducial volume ($r_{\text{fiducial}} = 574.6$ cm [55]), but still inside the active volume ($r_{\text{veto}} = 610$ cm), the whole event is vetoed. The remaining events are assigned a weight factor according to MiniBooNE’s e/γ efficiencies, published together with ref. [56].

If there are two or more photons left in an event, we need to determine whether they can be reconstructed separately or if they merge into one. We do so by applying a cut on the opening angle ϕ between pairs of photons. If ϕ is below a threshold ϕ_{thr} , the two photons are merged into one. If not, they are kept separate, and the algorithm continues to consider the next pair of photons. If, at the end of this procedure, exactly one photon is left, the event is considered as a fake ν_e event, contributing to the background in the $\nu_\mu \rightarrow \nu_e$ oscillation search. Otherwise, the event is discarded. We allow the threshold ϕ_{thr} to depend on the reconstructed neutrino energy E_{reco} that would be assigned to the event according to eq. (1) if the photon-induced electromagnetic shower was mis-interpreted as originating from an electron and the event was mis-reconstructed as a CCQE ν_e interaction, with all photons merged into one. In each E_{reco} bin, we choose $\phi_{\text{thr}}(E_{\text{reco}})$ such that our prediction for the NC π^0 background using the NUANCE generator agrees exactly with MiniBooNE’s prediction (likewise based on NUANCE) in that channel. We find that the resulting $\phi_{\text{thr}}(E_{\text{reco}})$ decreases with energy. With this procedure, our NUANCE prediction by construction matches exactly MiniBooNE’s, while our GENIE, NuWro, and GiBUU predictions differ from it,

highlighting the discrepancies between generators.

Our results for the NC π^0 background prediction are shown in fig. 2 (a). We see that differences between NUANCE, NuWro, and GENIE (with any tune) are small, while the GiBUU prediction is significantly lower. As mentioned already in section 2.1, this discrepancy between GiBUU’s predictions and MiniBooNE data on single-pion production is well known [50–52], but given that a similar discrepancy does not exist when comparing GiBUU to MINER ν A and T2K data, it is an open question whether it indicates a problem on the theory side or on the experimental side. Taking the GiBUU predictions at face value would even increase the significance of MiniBooNE’s low-E ν_e excess.

2.3 Neutral Current Single γ Production

We next discuss single photon events. Most of these arise from radiative decays of heavy hadronic resonances created in NC neutrino interactions, which is why this background is referred to as the $\Delta \rightarrow N\gamma$ background in MiniBooNE publications.⁴

We select simulated events that contain exactly one photon in the final state, no electrons or muons, and no other charged particles above the Čerenkov threshold. We apply the same energy and angular smearing as in section 2.2. Each event is then (mis)reconstructed as a CC ν_e interaction, misinterpreting the photon as an electron and applying eq. (1) to determine the reconstructed neutrino energy E_{reco} . We finally determine an E_{reco} -dependent reconstruction efficiency factor by demanding that, in each E_{reco} bin, our NUANCE prediction matches MiniBooNE’s prediction for the $\Delta \rightarrow N\gamma$ background. We find efficiency factors of order 10–20%, not too different from the e^\pm/γ efficiencies from the supplemental material of ref. [56].

The comparison between event generators is shown in fig. 3 (a). We find that GiBUU’s, NuWro’s, and GENIE’s single-photon spectra are 10–20% lower than MiniBooNE’s official, NUANCE-based, background prediction. The older GENIE tunes (G18_01a_02_11a and G18_01b_02_11a) seem to give the lowest single-photon yield.

In passing, let us note that for GENIE, we found an unusually large number of events with two photons from the decays of an η resonance. Other generators predict far fewer such events, a discrepancy which might be due to differences in hadronization models. Luckily, we find that the probability for missing one of the two photons from η decay is small. Therefore, even in the GENIE simulation, photons from η decay account only for a handful of background events in the ν_e appearance search.

3 DATA-DRIVEN BACKGROUND ESTIMATES

While our comparison of different Monte Carlo predictions for the MiniBooNE backgrounds in section 2 reveals important discrepancies between the various generators, MiniBooNE’s own background prediction is to some extent resilient to these discrepancies. This is because it is based on data-driven techniques, meaning that certain crucial aspects, for instance the π^0 or $\Delta(1232)$ production rate are directly measured in control samples rather than being predicted theoretically. However, theoretical input is still needed for translating measurements in the control regions to background predictions for the signal region. Therefore, even data-driven background estimation techniques are not fully immune to theoretical uncertainties. In the following, we will investigate these uncertainties for the π^0 and single photon backgrounds in MiniBooNE.

⁴ Recently, a novel but subdominant single photon production channel was presented in [57].

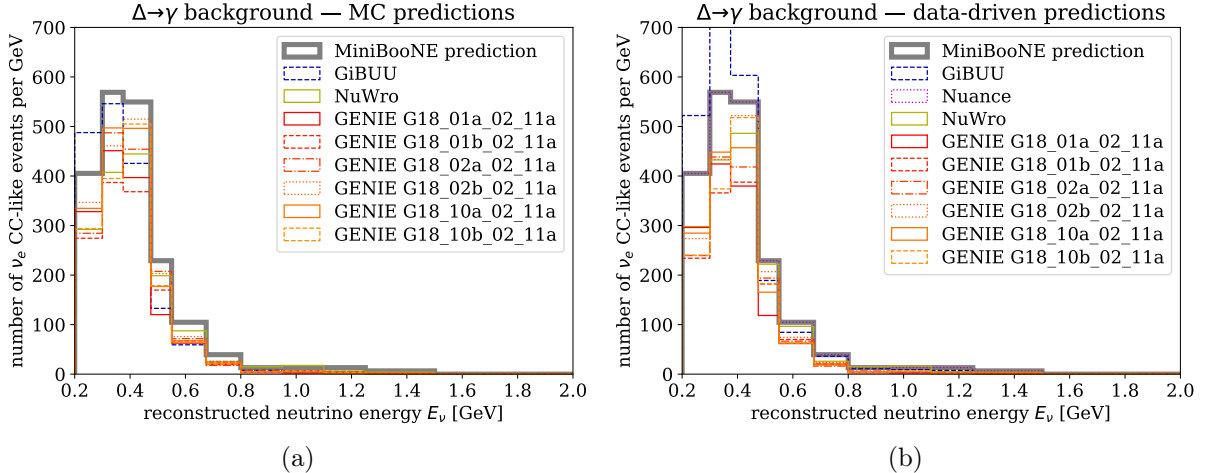


FIG. 3. (a) Monte Carlo-only predictions and (b) data-driven predictions for the NC single-photon background to MiniBooNE’s ν_e appearance search from different Monte Carlo event generators, in particular GiBUU (blue dashed), NUANCE (purple dotted), NuWro (yellow solid), and GENIE (orange/red) with different tunes as explained in table I. MiniBooNE’s own prediction for this background is shown as a thick gray histogram. We do not show Monte Carlo-only predictions from NUANCE because, given the way our analysis is designed, they agree exactly with MiniBooNE’s prediction.

3.1 Neutral Current π^0 Production

To emulate the data-driven estimation technique for the π^0 background, we select from our simulated events those containing exactly one π^0 . Other than that, events need to satisfy the same criteria as in section 2.2: no charged leptons (e^\pm or μ^\pm) and no charged particles above the Čerenkov threshold are allowed. We assume that the smearing kernel for pions is the same as for photons, see section 2.2.

We first compare the rate of π^0 production as a function of the π^0 momentum, p_{π^0} , to MiniBooNE’s measurement of the π^0 spectrum published in ref. [58]. To do so, we apply the efficiency factors given in fig. 5 of that reference. The result of the comparison is shown in fig. 4, where MiniBooNE’s prediction for the background (mostly mis-identified CC events, events with π^\pm , and multi-pion events) has been subtracted from the data. We see that Monte Carlo predictions vary by $\mathcal{O}(50\%)$, with NUANCE (purple), NuWro (yellow), and GENIE (orange/red) matching the data well, but not perfectly, while the event rate predicted by GiBUU is too low and the spectrum is too soft. In other words, we observe again the well-known discrepancy mentioned in sections 2.1 and 2.2 between GiBUU’s predictions and MiniBooNE data on single-pion production [50–52]. We also note that MiniBooNE’s own Monte Carlo prediction is on the low side, showing that even with a full detector simulation, the data on NC π^0 production is very difficult to understand theoretically.

To obtain a data-driven prediction for the π^0 background to the ν_e appearance search, we extract a signal sample of π^0 s faking a ν_e from the simulation using the same criteria as in section 2.2. The signal sample is then binned in p_{π^0} using the same bins as in the control sample, namely the 11 bins shown in fig. 4. Next, we reweight the signal sample in each of these bins with the ratio of observed to simulated single- π^0 events in the control sample. The thus reweighted signal sample is what we call our data-driven prediction.

The result of (mis)reconstructing the data-driven background sample as CCQE ν_e interactions is shown in fig. 2 (b) for the different event generators and tunes. We observe that the spread in our results is reduced compared to the Monte Carlo-only predictions in fig. 2 (a). In particular, the discrepancy between the GiBUU-based prediction and MiniBooNE’s own prediction is not quite as

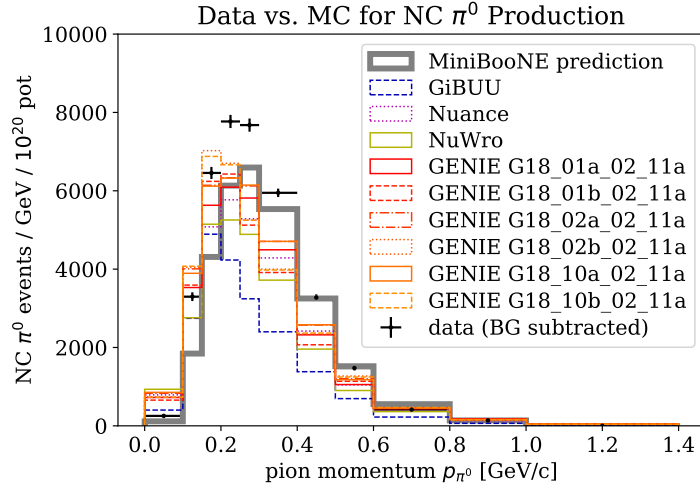


FIG. 4. Comparison of MiniBooNE’s data on NC π^0 production from ref. [58] to predictions from various Monte Carlo event generators as a function of the pion momentum p_{π^0} . The color scheme used here is the same as in figs. 2 and 3. MiniBooNE’s own prediction for the π^0 spectrum is shown as a thick gray histogram.

bad in panel (b) as in panel (a). Among the other generators, the spread is $\lesssim 10\%$.

3.2 Neutral Current Single γ Production

The procedure we follow to determine the impact of different Monte Carlo generators on the data-driven prediction for the single-photon background is very similar to the one employed in the case of the π^0 background above. The control sample consists once again of single- π^0 events, given that the neutral $\Delta(1232)$ resonance, which is responsible for most of the single-photon background, predominantly decays to pions. The signal sample is extracted from the simulation using the same criteria as in section 2.3. Both the π^0 control sample and the single-photon signal sample are then binned according to the π^0 and photon momentum, respectively, using the binning from fig. 4. Next, we reweight the simulated single-photon events in each of these bins with the ratio of observed to simulated single- π^0 events in the control sample. Finally, we (mis-)reconstruct the reweighted single-photon events as CC ν_e interactions to obtain our data-driven background prediction to the ν_e appearance search, binned in E_{reco} .

The result is shown in fig. 3 (b). In contrast to what we observed in section 3.1 for the π^0 background, we now find that the data-driven technique does *not* significantly decrease the spread between predictions from different generators. This indicates that, for the single-photon background, large theoretical uncertainties exist which are not related to the $\Delta(1232)$ production cross-section. Instead, they are due to other sources of single-photon events, not related to the Δ resonance, such as coherent photon production, decays of heavier resonances, etc. Reducing these uncertainties may be possible with more sophisticated data-driven methods separating different sources of single-photon events, and identifying suitable control samples for each of them. However, even if this was possible, we expect significant cross-contamination between the different control samples. For instance, it will be well-nigh impossible to fully separate different heavy baryonic resonances based only on their visible decay products. Therefore, any significant reduction in systematic uncertainties in the single photon channel will be very challenging. Our analysis thus shows that

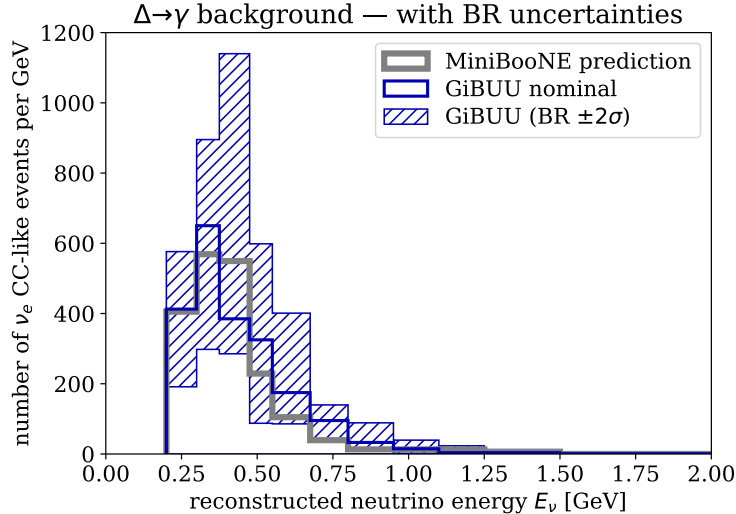


FIG. 5. Impact of branching ratio uncertainties on GiBUU’s prediction for the single-photon background in MiniBooNE. We have varied the radiative branching ratios of the $\Delta(1232)$ and $N(1440)$ baryon resonances within their respective 2σ confidence intervals [45], and we show the envelope of the resulting event spectra (blue hatched histogram). For comparison, we also show MiniBooNE’s NUANCE-based prediction for this background [58] (gray histogram).

systematic uncertainties in the single-photon channel could play an important role in understanding the MiniBooNE ν_e appearance anomaly.

4 UNCERTAINTIES IN THE RADIATIVE BRANCHING RATIOS OF HEAVY BARYONIC RESONANCES

We would finally like to emphasize an additional aspect that could contribute to the MiniBooNE anomaly: the branching ratios for radiative decays of heavy baryonic resonances ($\Delta(1232)$, $N(1440)$, etc.) have relatively large uncertainties. As an example, ref. [45] quotes $\text{BR}(\Delta(1232) \rightarrow N\gamma) = 0.55\%–0.65\%$, $\text{BR}(p(1440) \rightarrow p\gamma) = 0.035\%–0.048\%$, and $\text{BR}(n(1440) \rightarrow n\gamma) = 0.02\%–0.04\%$. None of these decays has been measured directly, and the quoted branching ratios are instead inferred from baryon–photon interaction amplitudes determined in pion–nucleon and photon–nucleon scattering, see for instance ref. [59]. In spite of the large uncertainties, state-of-the-art neutrino Monte Carlo generators work with fixed branching ratios, so the uncertainties need to be carefully accounted for a posteriori.⁵

To study the impact of these uncertainties, we have carried out a set of 25 GiBUU runs in which the radiative branching ratios of the $\Delta(1232)$ and $N(1440)$ resonances were scanned on a regular grid covering their 2σ uncertainty intervals from ref. [45]. We have analyzed the resulting event samples using the methods described in section 2.3. In fig. 5, we plot the envelope of the 25 single-photon event spectra as a function of the (mis-)reconstructed would-be neutrino energy (blue hatched region). We see that the event rate varies by a factor of 2–3, especially in the low-energy bins where MiniBooNE observes its anomaly.

We conclude that branching ratio uncertainties are crucial for making reliable background predictions in a MiniBooNE-like experiment, and may make an important contribution to the total

⁵ We also note that decay data in Monte Carlo generators is not always based on the latest version of the Particle Data Group’s compilation, which may add to the error on these data.

error budget. Crucially, they cannot be removed by using data-driven techniques – relating a π^0 control sample to the signal sample as in section 3.2 hinges on the knowledge of the radiative branching ratios.

To summarize our findings so far, we have identified a number of differences between different Monte Carlo predictions of the MiniBooNE backgrounds. Visual inspection of figs. 1 to 3 and 5 suggests that these discrepancies may alleviate the tension with the MiniBooNE ν_e data, but will probably not be large enough to fully explain away the observed event excess. In the following, we will quantify this statement by carrying out explicit fits in a $3 + 1$ sterile neutrino model and determining the confidence level at which the no oscillation hypothesis is excluded.

5 IMPACT ON STERILE NEUTRINO FITS

In fitting the MiniBooNE data, we will focus on the $3 + 1$ scenario, in which the Standard Model is extended by a single sterile neutrino whose mass is assumed to be at the eV scale. The only new interaction is thus a Yukawa coupling of the form

$$\mathcal{L} \supset y (i\sigma^2 H^*) LN, \quad (2)$$

where L is a Standard Model lepton doublet, N is the sterile neutrino field, H is the Standard Model Higgs doublet, σ^2 is the second Pauli matrix, and y is a dimensionless coupling constant. All fermion fields are interpreted as Weyl spinors here. In general, y and L carry flavor indices to allow for different mixing between N and each of the three Standard Model neutrino flavors. Equation (2) implies that the leptonic mixing matrix U is extended to a unitary 4×4 matrix, while apart from this modification the standard expression for the neutrino oscillation probabilities remains unchanged

$$P_{\alpha\beta} = \sum_{j,k} U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{-i\Delta m_{jk}^2 L/(2E)}. \quad (3)$$

Here, as usual $\Delta m_{jk}^2 \equiv m_j^2 - m_k^2$ is the difference between the squared masses of neutrino mass eigenstates j and k .

The $3 + 1$ scenario has been extensively studied in the context of the MiniBooNE anomaly (and the other short-baseline anomalies), see for instance refs. [37, 60–71]. These fits have revealed significant tension in the global data, caused mainly by the fact that MiniBooNE (and LSND) suggest relatively large mixing between the sterile neutrino ν_s and both ν_e and ν_μ . It is in particular the ν_s – ν_μ mixing that is strongly constrained by ν_μ disappearance searches such as the ones in MINOS/MINOS+ [72], IceCube [73–75], and MiniBooNE itself [76, 77]. The scenario is also constrained by cosmology, in particular by Big Bang Nucleosynthesis [78] and by CMB+structure formation data [79], though some or all of these constraints can be avoided in extended cosmological scenarios [7, 80–88]. In fact adding sterile neutrinos may even help alleviate the tension between local and cosmological determinations of the Hubble constant [89].

In the following, we address two important features of fits to the MiniBooNE data: the differences between a full four-flavor fit compared to the two-flavor fits presented for instance in refs. [2, 29], and the dependence of the fit results on the choice of event generator for the background predictions.

5.1 2-Flavor vs. 4-Flavor Fits to MiniBooNE Data

To fit the $3 + 1$ model to MiniBooNE data, we use an adapted version of the fitting code developed in refs. [7, 60, 63, 71]. It is based on the recommendations given by the MiniBooNE collaboration in the

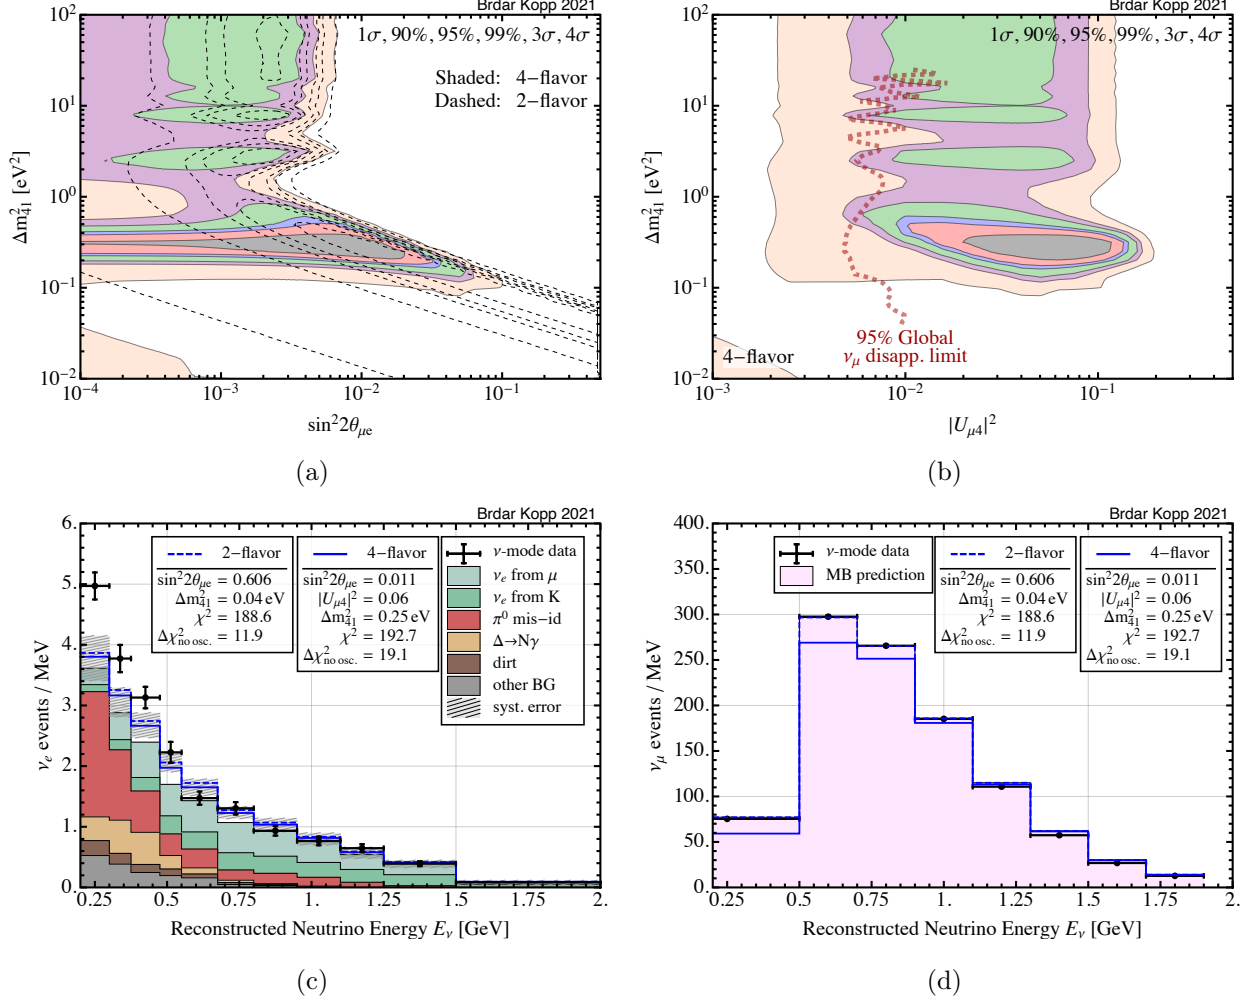


FIG. 6. Comparison between two-flavor and four-flavor fits to MiniBooNE data. In the two-flavor case, oscillations in the ν_e background and in the ν_μ control sample are negligible, while in the four-flavor case, $\nu_e \rightarrow \nu_s$ and $\nu_\mu \rightarrow \nu_s$ disappearance can occur. Because of the strong correlation between the ν_μ control sample and the ν_e signal sample, this leads to a shift in the best-fit region, with the four-flavor framework allowing for a good fit even in the absence of $\nu_\mu \rightarrow \nu_e$ oscillations, but with sizeable $|U_{\mu 4}|^2$, see shaded exclusion contours in panels (a) ($\sin^2 2\theta_{\mu e}$ vs. Δm_{41}^2) and (b) ($|U_{\mu 4}|^2$ vs. Δm_{41}^2). The ν_e spectrum at the four-flavor best fit point (solid blue histogram in panel (c)) is very close to the background prediction, while the ν_μ spectrum in panel (d) is suppressed by $\nu_\mu \rightarrow \nu_s$ disappearance. In the two-flavor case (dashed histograms), on the other hand, the fit is driven by $\nu_\mu \rightarrow \nu_e$ appearance. Note that the background prediction shown in panel (c) is from the four-flavor fit, but the one for the two-flavor scenario is practically identical. Panel (b) does not contain contours for the two-flavor case because $|U_{\mu 4}|^2$ is not defined in that scenario. We show, however, the global exclusion limit on ν_μ disappearance from ref. [71].

supplemental material of ref. [90] and uses the data released with ref. [34]. However, in contrast to the fits carried out in MiniBooNE's publications, we include the full impact of four-flavor oscillations on the signal and background prediction, as discussed in appendix A of ref. [7]. In particular, in a 3 + 1 model, explaining MiniBooNE's $\nu_\mu \rightarrow \nu_e$ oscillation signal requires mixing between the sterile state and both ν_e and ν_μ . $\nu_\mu \rightarrow \nu_e$ oscillations are thus necessarily accompanied by $\nu_e \rightarrow \nu_s$ and $\nu_\mu \rightarrow \nu_s$ disappearance. And while the probability for the appearance signal is proportional to $|U_{e4}|^2 |U_{\mu 4}|^2$, the disappearance probabilities are only suppressed by $|U_{e4}|^2$ and $|U_{\mu 4}|^2$. ν_e and ν_μ disappearance is thus a non-negligible effect which is not captured by the two-flavor fits employed

in the official MiniBooNE analyses. It has the following consequences:

1. **Oscillations in the ν_μ control sample.** MiniBooNE’s fit includes CC ν_μ events as a control sample to fix the unoscillated neutrino flux and spectrum in a data-driven way. An oscillation-induced ν_μ deficit, if not accounted for, will thus lead to too low a prediction for the intrinsic CC ν_e background and for the $\nu_\mu \rightarrow \nu_e$ signal. In MiniBooNE’s two-flavor fit, where a ν_μ deficit is only due to $\nu_\mu \rightarrow \nu_e$ oscillation, this effect can be neglected. But in a realistic four-flavor model, a much larger ν_μ deficit arises from $\nu_\mu \rightarrow \nu_s$ disappearance, leading to important corrections. We account for these corrections by scaling both the predicted oscillation signal and the CC ν_e background with the inverse of the ν_μ disappearance probability.⁶ That way, we compensate the bias that is introduced when MiniBooNE calibrate these backgrounds to the observed ν_μ rate. We do not rescale the other backgrounds (mostly π^0 and $\Delta \rightarrow \gamma$) because they are not normalized to the ν_μ control sample, but to single-pion control samples.
2. **ν_e disappearance.** By the same reasoning as for ν_μ , also the oscillation-induced deficit of ν_e is much larger in a 3 + 1 model than in MiniBooNE’s two-flavor scenario. This affects in particular the intrinsic CC ν_e background; we take this effect into account by rescaling said background with the ν_e disappearance probability.

An interesting outcome of the four-flavor fit – and in particular of the inclusion of ν_μ disappearance – is a significant distortion in MiniBooNE’s best fit regions, as illustrated in fig. 6 (a). In particular, because of the strong correlation between the ν_e signal sample and the ν_μ control sample that is used to normalize the signal, a good fit can be achieved not only by enhancing the ν_e flux, but also by suppressing the ν_μ flux. Therefore the four-flavor best fit regions (shaded contours in fig. 6) (a) extend to much smaller $\sin^2 2\theta_{\mu e} = 4 |U_{e4}|^2 |U_{\mu 4}|^2$ than in the two-flavor scenario (unshaded dashed contours), at the expense of relatively large $|U_{\mu 4}|^2$, see panel (b). Correspondingly, the ν_e spectrum at the four-flavor best-fit point (solid blue histogram in fig. 6 (c)) is closer to the background prediction, while the ν_μ spectrum in panel (d) is suppressed. In the two-flavor scenario, on the other hand, the ν_e flux needs to be enhanced as the ν_μ flux remains unsuppressed.

Note that the four-flavor treatment *increases* the significance of the anomaly: our two-flavor fit disfavors the no-oscillation hypothesis at 3.0σ , while the four-flavor fit disfavors it at 4.0σ . The lower significance of our own fit compared to MiniBooNE’s – even when following MiniBooNE’s recommended approach (including the assumption of two-flavor oscillations) as given in the supplemental material to refs. [34, 90] – has been noted before [60, 63, 71]. It implies that our results will be erring slightly on the side of being conservative.

It should be kept in mind that ν_μ disappearance is strongly constrained by dedicated measurements, including measurements by MiniBooNE itself [76, 77, 91], as well as IceCube [73–75], DeepCore [92, 93], CDHS [94], SuperKamiokande [95, 96], NO ν A, [97], and MINOS/MINOS+ [72]. In fig. 6 (b), we show as a dotted red line the combined ν_μ disappearance limit from ref. [71].

In the following, we will always use the full four-flavor framework when fitting MiniBooNE data.

5.2 Dependence on Background Predictions

We now investigate how the interpretation of MiniBooNE’s results in the context of the 3 + 1 scenario depends on the event generator used for predicting the backgrounds. In fig. 7, we show the resulting best-fit spectra in the ν_e channel (left) and the ν_μ channel (right) for different generators.

⁶ In fact, following MiniBooNE’s normalization strategy, the appropriate rescaling factor for ν_e from muon decay is the ν_μ disappearance probability at MiniBooNE, while for ν_e from kaon decay, it is the corresponding probability at the SciBooNE baseline of 100 m.

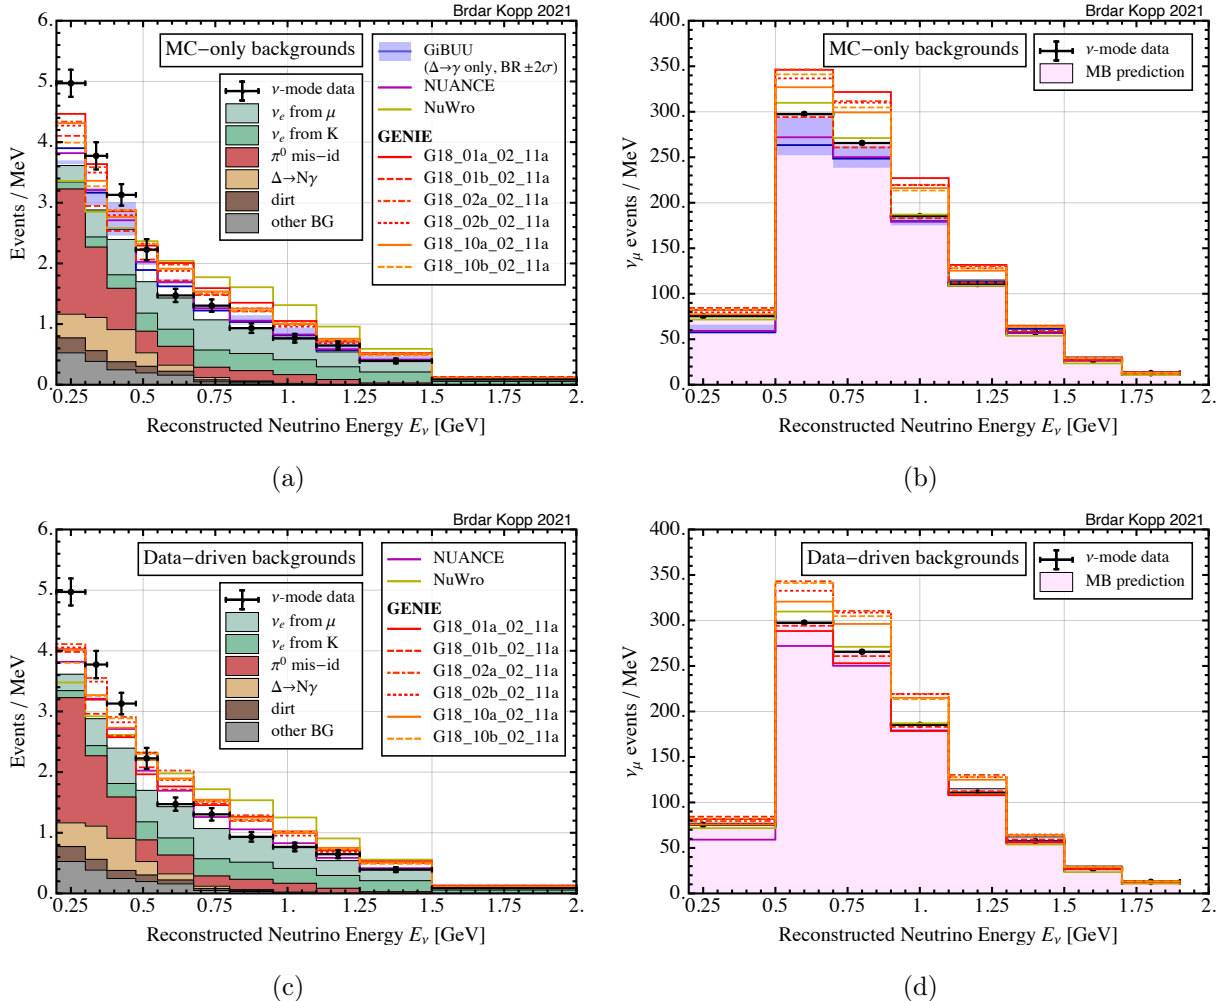


FIG. 7. Predicted MiniBooNE event spectra at the best-fit point of a $3 + 1$ model, using different event generators and tunes to predict the ν_e , π^0 and $\Delta \rightarrow N\gamma$ backgrounds. The panels on the left are for ν_e -like events, while the panels on the right show the ν_μ spectra. In the top panels (a) and (b), background predictions are based on the Monte Carlo-only simulations described in section 2, while panels (c) and (d) at the bottom are based on the data-driven predictions from section 3.

In the fits, the ν_e , π^0 , and single-photon (“ $\Delta \rightarrow N\gamma$ ”) backgrounds in the ν_e channel, as well as the prediction for the ν_μ channel, are based on our own prediction discussed in the previous sections.⁷ For the other backgrounds we use MiniBooNE’s official predictions, see ref. [34]. An exception is GiBUU: as discussed in section 2, GiBUU underpredicts the rate of pion production in MiniBooNE (but, curiously, not in other experiments), which makes its predictions for the CCQE channels, for the NC π^0 backgrounds, and for all data-driven predictions anchored to the NC π^0 rate, problematic. In view of this, we do not show fits with data-driven backgrounds from GiBUU at all, and for fits with Monte Carlo-only backgrounds, we use GiBUU only for the $\Delta \rightarrow N\gamma$ channel, and MiniBooNE’s own predictions in all other channels. On the other hand, we indicate as a blue band how GiBUU’s best fit spectrum varies if the radiative branching ratios of the $\Delta(1232)$ and $N(1440)$ resonances are varied within their 2σ limits, as discussed in section 4. Note that the backgrounds shown in fig. 7 as colored histograms are MiniBooNE’s; see figs. 1 to 3 for comparisons

⁷ Note that we do not take into account the possible impact of misreconstruction of CCQE events on the signal events. It has been shown in ref. [33] that this effect is small.

Monte Carlo-only background predictions

Generator	Tune	Δm_{41}^2 [eV ²]	$\sin^2 2\theta_{\mu e}$	$ U_{\mu 4} ^2$	χ^2/dof	$\Delta\chi_{\text{no osc.}}^2$	Significance
MB official		0.25	0.01	0.062	12.0	19.1	4.0 σ
GiBUU	default	0.25	0.01	0.076	12.0	24.6	4.6 σ
	BR($\Delta \rightarrow \gamma$) - 2 σ	0.32	0.0063	0.090	12.4	32.3	5.3 σ
	BR($\Delta \rightarrow \gamma$) + 2 σ	2.5	0.0020	0.0056	12.1	8.11	2.4 σ
NUANCE	–	0.32	0.0079	0.051	12.3	19.3	4.0 σ
NuWro	–	3.2	0.0020	0.040	13.7	15.6	3.5 σ
GENIE	G18_01a_02_11a	0.13	0.079	0.16	12.2	21.6	4.3 σ
	G18_01b_02_11a	0.79	0.0001	0.12	12.2	16.1	3.6 σ
	G18_02a_02_11a	0.13	0.050	0.16	12.0	15.1	3.5 σ
	G18_02b_02_11a	0.13	0.050	0.18	12.1	15.0	3.5 σ
	G18_10a_02_11a	0.25	0.016	0.051	12.1	11.2	2.9 σ
	G18_10b_02_11a	0.40	0.013	0.016	12.1	17.9	3.8 σ

data-driven backgrounds

Generator	Tune	Δm_{41}^2	$\sin^2 2\theta_{\mu e}$	$ U_{\mu 4} ^2$	χ^2/dof	$\Delta\chi_{\text{no osc.}}^2$	Significance
MB official		0.25	0.01	0.062	12.0	19.1	4.0 σ
NUANCE	–	0.32	0.0079	0.051	12.3	19.3	4.0 σ
NuWro	–	3.2	0.0016	0.040	13.3	12.7	3.1 σ
GENIE	G18_01a_02_11a	0.79	0.00020	0.14	12.2	23.3	4.4 σ
	G18_01b_02_11a	0.79	0.0001	0.12	12.2	15.5	3.5 σ
	G18_02a_02_11a	0.13	0.063	0.18	12.2	19.2	4.0 σ
	G18_02b_02_11a	0.13	0.050	0.20	12.3	16.9	3.7 σ
	G18_10a_02_11a	0.25	0.016	0.062	12.3	15.1	3.5 σ
	G18_10b_02_11a	0.40	0.013	0.016	12.1	19.5	4.0 σ

TABLE II. Results of fitting a 3 + 1 sterile neutrino model to MiniBooNE data, using different event generators and tunes to predict the ν_e , π^0 , and single-photon backgrounds. Besides the parameter values at the respective best-fit points, we also list the $\Delta\chi^2$ at which the no-oscillation hypothesis is excluded, and we convert this number into a statistical significance for an anomaly, assuming a χ^2 distribution with one degree of freedom. Note that we do not use GiBUU’s Monte Carlo-only predictions for CCQE-like events and NC π^0 events, which appear problematic because GiBUU predicts fewer pions than observed. For the same reason, we also do not show result for GiBUU with data-driven backgrounds at all. Here, GiBUU’s pion deficit would affect the π^0 control sample.

between MiniBooNE’s background predictions and ours for individual generators. Finally, as we do not have a data-driven prediction for the intrinsic ν_e backgrounds (green shaded histograms in fig. 7), we use the Monte Carlo-only one even for the fits labeled “data-driven”.

We observe significant spread between the best-fit spectra from different generators. Importantly, this is the case not only for background predictions based on Monte Carlo simulations alone (top panels), but also for our data-driven predictions (bottom panels). For the latter, it is driven by the residual uncertainty in the single-photon channel, see section 3.2. Some generators, in particular some GENIE tunes, are able to accommodate fairly large event rates in the low-energy bins, consistent with the observed excess. However, they tend to also overpredict the rate at higher energies a bit, suggesting that the goodness of fit will not be too different compared to MiniBooNE’s official fit.

Table II reveals that differences between event generators translate into some differences in the best-fit points and in the significance of the anomaly. However, the latter remains around 3 – 4 σ , similar to the significance we obtain when using MiniBooNE’s background predictions. We find the

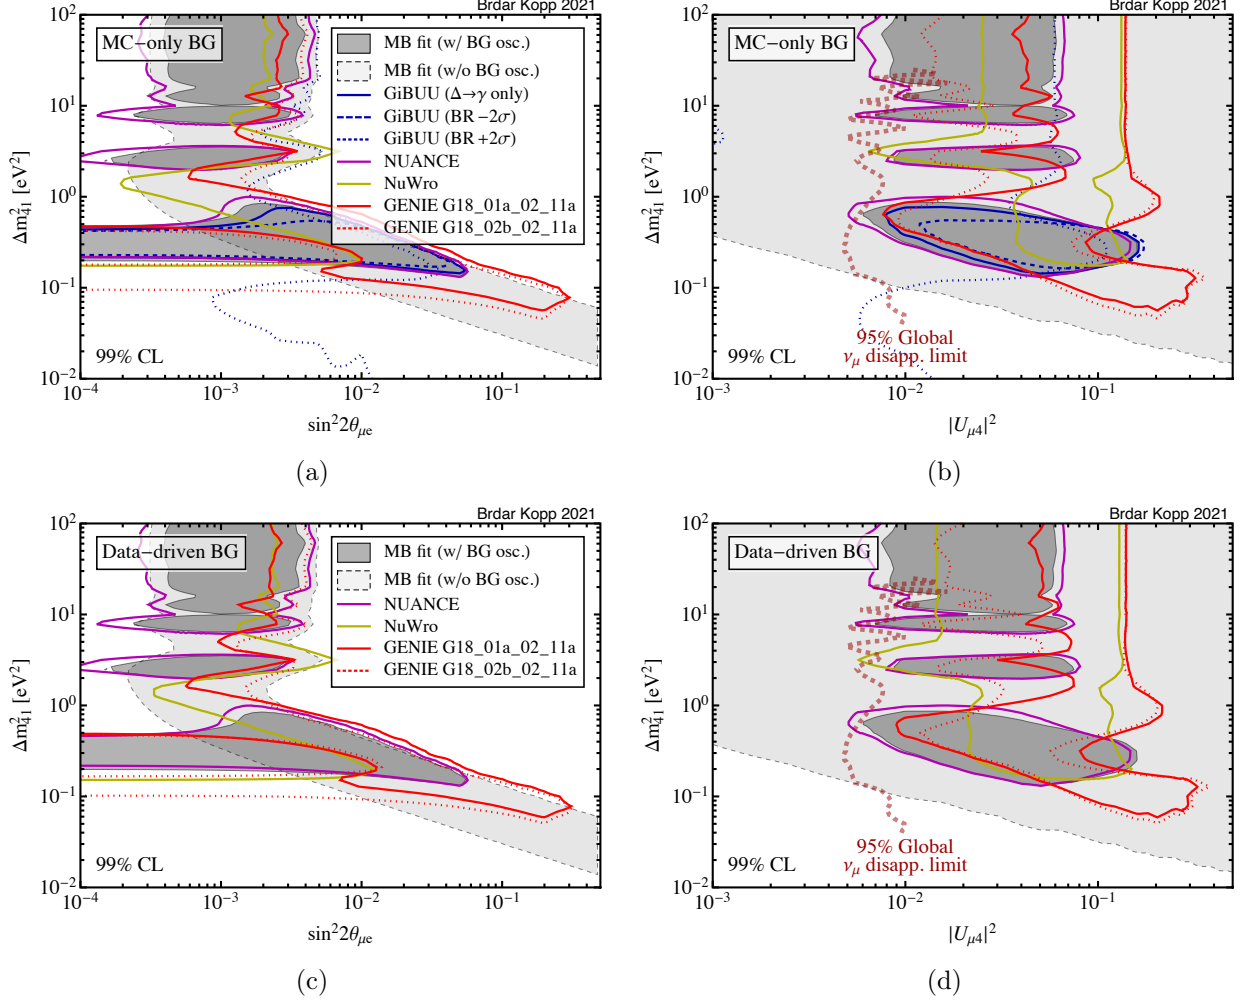


FIG. 8. MiniBooNE 99% exclusion contours for a 3+1 sterile neutrino model using different event generators and tunes for the prediction of the ν_e , π^0 and $\Delta \rightarrow N\gamma$ backgrounds. (For clarity, we show results for only two GENIE tunes here; exclusion contours for the others can be found in appendix A.) The panels on the left show the $\sin^2 2\theta_{\mu e} - \Delta m_{41}^2$ plane (ν_e appearance); the ones on the right show the $|U_{\mu 4}|^2 - \Delta m_{41}^2$ plane (ν_μ disappearance). In the top panels (a) and (b), background predictions are based on the Monte Carlo-only calculations described in section 2, while panels (c) and (d) at the bottom are based on the data-driven predictions from section 3.

lowest significance (2.4σ) when using GiBUU to predict the $\Delta \rightarrow N\gamma$ background and increasing the $\Delta(1232)$ and $N(1440)$ radiative decay widths to the upper end of their 2σ -allowed intervals from ref. [45].

This is also illustrated in fig. 8, which compares the parameter space exclusion regions in the 3 + 1 scenario between analyses based on different generators. This is once again done both for background predictions based on Monte Carlo simulations alone (top panels) and for data-driven background predictions in the π^0 and single-photon channels (bottom panels). We show projections onto the $\sin^2 2\theta_{\mu e} - \Delta m_{41}^2$ plane (panels (a) and (c)) and onto the $|U_{\mu 4}|^2 - \Delta m_{41}^2$ plane (panels (b) and (d)).

We observe that the contours based on the NUANCE generator (purple) are, as expected, in excellent agreement with those based on MiniBooNE’s official background predictions (which also rely on NUANCE). Significant deviations are seen for fits using NuWro, GENIE, and GiBUU predic-

tions, though. (As before, we include GiBUU only in the comparison Monte Carlo-only predictions, and even there we use its predictions only for the single-photon channel.) Notably, NuWro and GENIE allow $\sin^2 2\theta_{\mu e} = 0$ at the 99% CL over a wide range of Δm_{41}^2 values, while GiBUU, as well as the fit using MiniBooNE’s official background predictions, do so only in a narrow window around $\Delta m_{41}^2 \sim 0.3 \text{ eV}^2$. Remember that allowing $\sin^2 2\theta_{\mu e} = 0$ does not mean that the anomaly is resolved – the fits then require non-zero $|U_{\mu 4}|^2$, see the right-hand panels of fig. 8. As explained in section 5.1, this can be understood from the strong correlations between the ν_e and ν_μ data. An excess of ν_μ events compared to the theory prediction may thereby be sufficient to explain an excess also in the ν_e channel even without explicit ν_e appearance. Note, however, that the values of $|U_{\mu 4}|^2$ required to accommodate the MiniBooNE anomaly are still in tension with the global exclusion limit on ν_μ disappearance (dotted red line in the right-hand panels of fig. 8). The one exception is the GiBUU fit with increased radiative branching ratios (blue dotted lines). It is consistent with the no-oscillation hypothesis in both the ν_e and ν_μ channel at the 99% CL, in agreement with the low significance of the anomaly we found in this scenario in table II.

6 SUMMARY AND CONCLUSIONS

The MiniBooNE anomaly is still one of the biggest mysteries in neutrino physics. In this paper we have revisited the most relevant backgrounds for the MiniBooNE ν_e appearance analysis in which the anomalous event excess is observed. We have in particular studied CC interactions of beam ν_e , NC π^0 production, and single-photon production. We have predicted the event rates in these channels using different event generators – namely NUANCE, GiBUU, GENIE, and NuWro – and have compared the results to estimate the theoretical uncertainties associated with our predictions. For the π^0 background, we have found that generators agree at the 10% level (with the exception of GiBUU, which underpredicts the π^0 production rate by almost a factor of two). For CC ν_e and single-photon events, discrepancies are somewhat larger, with predictions differing by $\mathcal{O}(30\%)$.

The situation improves only slightly when we attempt to predict the π^0 and single- γ backgrounds in a more data-driven way by normalizing them to MiniBooNE’s measured π^0 production rate. (The deficit of π^0 -induced events in GiBUU is, however, almost entirely removed that way.)

In addition, we have pointed out that the radiative branching ratios of heavy hadronic resonances, most importantly the $\Delta(1232)$, come with large errors of order tens of per cent that affect in particular the single-photon background.

In the final part of the paper, we have discussed fits to the MiniBooNE data in the context of sterile neutrino models. We have highlighted the important differences between fits in a two-flavor framework, which are often shown in the literature, and a more careful fit that takes into account full four-flavor oscillations. In the four-flavor case, oscillations of the ν_e background and in the ν_μ control sample play a crucial role. Most notably, the anomaly could be entirely explained by ν_μ disappearance alone, thanks to the strong correlation between the ν_μ and ν_e samples in the fit. The tension with the non-observation of ν_μ disappearance in other experiments would still persist, though.

We have then studied how the choice of event generator affects the fit in a $3 + 1$ sterile neutrino model, and in particular the significance of the anomaly. We have found that all generators roughly agree on the significance of the anomaly, but that allowing for some freedom ($\pm 2\sigma$) in the branching ratios of the (never directly observed) radiative decays of heavy hadronic resonances like $\Delta(1232)$ and $N(1440)$ can bring the anomaly down from $\sim 4\sigma$ to less than 3σ , without completely resolving it though.

We conclude that theoretical uncertainties in MiniBooNE’s background predictions certainly deserve further study. However, it seems that with our current understanding of neutrino interaction

physics – as implemented in state-of-the-art event generators, the anomaly is robust. Not even an “Altarelli cocktail” of several small deviations in different channels that add up to a potentially much bigger overall discrepancy seems to be able to fully explain the event excess.

ACKNOWLEDGMENTS

We would like to thank David Caratelli, Dave Casper, Teppei Katori, Bill Louis, Ulrich Mosel, and Jan Sobczyk for very useful discussions, as well as Pedro Machado for collaboration in the early stages of this work. Fermilab is operated by Fermi Research Alliance, LLC under contract No. DE-AC02-07CH11359 with the United States Department of Energy. JK has been partly funded by the German Research Foundation (DFG) in the framework of the PRISMA+ Cluster of Excellence and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 637506, “ ν Directions”).

A PARAMETER SCANS IN THE 3 + 1 STERILE NEUTRINO MODEL

In this appendix, we supplement the discussion in section 5.2 by providing parameter space exclusion plots for all Monte Carlo generators and tunes studied in this paper in the 3 + 1 sterile neutrino scenario. In particular, we show in fig. 9 our fit results using Monte Carlo-only background predictions, that is predictions which are not tuned to MiniBooNE’s own π^0 production data. In fig. 10 we do the same using data-driven prediction backgrounds for the π^0 and single-photon channels, which have been tuned to π^0 data. For each background model, we show the projections of the 3 + 1 model’s parameter space onto the $\sin^2 2\theta_{\mu e} - \Delta m_{41}^2$ plane and onto the $|U_{\mu 4}|^2 - \Delta m_{41}^2$ plane. (The SM oscillation parameters are irrelevant here because MiniBooNE’s baseline is too short for SM oscillations to develop, and the mixing matrix element $U_{\tau 4}$ is irrelevant due to the absence of τ neutrinos.)

-
- [1] **BooNE Collaboration**, E. Church *et al.*, *A proposal for an experiment to measure $\nu_\mu \rightarrow \nu_e$ oscillations and ν_μ disappearance at the Fermilab Booster: BooNE*, .
 - [2] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *Updated MiniBooNE Neutrino Oscillation Results with Increased Data and New Background Studies*, [2006.16883](#).
 - [3] J. A. Formaggio and G. P. Zeller, *From eV to EeV: Neutrino Cross Sections Across Energy Scales*, *Rev. Mod. Phys.* **84** (2012) 1307–1341, [[1305.7513](#)].
 - [4] O. Fischer, A. Hernández-Cabezudo, and T. Schwetz, *Explaining the MiniBooNE excess by a decaying sterile neutrino with mass in the 250 MeV range*, [1909.09561](#).
 - [5] S. N. Gninenko, *The MiniBooNE anomaly and heavy neutrino decay*, *Phys. Rev. Lett.* **103** (2009) 241802, [[0902.3802](#)].
 - [6] E. Bertuzzo, S. Jana, P. A. N. Machado, and R. Zukanovich Funchal, *Dark Neutrino Portal to Explain MiniBooNE excess*, *Phys. Rev. Lett.* **121** (2018), no. 24 241801, [[1807.09877](#)].
 - [7] M. Dentler, I. Esteban, J. Kopp, and P. Machado, *Decaying Sterile Neutrinos and the Short Baseline Oscillation Anomalies*, [1911.01427](#).
 - [8] P. Ballett, S. Pascoli, and M. Ross-Lonergan, *$U(1)$ ’ mediated decays of heavy sterile neutrinos in MiniBooNE*, *Phys. Rev.* **D99** (2019) 071701, [[1808.02915](#)].
 - [9] A. de Gouvêa, O. L. G. Peres, S. Prakash, and G. V. Stenico, *On The Decaying-Sterile Neutrino Solution to the Electron (Anti)Neutrino Appearance Anomalies*, [1911.01447](#).
 - [10] W. Abdallah, R. Gandhi, and S. Roy, *Understanding the MiniBooNE and the muon $g - 2$ anomalies with a light Z' and a second Higgs doublet*, [2006.01948](#).

Monte Carlo Backgrounds

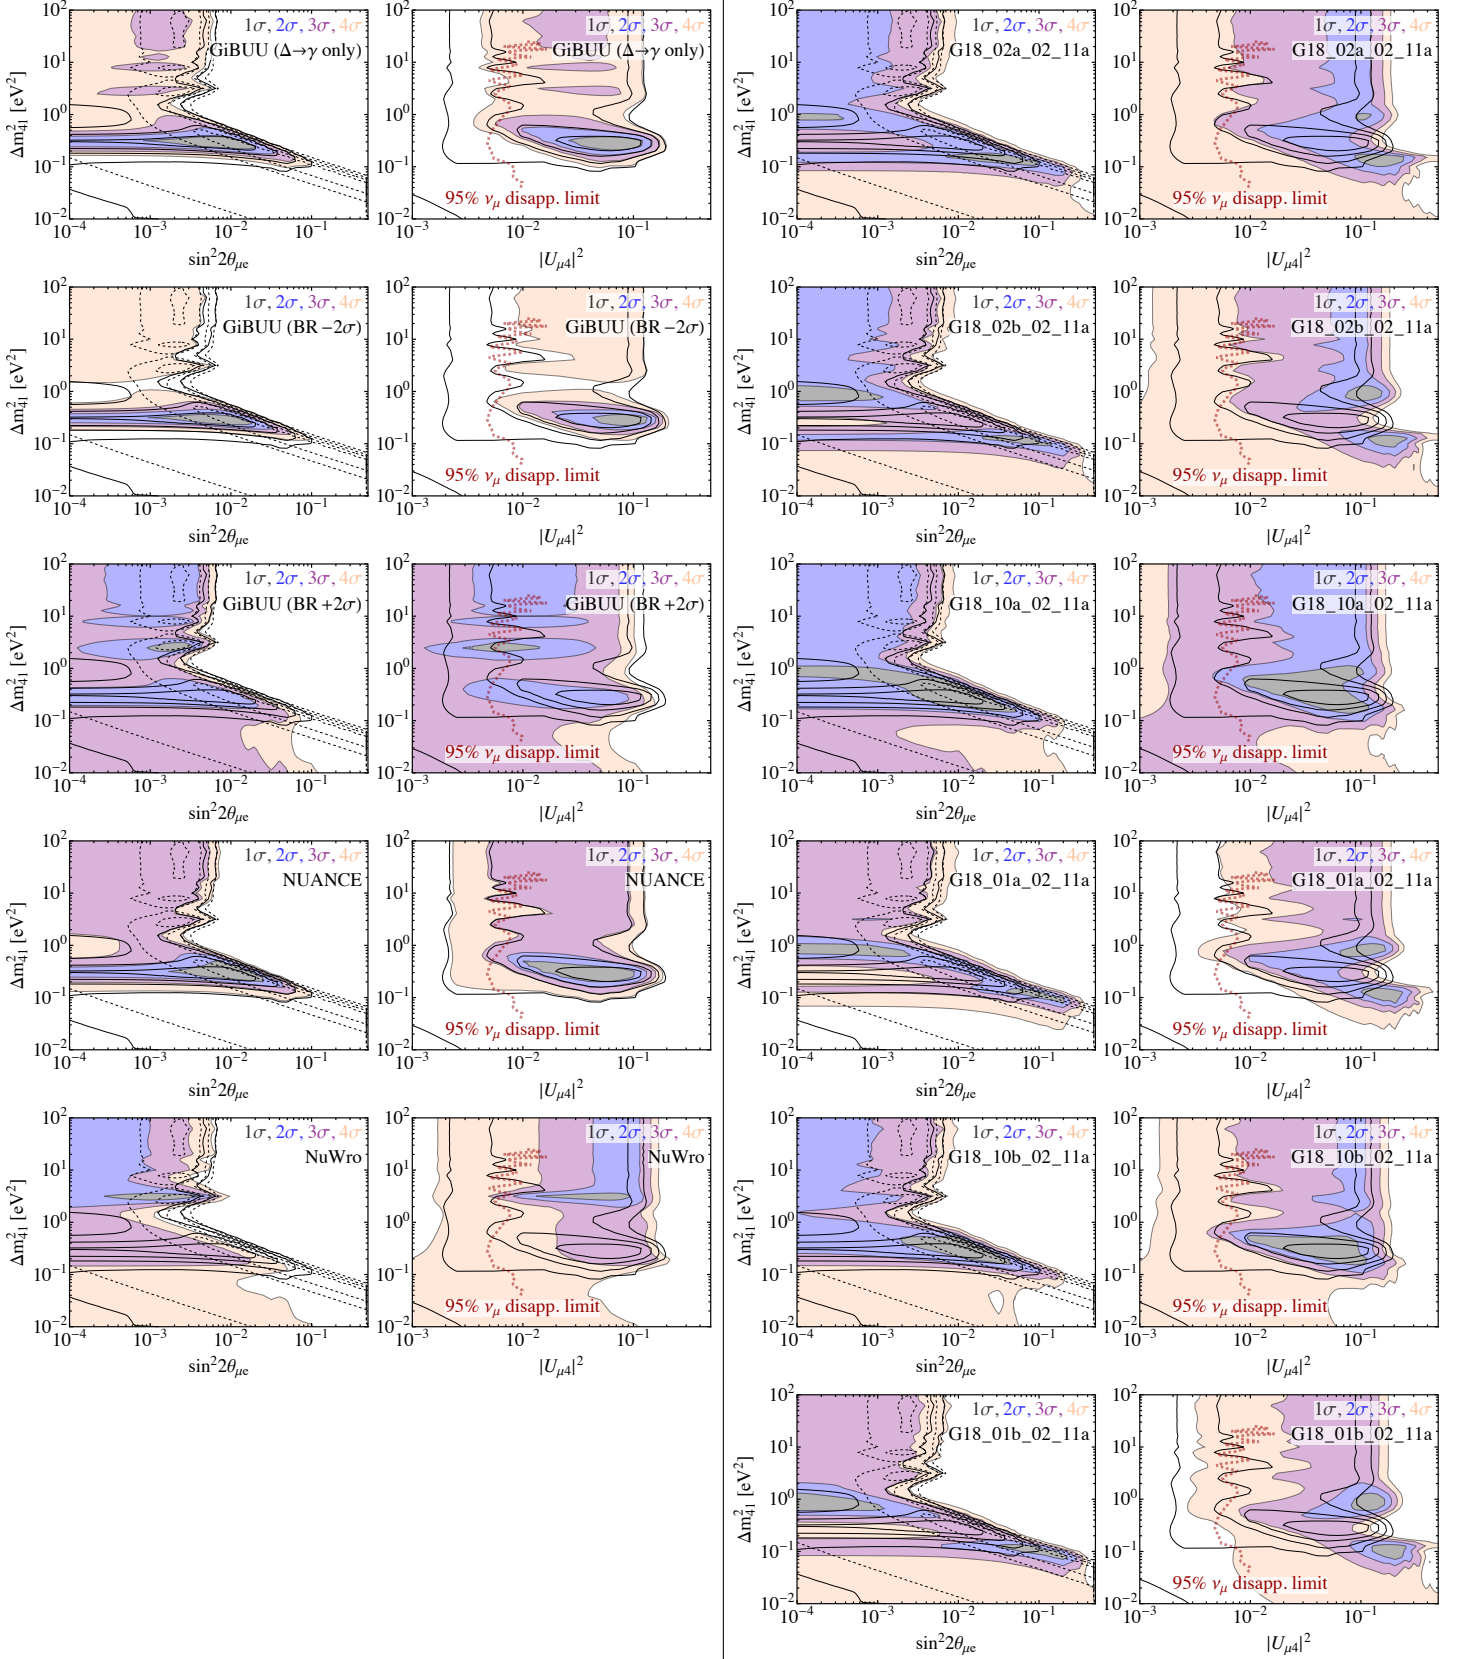


FIG. 9. Shaded contours: MiniBooNE exclusion contours in the $\sin^2 2\theta_{\mu e} - \Delta m_{41}^2$ plane and in the $|U_{\mu 4}|^2 - \Delta m_{41}^2$ plane for different event generators using Monte Carlo predictions *without* tuning to MiniBooNE π^0 data. Unfilled black contours: fit results using MiniBooNE's official background predictions. Red dotted line in the $|U_{\mu 4}|^2$ -vs.- Δm_{41}^2 panels: global 95% ν_{μ} disappearance limit from ref. [71].

Data-Driven Backgrounds

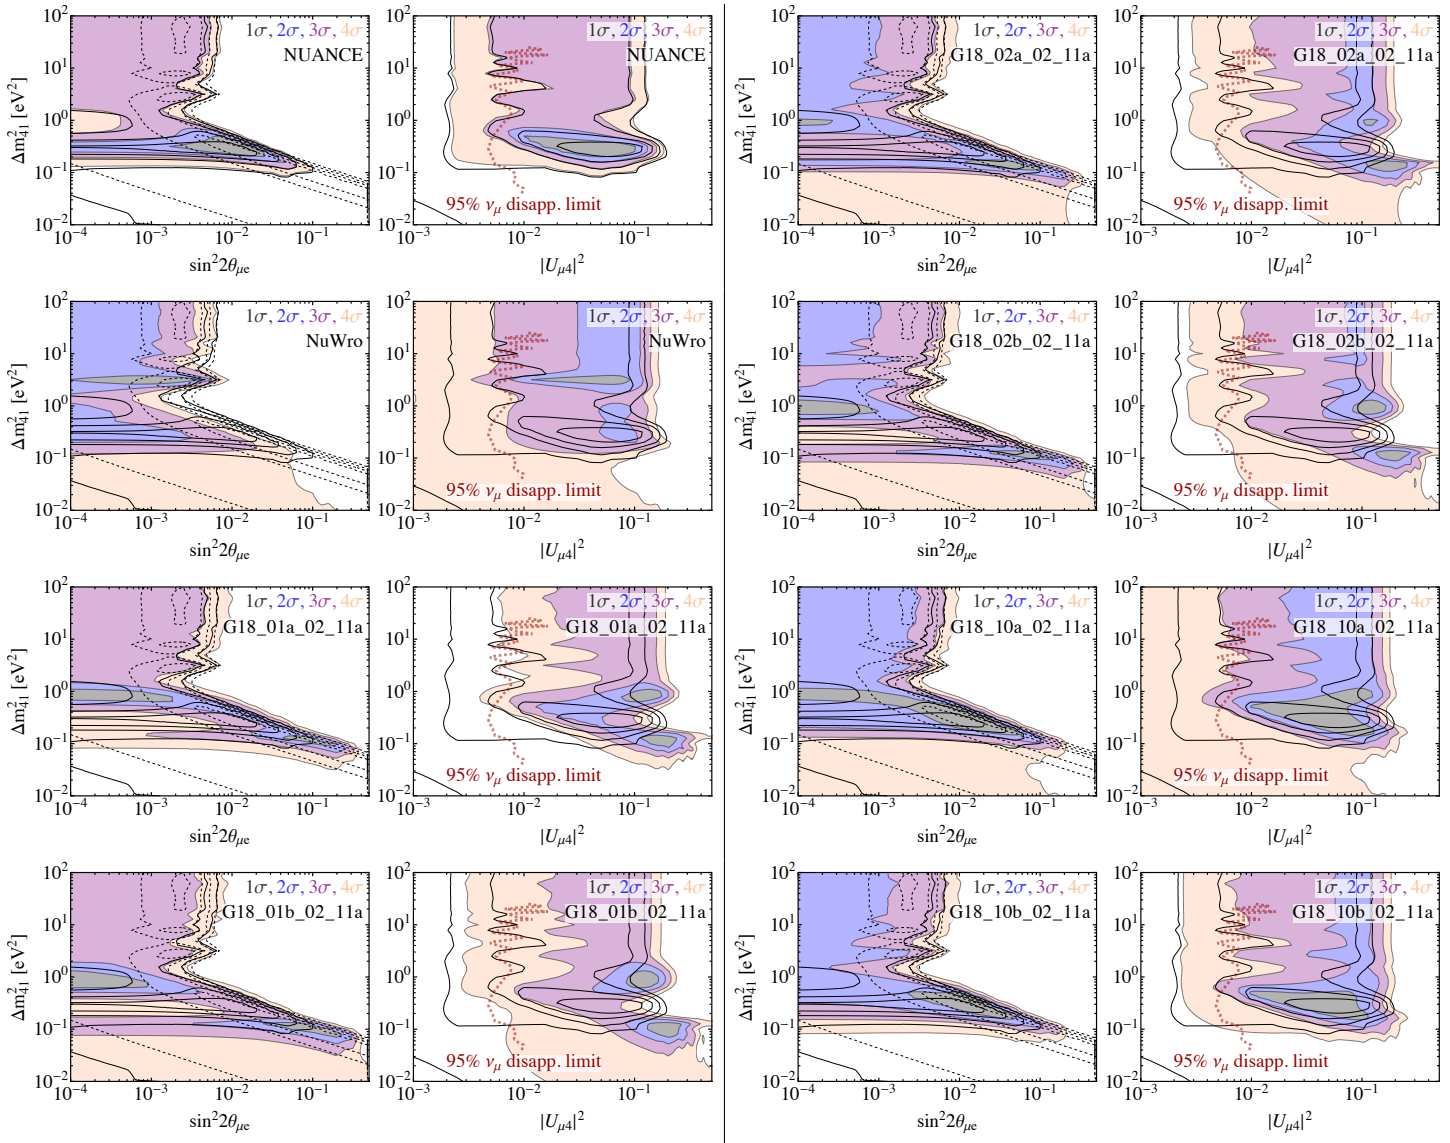


FIG. 10. Same as fig. 9, but for data-driven background predictions in the π^0 and single-photon channels.

- [11] B. Dutta, S. Ghosh, and T. Li, *Explaining $(g - 2)_{\mu,e}$, KOTO anomaly and MiniBooNE excess in an extended Higgs model with sterile neutrinos*, [2006.01319](#).
- [12] A. Datta, S. Kamali, and D. Marfatia, *Dark sector origin of the KOTO and MiniBooNE anomalies*, [2005.08920](#).
- [13] A. Abdullahi, M. Hostert, and S. Pascoli, *A Dark Seesaw Solution to Low Energy Anomalies: MiniBooNE, the muon $(g - 2)$, and BaBar*, [2007.11813](#).
- [14] V. Brdar, O. Fischer, and A. Y. Smirnov, *Model Independent Bounds on the Non-Oscillatory Explanations of the MiniBooNE Excess*, [2007.14411](#).
- [15] **MicroBooNE Collaboration**, R. Acciarri *et al.*, *Design and Construction of the MicroBooNE Detector*, *JINST* **12** (2017), no. 02 P02017, [[1612.05824](#)].
- [16] **SBND Collaboration**, N. McConkey, *SBND: Status of the Fermilab Short-Baseline Near Detector*, *J. Phys. Conf. Ser.* **888** (2017), no. 1 012148.
- [17] C. Rubbia *et al.*, *Underground operation of the ICARUS T600 LAr-TPC: first results*, *JINST* **6** (2011) P07011, [[1106.0975](#)].

- [18] P. A. Machado, O. Palamara, and D. W. Schmitz, *The Short-Baseline Neutrino Program at Fermilab*, *Ann. Rev. Nucl. Part. Sci.* **69** (2019) 363–387, [[1903.04608](#)].
- [19] R. K. Ellis, *Guido Altarelli and the evolution of QCD*, *Nuovo Cim. C* **39** (2017), no. 4 355, [[1608.05574](#)].
- [20] M. Martini, M. Ericson, G. Chanfray, and J. Marteau, *A Unified approach for nucleon knock-out, coherent and incoherent pion production in neutrino interactions with nuclei*, *Phys. Rev.* **C80** (2009) 065501, [[0910.2622](#)].
- [21] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas, *Inclusive Charged-Current Neutrino-Nucleus Reactions*, *Phys. Rev. C* **83** (2011) 045501, [[1102.2777](#)].
- [22] J. T. Sobczyk, *Multinucleon ejection model for Meson Exchange Current neutrino interactions*, *Phys. Rev.* **C86** (2012) 015504, [[1201.3673](#)].
- [23] A. Meucci and C. Giusti, *Relativistic descriptions of final-state interactions in charged-current quasielastic antineutrino-nucleus scattering at MiniBooNE kinematics*, *Phys. Rev.* **D85** (2012) 093002, [[1202.4312](#)].
- [24] O. Lalakulich, K. Gallmeister, and U. Mosel, *Many-Body Interactions of Neutrinos with Nuclei - Observables*, *Phys. Rev.* **C86** (2012), no. 1 014614, [[1203.2935](#)]. [Erratum: *Phys. Rev. C* **90**, no. 2, 029902 (2014)].
- [25] D. Meloni and M. Martini, *Revisiting the T2K data using different models for the neutrino-nucleus cross sections*, *Phys. Lett.* **B716** (2012) 186–192, [[1203.3335](#)].
- [26] J. Nieves, F. Sanchez, I. Ruiz Simo, and M. J. Vicente Vacas, *Neutrino Energy Reconstruction and the Shape of the CCQE-like Total Cross Section*, *Phys. Rev.* **D85** (2012) 113008, [[1204.5404](#)].
- [27] O. Lalakulich, U. Mosel, and K. Gallmeister, *Energy reconstruction in quasielastic scattering in the MiniBooNE and T2K experiments*, *Phys. Rev.* **C86** (2012) 054606, [[1208.3678](#)].
- [28] M. Martini, M. Ericson, and G. Chanfray, *Energy reconstruction effects in neutrino oscillation experiments and implications for the analysis*, *Phys. Rev. D* **87** (2013), no. 1 013009, [[1211.1523](#)].
- [29] **MiniBooNE Collaboration**, A. Aguilar-Arevalo *et al.*, *Improved Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment*, *Phys. Rev. Lett.* **110** (2013) 161801, [[1303.2588](#)].
- [30] P. Coloma and P. Huber, *Impact of nuclear effects on the extraction of neutrino oscillation parameters*, *Phys. Rev. Lett.* **111** (2013), no. 22 221802, [[1307.1243](#)].
- [31] U. Mosel, O. Lalakulich, and K. Gallmeister, *Energy reconstruction in the Long-Baseline Neutrino Experiment*, *Phys. Rev. Lett.* **112** (2014) 151802, [[1311.7288](#)].
- [32] G. D. Megias *et al.*, *Meson-exchange currents and quasielastic predictions for charged-current neutrino- ^{12}C scattering in the superscaling approach*, *Phys. Rev.* **D91** (2015), no. 7 073004, [[1412.1822](#)].
- [33] M. Ericson, M. V. Garzelli, C. Giusti, and M. Martini, *Assessing the role of nuclear effects in the interpretation of the MiniBooNE low-energy anomaly*, *Phys. Rev.* **D93** (2016), no. 7 073008, [[1602.01390](#)].
- [34] **MiniBooNE Collaboration**, A. Aguilar-Arevalo *et al.*, *Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment*, *Phys. Rev. Lett.* **121** (2018), no. 22 221801, [[1805.12028](#)].
- [35] P. Stowell *et al.*, *NUISANCE: a neutrino cross-section generator tuning and comparison framework*, *JINST* **12** (2017), no. 01 P01016, [[1612.07393](#)].
- [36] A. Ioannian, *A Standard Model explanation for the excess of electron-like events in MiniBooNE*, [1909.08571](#).
- [37] C. Giusti, A. Ioannian, and G. Ranucci, *A new analysis of the MiniBooNE low-energy excess*, [1912.01524](#).
- [38] W. C. Louis, 2019. private communication.
- [39] C. Andreopoulos, C. Barry, S. Dytman, H. Gallagher, T. Golan, R. Hatcher, G. Perdue, and J. Yarba, *The GENIE Neutrino Monte Carlo Generator: Physics and User Manual*, [1510.05494](#).
- [40] D. Casper, *The Nuance neutrino physics simulation, and the future*, *Nucl. Phys. B Proc. Suppl.* **112** (2002) 161–170, [[hep-ph/0208030](#)].
- [41] T. Golan, C. Juszczak, and J. T. Sobczyk, *Final State Interactions Effects in Neutrino-Nucleus Interactions*, *Phys. Rev. C* **86** (2012) 015505, [[1202.4197](#)].
- [42] T. Leitner, O. Buss, L. Alvarez-Ruso, and U. Mosel, *Electron- and neutrino-nucleus scattering from the quasielastic to the resonance region*, *Phys. Rev. C* **79** (2009) 034601, [[0812.0587](#)].

- [43] **MiniBooNE Collaboration**, A. Aguilar-Arevalo *et al.*, *First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section*, *Phys. Rev. D* **81** (2010) 092005, [[1002.2680](#)].
- [44] “Genie comprehensive model configurations and tunes.” <http://tunes.genie-mc.org/>.
- [45] **Particle Data Group Collaboration**, P. A. Zyla *et al.*, *Review of Particle Physics*, *PTEP* **2020** (2020), no. 8 083C01.
- [46] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *The Neutrino Flux prediction at MiniBooNE*, *Phys. Rev. D* **79** (2009) 072002, [[0806.1449](#)].
- [47] T. Katori, *A Measurement of the muon neutrino charged current quasielastic interaction and a test of Lorentz violation with the MiniBooNE experiment*. PhD thesis, Indiana U., 2008.
- [48] E. J. Moniz, I. Sick, R. R. Whitney, J. R. Ficenec, R. D. Kephart, and W. P. Trower, *Nuclear fermi momenta from quasielastic electron scattering*, *Phys. Rev. Lett.* **26** (Feb, 1971) 445–448.
- [49] M. Martini, M. Ericson, and G. Chanfray, *Neutrino energy reconstruction problems and neutrino oscillations*, *Phys. Rev. D* **85** (2012) 093012, [[1202.4745](#)].
- [50] T. Leitner, O. Buss, U. Mosel, and L. Alvarez-Ruso, *Neutrino induced pion production at MiniBooNE and K2K*, *Phys. Rev. C* **79** (2009) 038501, [[0812.1787](#)].
- [51] O. Lalakulich and U. Mosel, *Pion production in the MiniBooNE experiment*, *Phys. Rev. C* **87** (2013), no. 1 014602, [[1210.4717](#)].
- [52] U. Mosel, *Neutrino Interactions with Nucleons and Nuclei: Importance for Long-Baseline Experiments*, *Ann. Rev. Nucl. Part. Sci.* **66** (2016) 171–195, [[1602.00696](#)].
- [53] “Miniboone collaboration: talk on april 11, 2007, slide 24.” http://www-boone.fnal.gov/publicpages/First_Results.pdf.
- [54] <http://pdg.lbl.gov/2019/reviews/rpp2019-rev-passage-particles-matter.pdf>.
- [55] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *The MiniBooNE Detector*, *Nucl. Instrum. Meth.* **A599** (2009) 28–46, [[0806.4201](#)].
- [56] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *A Combined $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillation Analysis of the MiniBooNE Excesses*, 7, 2012. [1207.4809](#). accompanying data release at https://www-boone.fnal.gov/for_physicists/data_release/nue_nuebar_2012/.
- [57] G. Chanfray and M. Ericson, *Gamma production in neutrino interaction*, [2105.02505](#).
- [58] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *Measurement of ν_μ and $\bar{\nu}_\mu$ induced neutral current single π^0 production cross sections on mineral oil at $E_\nu \sim \mathcal{O}(1\text{GeV})$* , *Phys. Rev. D* **81** (2010) 013005, [[0911.2063](#)].
- [59] D. Rönchen, M. Döring, H. Habermann, J. Haidenbauer, U. G. Meißner, and K. Nakayama, *Eta photoproduction in a combined analysis of pion- and photon-induced reactions*, *Eur. Phys. J. A* **51** (2015), no. 6 70, [[1504.01643](#)].
- [60] J. Kopp, M. Maltoni, and T. Schwetz, *Are there sterile neutrinos at the eV scale?*, *Phys.Rev.Lett.* **107** (2011) 091801, [[1103.4570](#)].
- [61] J. Conrad, C. Ignarra, G. Karagiorgi, M. Shaevitz, and J. Spitz, *Sterile Neutrino Fits to Short Baseline Neutrino Oscillation Measurements*, *Adv.High Energy Phys.* **2013** (2013) 163897, [[1207.4765](#)].
- [62] M. Archidiacono, N. Fornengo, C. Giunti, S. Hannestad, and A. Melchiorri, *Sterile Neutrinos: Cosmology vs Short-Baseline Experiments*, [1302.6720](#).
- [63] J. Kopp, P. A. N. Machado, M. Maltoni, and T. Schwetz, *Sterile Neutrino Oscillations: The Global Picture*, *JHEP* **1305** (2013) 050, [[1303.3011](#)].
- [64] A. Mirizzi, G. Mangano, N. Saviano, E. Borriello, C. Giunti, *et al.*, *The strongest bounds on active-sterile neutrino mixing after Planck data*, [1303.5368](#).
- [65] C. Giunti, M. Laveder, Y. Li, and H. Long, *Pragmatic View of Short-Baseline Neutrino Oscillations*, *Phys.Rev.* **D88** (2013) 073008, [[1308.5288](#)].
- [66] S. Gariazzo, C. Giunti, and M. Laveder, *Light Sterile Neutrinos in Cosmology and Short-Baseline Oscillation Experiments*, *JHEP* **1311** (2013) 211, [[1309.3192](#)].
- [67] G. H. Collin, C. A. Argüelles, J. M. Conrad, and M. H. Shaevitz, *Sterile Neutrino Fits to Short Baseline Data*, [1602.00671](#).
- [68] S. Gariazzo, C. Giunti, M. Laveder, and Y. F. Li, *Updated Global 3+1 Analysis of Short-Baseline Neutrino Oscillations*, [1703.00860](#).

- [69] C. Giunti, X. P. Ji, M. Laveder, Y. F. Li, and B. R. Littlejohn, *Reactor Fuel Fraction Information on the Antineutrino Anomaly*, [1708.01133](#).
- [70] M. Dentler, A. Hernández-Cabezudo, J. Kopp, M. Maltoni, and T. Schwetz, *Sterile Neutrinos or Flux Uncertainties? - Status of the Reactor Anti-Neutrino Anomaly*, [1709.04294](#).
- [71] M. Dentler, A. Hernández-Cabezudo, J. Kopp, P. A. N. Machado, M. Maltoni, I. Martinez-Soler, and T. Schwetz, *Updated global analysis of neutrino oscillations in the presence of eV-scale sterile neutrinos*, [1803.10661](#).
- [72] **MINOS Collaboration**, P. Adamson *et al.*, *Search for sterile neutrinos in MINOS and MINOS+ using a two-detector fit*, Submitted to: *Phys. Rev. Lett.* (2017) [[1710.06488](#)].
- [73] **IceCube Collaboration**, M. G. Aartsen *et al.*, *Searches for Sterile Neutrinos with the IceCube Detector*, *Phys. Rev. Lett.* **117** (2016), no. 7 071801, [[1605.01990](#)].
- [74] B. J. P. Jones, *Sterile neutrinos in cold climates*. PhD thesis, Massachusetts Institute of Technology, 2015. available from <http://hdl.handle.net/1721.1/101327>.
- [75] C. A. Argüelles, *New Physics with Atmospheric Neutrinos*. PhD thesis, University of Wisconsin, Madison, 2015. available from <https://docushare.icecube.wisc.edu/dsweb/Get/Document-75669/tesis.pdf>.
- [76] **MiniBooNE Collaboration**, A. A. Aguilar-Arevalo *et al.*, *A Search for muon neutrino and antineutrino disappearance in MiniBooNE*, *Phys.Rev.Lett.* **103** (2009) 061802, [[0903.2465](#)].
- [77] **MiniBooNE Collaboration**, **SciBooNE Collaboration**, G. Cheng *et al.*, *Dual baseline search for muon antineutrino disappearance at $0.1\text{eV}^2 < \Delta m^2 < 100\text{eV}^2$* , *Phys.Rev.* **D86** (2012) 052009, [[1208.0322](#)].
- [78] R. H. Cyburt, B. D. Fields, K. A. Olive, and T.-H. Yeh, *Big Bang Nucleosynthesis: 2015*, *Rev. Mod. Phys.* **88** (2016) 015004, [[1505.01076](#)].
- [79] **Planck Collaboration**, P. A. R. Ade *et al.*, *Planck 2015 results. XIII. Cosmological parameters*, [1502.01589](#).
- [80] B. Dasgupta and J. Kopp, *A ménage à trois of eV-scale sterile neutrinos, cosmology, and structure formation*, *Phys.Rev.Lett.* **112** (2014) 031803, [[1310.6337](#)].
- [81] S. Hannestad, R. S. Hansen, and T. Tram, *How secret interactions can reconcile sterile neutrinos with cosmology*, *Phys.Rev.Lett.* **112** (2014) 031802, [[1310.5926](#)].
- [82] X. Chu, B. Dasgupta, M. Dentler, J. Kopp, and N. Saviano, *Sterile Neutrinos with Secret Interactions – Cosmological Discord?*, [1806.10629](#).
- [83] C. E. Yaguna, *Sterile neutrino production in models with low reheating temperatures*, *JHEP* **06** (2007) 002, [[0706.0178](#)].
- [84] N. Saviano, A. Mirizzi, O. Pisanti, P. D. Serpico, G. Mangano, *et al.*, *Multi-momentum and multi-flavour active-sterile neutrino oscillations in the early universe: role of neutrino asymmetries and effects on nucleosynthesis*, *Phys.Rev.* **D87** (2013) 073006, [[1302.1200](#)].
- [85] M. Giovannini, H. Kurki-Suonio, and E. Sihvola, *Big bang nucleosynthesis, matter antimatter regions, extra relativistic species, and relic gravitational waves*, *Phys. Rev.* **D66** (2002) 043504, [[astro-ph/0203430](#)].
- [86] F. Bezrukov, A. Chudaykin, and D. Gorbunov, *Hiding an elephant: heavy sterile neutrino with large mixing angle does not contradict cosmology*, [1705.02184](#).
- [87] Y. Farzan, *Ultra-light scalar saving the 3+1 neutrino scheme from the cosmological bounds*, [1907.04271](#).
- [88] J. M. Cline, *Viable secret neutrino interactions with ultralight dark matter*, [1908.02278](#).
- [89] M. Archidiacono, S. Gariazzo, C. Giunti, S. Hannestad, and T. Tram, *Sterile neutrino self-interactions: H_0 tension and short-baseline anomalies*, [2006.12885](#).
- [90] **MiniBooNE Collaboration**, A. Aguilar-Arevalo *et al.*, *Event Excess in the MiniBooNE Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations*, *Phys.Rev.Lett.* **105** (2010) 181801, [[1007.1150](#)].
- [91] **MiniBooNE Collaboration**, “Data release for arXiv:0903.2465.” http://www-boone.fnal.gov/for_physicists/data_release/numu_numubar/.
- [92] **IceCube Collaboration**, M. Aartsen *et al.*, *Determining neutrino oscillation parameters from atmospheric muon neutrino disappearance with three years of IceCube DeepCore data*, [1410.7227](#).
- [93] **IceCube Collaboration**, J. P. Yañez *et al.*, “IceCube Oscillations: 3 years muon neutrino disappearance data.” http://icecube.wisc.edu/science/data/nu_osc.

- [94] F. Dydak, G. Feldman, C. Guyot, J. Merlo, H. Meyer, *et al.*, *A Search for Muon-neutrino Oscillations in the Delta m^2 Range 0.3-eV^2 to 90-eV^2* , *Phys.Lett.* **B134** (1984) 281.
- [95] **Super-Kamiokande Collaboration**, R. Wendell *et al.*, *Atmospheric neutrino oscillation analysis with sub-leading effects in Super-Kamiokande I, II, and III*, *Phys.Rev.* **D81** (2010) 092004, [[1002.3471](#)]. Long author list - awaiting processing.
- [96] **Super-Kamiokande Collaboration**, R. Wendell, *Atmospheric Results from Super-Kamiokande*, *AIP Conf. Proc.* **1666** (2015), no. 1 100001, [[1412.5234](#)].
- [97] **NOvA Collaboration**, P. Adamson *et al.*, *Search for active-sterile neutrino mixing using neutral-current interactions in NOvA*, [1706.04592](#).