

Machine Learning-Informed 3+1 Sterile Neutrino Global Fits using Posterior Density Estimation of Electron Disappearance Data

Joshua Villarreal ^{1*}, Julia Woodward ¹, John Hardin ¹, Janet Conrad ¹

¹Department of Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, 02138, MA, USA.

*Corresponding author(s). E-mail(s): villaj@mit.edu;

Contributing authors: julia785@mit.edu; jmhardin@mit.edu; conrad@mit.edu;

Abstract

Global analyses of particle physics data are integral for validating and scrutinizing published results of experiments. Global fits of anomalous oscillation data which search for one or more eV-scale sterile neutrinos are particularly challenging both to evaluate and to reconcile in the global picture. Fits (especially joint ones) to oscillation data suffer from significant computational burdens, such as likelihood intractability, making traditional Markov Chain-Monte Carlo all but impossible. Given evidence both supporting and challenging beyond Standard Model physics across neutrino experiments of various baselines, energies, and detection techniques, the global search for sterile neutrinos requires additional tools in order to determine whether sterile neutrinos remain a viable solution to unexplained anomalies. Furthermore, both a Bayesian and frequentist interpretation of sterile neutrino data is needed for a complete assessment of longstanding tensions in the field. Techniques from the machine learning subfield of simulation-based inference have a natural application to such a problem. In this contribution, we illustrate some of the outstanding questions of the global picture of light sterile neutrinos by focusing on experiments searching with the disappearance of electron (anti)neutrinos, and look to posterior density estimation strategies to craft answers, including comparisons to a machine-learning-based frequentist approach.

Keywords: Sterile neutrino, simulation based inference, global fit, machine learning

1 Introduction

Within the field of particle physics, progress often comes from analyzing disparate data sets to extract subtle signals. Searches for light, eV-scale sterile neutrinos are a prime example of this “global fit” approach [1–5]. These analyses combine short-baseline data from neutrino experiments with different baselines, energies, and detection technologies, each with unique systematic

uncertainties and sensitivities, to identify the most interesting parameter space for future searches and to assess the probability that anomalous signals are due to specific new physics models. Most often, the models add a single non-interacting, hence “sterile” neutrino to the three neutrinos that interact via the Standard Model’s weak force; hence this class is called 3+1 models.

Within 3+1 studies, and in the case of many global fits to particle physics data, fitters most

often rely on frequentist statistical tests of compatibility of the data. In the frequentist paradigm, a global fit typically proceeds by evaluating a test statistic over a physical model’s parameter space and interpreting its profile by quoting a *confidence level* (CL). The standard recipe profiles the likelihood ratio between a given hypothesis and the best-fit case, assuming that the resulting test statistic follows a χ^2 distribution as described by Wilks’ theorem [6]. While this approximation enables rapid computation of CLs, it relies on a set of assumptions that are not met by searches for small oscillatory signals (such as 3+1) because effective degrees of freedom vary over the parameter space and systematic uncertainties can deviate from Gaussianity [7]. In these situations, the nominal coverage of the inferred CLs can fail consistency checks under repeated pseudoexperiments [8], motivating alternative approaches to inference for complex or non-Gaussian global analyses. A very common example of this is the setup of trials-based tests first described by Feldman and Cousins [9]. However, while full Feldman–Cousins constructions of CLs are exact, they are computationally infeasible to compute at the scale of modern sterile-neutrino data. In previous work [10], we introduced a trials-based framework that uses machine learning (in particular, simulation-based inference [SBI]) to construct well-calibrated CLs without relying on Wilks’ theorem. That work demonstrates that frequentist global fits for sterile-neutrino scenarios can be made both principled and tractable.

This paper provides the Bayesian counterpart to our frequentist method [10]. As our example, we focus on electron-flavor disappearance data in a 3+1 framework; formulating the global analysis by once again leveraging recent developments in SBI to use the same simulator to efficiently generate a large number of pseudoexperiments over a variety of physics hypothesis, and perform density estimation to directly and efficiently estimate the posterior distribution over the parameters. In this work, we compare several state-of-the-art posterior density estimation strategies, based on discrete and continuous normalizing flows and diffusion models, and the extracted posterior from the neural likelihood ratio estimation strategy utilized in Ref. [10]. We then use these learned posteriors to draw parameter constraints and credibility

regions (the Bayesian counterpart to frequentist CLs) from the published reactor-, source-, and accelerator-based $\nu_e/\bar{\nu}_e$ disappearance data. The result is a Bayesian analysis said to be “likelihood free”, and does not rely on Markov Chain Monte Carlo sampling methods to gather posterior estimates leading to significant improvements in run time at inference.

The Bayesian treatment of sterile neutrino data is an important companion to the frequentist version of the analysis. The statistical questions asked by each analysis are not redundant, but complementary methods of model fitting and uncertainty quantification. Answers to both analysis philosophies are informative in a field where results from different experiments can be in tension; in recent times, both fits are commonly reported (e.g., Refs. [11, 12]).

Also, the simulator-level approach lets us propagate non-Gaussian and experiment-specific systematics by construction. In the Bayesian setting, these effects can be marginalized hierarchically and reported transparently via full posterior structure (including multi-modality and parameter degeneracies). Credible regions (the highest posterior density [HPD]) and posterior predictive checks provide diagnostics that complement frequentist coverage studies.

The two approaches do not represent double the effort. The paradigms can share the same pseudoexperiment surrogates and common algorithmic pipelines. In particular, amortized posteriors enable fast maximum-*a posteriori* (MAP) estimates, which can satisfy the ranking criterion of the best-fit point estimates required by Feldman–Cousins-like test statistics [9, 10]. As we show, some posterior estimators (notably SNPE-C in our benchmarks) deliver MAPs with low error at very low latency, making them attractive plug-and-play upgrades for trials-based approaches like Ref. [10].

The methods we recommend in this work can be applicable across many types of global (or even single-experiment) fits within particle physics. In this work, we use the specific example of electron-flavor disappearance in a 3+1 model: a relatively simple case that, nonetheless, tests a physical model dependent on more than one parameter and relies on three qualitatively different data sets taken at reactors (which infer oscillations of

$\bar{\nu}_e$), or radioactive sources and accelerators (oscillations of ν_e). These experiments combine data from particles (ν_e) and antiparticles ($\bar{\nu}_e$) in experiments with complementary baselines, energies, and detector designs, leading to very different systematic uncertainties, emblematic of the issues that global fitters of any particle physics data sets will face.

The goal of the analysis is to determine the best pair of parameters of the 3+1 model for electron neutrino disappearance found by estimating the posterior distribution and constructing credible regions over a grid. Our method trains density estimation strategies on data generated by the simulation code described in Ref. [10], which produces pseudoexperiments capturing experiments’ reported systematic and statistical uncertainties. We then evaluate calibration and accuracy using withheld test-set samples, and use the trained model to fit the published data to obtain HPD regions of credibility that can be compared (qualitatively) with the confidence regions computed from a frequentist framework.

Taken together, the frequentist and Bayesian analyses present a coherent, simulation-driven view of the electron-disappearance sector in 3+1 models. Agreement between HPD regions and trials-based confidence regions corroborates that the conclusions of the global fit are not artifacts. Differences between results can help localize sensitivity to priors, non-Gaussian systematics, ordering rules, or even inconsistencies in the simulator itself. The result is a more robust global picture.

The remainder of this paper is structured as follows. Section 2 briefly reviews the data sets and the 3+1 model used in the example global fit. Section 3 details the simulator construction, training data preparation, and the various posterior estimation strategies evaluated. Section 4 presents the benchmark results for each estimator and applies the chosen method to the ν_e and $\bar{\nu}_e$ disappearance data. Section 5 details the integration of the posterior estimators with the frequentist framework of Ref. [10]. Section 6 discusses the implications of the Bayesian fits and their limitations. Finally, Section 7 summarizes our conclusions.

2 About Electron Neutrino Disappearance

Throughout this paper we will use the 3+1 model for electron flavor neutrino disappearance as our example, although our methods have application to most global fits performed in particle physics. This section provides a very brief introduction to the model we will use and the data sets that will be fit.

The 3+1 model describes potential neutrino oscillations between the three standard light flavors and a sterile neutrino. Our example focuses on whether a neutrino produced as electron-flavored (ν_e) will interact downstream from the production point as a ν_e or whether it will become non-interacting, hence sterile (ν_s). The probability that a produced ν_e continues to interact as a ν_e downstream has an oscillatory dependence, given by:

$$P_{\nu_e \rightarrow \nu_e} = 1 - 4|U_{e4}|^2(1 - |U_{e4}|^2) \sin^2(\Delta_{41}L/E). \quad (1)$$

where L [m], the distance traveled (often called the “baseline”) and E [MeV] the energy of the neutrino, are experimentally chosen parameters. From this formula, one can extract two underlying parameters of the 3+1 model: $|U_{e4}|$, associated with the amplitude of the oscillation; and $\Delta_{41} \equiv 1.27\Delta m_{41}^2$, the frequency of the oscillations in eV^2 . In principle, to keep $P_{\nu_e \rightarrow \nu_e} \in [0, 1]$, we require $|U_{e4}| \in [0, 1]$; however, in our studies we will use $|U_{e4}| \leq 1/\sqrt{2}$ which turns out to fully explore the electron-flavor disappearance parameter space. In a 3+1 model, the formula for the antineutrino disappearance probability is identical to that of neutrinos (Eq. 1).

Table 1 summarizes the experiments used in our electron-flavor global fit. Published data for each are reproduced in the appendix (see Fig. 8).

2.1 Reactor Experiments

The reactor experiments with the most extensive reach in 3+1 parameter space are STEREO [13], PROSPECT [14], DANSS [15], and NEOS [16]. Reactors emit anti-electron neutrinos ($\bar{\nu}_e$) detected via inverse beta decay interactions $\bar{\nu}_e(p, n)e^+$, allowing precise reconstruction of baseline L and antineutrino energy E . To detect sterile neutrino oscillations, the measured event

Table 1: Relevant parameters of experiments used in this study. Characteristic L and E values are approximate, see reference (column 2) for more details. Column 5 provides the reported anomaly significance of a 3+1 signal in numbers of standard deviations (or NR if the best fit point is not reported), though each provides a best fit point regardless of 3+1 preference; for these, see Appendix B.

Name	Ref.	L [m]	E [MeV]	Significance	Description
Reactor Experiments ($\bar{\nu}_e$)					
STEREO	[13]	9–11	1.625–7.125	NR	ILL research reactor, extended detector
PROSPECT	[14]	7–10	0.8–7.2	NR	ORNL research reactor, extended detector
DANSS	[15]	10–12	1–7	NR	Kalinin power reactor, moving baseline system
NEOS/RENO	[16]	224/294	1–10	NR	Hanbit power reactor, two detectors with RENO setting SM flux
Source Experiments (ν_e)					
SAGE & GALLEX	[17–21]	0–0.7 & 0–1.9	0.426, 0.747, 0.752, 0.811, 0.813	2.3σ	^{51}Cr and ^{37}Ar sources; single segment ^{71}Ga target
BEST	[22, 23]	0–1.7	0.426, 0.747, 0.752	$>5\sigma$	^{51}Cr source; two-segment ^{71}Ga target
Accelerator Experiments (ν_e)					
KARMEN/LSND	[24]	17.7/29.8	17.3–52.8	NR	ν_e -carbon cross section measurements from pion decay-at-rest ν_e

rate as a function of L/E is compared to Standard Model (SM) predictions for the event rate given the calculated reactor antineutrino flux. Evidence for a sterile neutrino is a deficit of measured $\bar{\nu}_e$ compared to flux expectations, hence a reduction in rate. The SM-based reactor flux prediction used in this paper is the Huber-Mueller (HM) flux [25], and is a common choice. However, this flux prediction is known to have two unresolved issues when compared with data.

The first issue is that the HM flux over-predicts the total rate of neutrino interactions, an effect called the Reactor Antineutrino Anomaly (RAA) that was initially reported in 2011 [26]. Since that time, multiple new flux models addressing this issue, summarized in Ref. [27], have been put forward. These models range in data-to-prediction ratios from $0.925^{+0.025}_{-0.023}$ for the HKSS model to $0.975^{+0.022}_{-0.021}$ for the KI model, but there is no consensus “best choice” for among these for a new flux prediction [27]. In principle, we can assign a theoretical uncertainty to the HM flux and proceed with our fits. In practice, however, SM explanations for this overall deficit are degenerate with a signature for sterile neutrino oscillations in certain parameter space. Examining Eq. 1, one sees that as Δm^2 becomes large relative to the L/E of a given experiment, oscillations become rapid. Due to experimental resolution, the oscillatory behavior becomes unobservable and the term depending on Δm^2 averages to 1/2, consistent with an overall deficit. To remove the confusion between an SM or sterile neutrino source of a deficit and to reduce the theoretical uncertainty from the flux models, we remove the overall normalization mismatch to perform “shape-only” fits that demand oscillatory behavior. This reflects the view that neutrino oscillations should explicitly be seen to oscillate, but this conservative choice does diminish sensitivity to high values of Δm^2_{41} ($\gtrsim \mathcal{O}(10\text{eV}^2)$).

The second data-to-simulation comparison issue for the HM reactor flux is a deviation in the predicted reactor flux at energies ≈ 5 MeV. Known as the “5 MeV bump”, the $\sim 10\%$ excess is seen in all modern reactor data sets across baselines from ~ 1 m to 1000 m. The source of the effect is not understood, but it is not due to 3+1 oscillations, since that would produce a deficit rather than an excess and also would

show L dependence (see Eq. 1). First-principles models based on reactor fuel ^{235}U and ^{239}Pu consumption aimed at explaining the excess have not successfully fit all of the reactor data, even when introducing an empirical parameter [28]. In order to address this issue, short baseline (SBL) reactor experiments employ two types of mitigation. The first, used by DANSS and NEOS/RENO, fits ratios of data sets across the same E range but at different L -values, which will remove the “bump effect”. DANSS took ratios of data from the same detector at different locations. NEOS and RENO are experiments located at a short (NEOS) and long (RENO) baseline at the same reactor complex, and in the discussion below are always treated as an “experimental pair” [16]. The second, used by STEREO and PROSPECT, fits for the L dependence in each energy bin separately. Due to these mitigations, the reactor results that are used in the fits in this paper have negligible bias from the 5 MeV bump.

2.2 Source Experiments

The SAGE [17, 20], GALLEX [19] and BEST [22, 23] experiments use fluxes produced by radioactive sources with MCi-level decay rates. These sources emit monoenergetic electron neutrinos (ν_e) produced through electron capture in ^{51}Cr and ^{37}Ar sources, and are paired with gallium-filled detectors. The ν_e interactions convert gallium nuclei to germanium ($\nu_e(^{71}\text{Ga}, ^{71}\text{Ge})e^-$), which is periodically removed and counted to compare against well-predicted rates [29]. While low normalization uncertainty allows high sensitivity to non-zero $|U_{e4}|$, poor segmentation limits Δm_{41}^2 sensitivity, restricting measurements to $\Delta m_{41}^2 \gtrsim \mathcal{O}(10 \text{ eV}^2)$. Thus, the reactor and source data sets are complementary in multiple ways.

2.3 Accelerator-based Decay-at-rest Experiments

The third class of experiments for our example fits uses fluxes produced at two accelerator-based decay-at-rest sources. These sources impinge 800 MeV protons on a target, producing pions that stop and produce a decay-at-rest chain resulting in 4π emission of ν_μ , $\bar{\nu}_\mu$, and ν_e in equal fractions with a well-understood energy spectrum and a normalization known to the $\sim 10\%$ level.

The highly reproducible nature of the decay-at-rest sources allows cross comparison of widely-separated experiments, even at facilities on separate continents. This study compares data from the KARMEN detector [30] at Rutherford Lab in the UK to the LSND detector [31] at LAMPF in the US. The experiments were located at different L from their respective decay-at-rest sources: KARMEN is at baseline of 17.7 m and LSND is at baseline of 29.8 m. Thus, comparison of the relative rate of ν_e interactions in the two experiments, published in the form of energy-dependent cross-sections, allows a search for the L dependent effects of ν_e 3+1 oscillations. In the global fit presented here, we use the comparison study presented in Ref. [24]. In the fits and discussion below, like the NEOS/RENO ratio, KARMEN and LSND will be treated as an inseparable experimental pair.

3 Training, Evaluation, and Selection of Posterior Density Estimator

3.1 Training data preparation

The Monte Carlo simulation used in this work, `sblmc`, is described in detail in Refs. [1, 2], and is the same as used in the frequentist SBI global fit approach paper in Ref. [10]. In short, the codebase is designed for calculations of model log-likelihoods based on experimental data, requiring either a covariance matrix or a set of pull parameters provided by experiments during result releases. In addition to a set of model expectations, the code can be used to generate a set of realizations for any sterile neutrino model for any combination of experiments.

To generate the training- and test-set data used in this study, we define an evenly spaced 50×50 grid in $\log_{10} U_{e4}$, $\log_{10} \Delta m_{41}^2$ for $U_{e4} \in [0.001, 1/\sqrt{2}]$ and $\Delta m_{41}^2 \in [0.01, 100] \text{ eV}^2$. For each of the experiments listed in Tab. 1, we generate 1000 mock realizations per parameter point resulting in 2,500,000 total pseudoexperiments.

3.2 Posterior Estimation Strategies Considered Here

To estimate the posterior density, we must begin with the following:

- Model parameters θ , which, in this contribution, are shorthand for the sterile oscillation parameters $\theta \equiv (U_{e4}, \Delta m_{41}^2)$ in Eq. 1,
- A pre-defined prior distribution $\mathbb{P}(\theta)$, and
- A simulator capable of sampling experimental data $\mathbf{x} \sim \mathbb{P}(\mathbf{x}|\theta)$.

Given these inputs, the strategy is to infer, for observed experimental data $\mathbf{x}_{\text{obs}}$, the posterior probability density $\mathbb{P}(\theta|\mathbf{x}_{\text{obs}})$.

Conditional density estimators approximate the posterior distribution with neural networks. That is, the true posterior $\mathbb{P}(\theta|\mathbf{x}) \approx q_{\psi}(\theta)$, with parameters $\psi \approx F(\mathbf{x}|\phi)$ for F the neural network having parameters ϕ trained with loss criterion $\mathcal{L}(\phi) = -\sum_i \log q_{F(\mathbf{x}_i|\phi)}(\theta_i)$. In the case where the neural network F is sufficiently expressive and the size of the training data is large enough, as $\mathcal{L}$ is minimized, $q_{F(\mathbf{x}_i|\phi)} \rightarrow \mathbb{P}(\theta|\mathbf{x}_i)$. Such a conditional density estimator can then be used to approximate the posterior probability density as $\mathbb{P}(\theta|\mathbf{x}_{\text{obs}}) \approx q_{F(\mathbf{x}_{\text{obs}}|\phi)}(\theta)$. Presently, neural networks called *normalizing flows* which learn a sequence of invertible transformations between two probability distributions are used for such tasks and have shown promise in a variety of applications [32].

All of the strategies except for the direct amortized neural likelihood ratio estimation were investigated using implementations provided by the `sbi` Python package [33], itself a heavy user of `PyTorch` [34]. The method described in Sec. 3.2.4 was implemented using `tensorflow` [35],

3.2.1 Sequential Neural Posterior Estimation

For a complete overview, refer to Ref. [36]. In *sequential neural posterior estimation*, a conditional density estimator is refined across multiple rounds of training to approximate the true posterior distribution $\mathbb{P}(\theta|\mathbf{x}_{\text{obs}})$. In the first round, parameters $\theta \sim \mathbb{P}(\theta)$ are drawn from the prior and used to generate simulated data $\mathbf{x}$. The neural network $F(\mathbf{x}|\phi)$ is then trained such that the induced density $q_F(\theta)$ approximates the posterior distribution $\mathbb{P}(\theta|\mathbf{x})$. In subsequent rounds of training, parameters are sampled instead from a proposal distribution $\tilde{\mathbb{P}}(\theta)$ which is meant to concentrate probability mass in regions more informative about observed data $\mathbf{x}_{\text{obs}}$.

However, training on these samples can bias the estimator. In particular, when $\theta \sim \tilde{\mathbb{P}}(\theta)$ in the second step of the training sequence, minimizing the loss function $\mathcal{L}(\phi) = -\sum_i \log q_{F(\mathbf{x}_i|\phi)}(\theta_i)$ over these samples forces $q_F(\theta) \rightarrow \mathbb{P}(\theta|\mathbf{x}) \frac{\tilde{\mathbb{P}}(\theta)\mathbb{P}(\mathbf{x})}{\mathbb{P}(\theta)\tilde{\mathbb{P}}(\mathbf{x})}$ (for $\tilde{\mathbb{P}}(\mathbf{x}) = \int_{\theta} \tilde{\mathbb{P}}(\theta)\mathbb{P}(\mathbf{x}|\theta)$). *Automatic posterior transformation* (which, based on precedent, is known as **SNPE-C**, indicating the third generation of previously-developed *sequential neural posterior estimation* strategies) corrects for this bias by modifying the loss function, training the network to minimize $\tilde{\mathcal{L}}(\phi) = -\sum_i \log[q_{F(\mathbf{x}_i|\phi)}\tilde{\mathbb{P}}(\theta)/\mathbb{P}(\theta)]$. Over multiple rounds of training, the network weights ϕ are modified by selecting simulations from sequentially updated proposals $\tilde{\mathbb{P}}(\theta)$. As the network sends $q_{F(\mathbf{x}_i|\phi)}(\theta) \rightarrow \mathbb{P}(\theta|\mathbf{x})$, $\mathbb{E}[\tilde{\mathcal{L}}]$ is minimized, agnostic to the proposal $\tilde{\mathbb{P}}(\theta)$. Therefore, an advantage of **SNPE-C** is that the network can leverage data from multiple training rounds simultaneously by simply summing their respective losses. With some additional technical details (to which we refer the reader to Ref. [36]), one can use normalizing flows (parameterized by ψ) for the prior, proposal, and posterior estimates, taking advantage of advances in machine learning.

3.2.2 Flow Matching

Flow matching posterior estimation (FMPE) replaces the discrete normalizing flow in the preceding section with a continuous normalizing flow, where the transformation from a simple base distribution to the posterior distribution is defined by the solution of an ordinary differential equation. First introduced in Refs. [37, 38], we describe the general idea below for completeness.

A *continuous normalizing flow*, parameterized by a continuous “time” variable $t \in [0, 1]$, much like a discrete normalizing flow, also maps a base to a target distribution. For base distribution q_0 and target distribution q , and for invertible mapping $\psi_x : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the a velocity field $v_{t,x}$ governs this transformation with an ordinary differential equation:

$$\frac{d}{dt}\psi_{t,x}(\theta) = v_{t,x}(\psi_{t,x}(\theta)); \psi_{0,x}(\theta) = \theta.$$

The final density can be found by integration:

$$q(\theta|x) = q_0(\theta) \exp \left(- \int_0^1 \operatorname{div} v_{t,x}(\theta_t) dt \right).$$

The goal of flow matching is to use a neural network to approximate the velocity field $v_{t,x}$. To ensure the target probability path is found, a well-constructed loss function does the trick; we refer the reader to Refs. [37, 38] for the details. The key insight is that discrete normalizing flows depend on a prespecified density estimator architecture, whereas continuous normalizing flows can be more expressive by learning arbitrary probability paths.

3.2.3 Neural Posterior Score Estimation

Neural Posterior Score Estimation (NPSE) formulates simulation-based inference as a conditional diffusion process that learns to approximate the score of the posterior distribution [39, 40]. Details are listed here. NPSE introduces a forward diffusion process that gradually perturbs θ toward a tractable reference distribution (e.g., a standard Gaussian) according to the stochastic differential equation (SDE)

$$d\theta_t = f(\theta_t, t) dt + g(t) dw_t, \quad (2)$$

where f and g are drift and diffusion coefficients, and $(w_t)_{t \geq 0}$ is a standard Brownian motion. The reverse-time process defining the generative model evolves as

$$d\bar{\theta}_t = [-f(\bar{\theta}_t, T-t) + g^2(T-t)\nabla_{\theta} \log p_{T-t}(\bar{\theta}_t|x)]dt + g(T-t) dw_t. \quad (3)$$

Under this formulation, generating posterior samples corresponds to simulating the reverse diffusion process conditioned on the observation x_{obs} . Since the true score $\nabla_{\theta} \log p_t(\theta|x)$ is intractable, NPSE trains a neural network $s_{\psi}(\theta, x, t)$ to approximate it by minimizing a weighted Fisher divergence, with the internal expectation approximated via Monte Carlo sampling from the prior and the forward diffusion process.

For deterministic inference or density evaluation, NPSE can alternatively employ the probability flow ODE associated with Eq. (2),

$$\frac{d\theta_t}{dt} = f(\theta_t, t) - \frac{1}{2}g^2(t)\nabla_{\theta} \log p_t(\theta|x), \quad (4)$$

which shares the same marginal distributions as the stochastic process. Once the score network is trained, posterior samples are drawn by simulating the reverse-time SDE or its deterministic counterpart, initialized from the reference distribution. This provides a flexible and likelihood-free approach to posterior inference with strong generative capabilities.


3.2.4 Direct Amortized Neural Likelihood Ratio Estimation

SNPE-C, NPSE, and FM are learning tasks in which the goal is to find a neural network estimate for the posterior distribution $\mathbb{P}(\theta|\mathbf{x})$. For most practical applications, the true posterior distribution is unknown, so optimizing a machine learning model to estimate the posterior is known as an unsupervised learning task. On the other hand, supervised learning tasks require knowledge of the true labels or target values at training time (not necessarily at the inference step), and tend to be more straightforward.


One such strategy, Direct Amortized Neural Likelihood Ratio Estimation (DNRE) [41], has been used to conduct global fits of anomalous muon neutrino disappearance in a frequentist setting [10]. The details of implementation are described therein, but are summarized briefly for completeness here. Samples $\{(\mathbf{x}_i, \theta_i)\}_{i=1}^N$ are collected via the method of Sec. 3.1, and in addition, M additional prior-sampled model parameters are collected: $\{\theta'_j\}_{j=1}^N$. The training set is constructed by two steps of matrix augmentation: samples $\{(\mathbf{x}_i, \theta_i, \theta'_i)\}_{i=1}^N$ are assigned the label 1, while samples $\{(\mathbf{x}_i, \theta'_i, \theta_i)\}_{i=1}^N$ are assigned the label 0. A classifier $\mathbf{d}_{\phi}(\mathbf{x}, \theta, \theta')$ is trained using binary cross-entropy loss to predict the correct label given $(\mathbf{x}, \theta, \theta')$; its goal is to determine whether the experimental (Monte Carlo) data was generated from θ or θ' . The network's output can be transformed into an estimate of the likelihood ratio,

$$\frac{\mathbb{P}(\mathbf{x}|\theta)}{\mathbb{P}(\mathbf{x}|\theta')} \approx \hat{r}(\mathbf{x}|\theta, \theta') = \frac{\mathbf{d}_{\phi}(\mathbf{x}, \theta, \theta')}{1 - \mathbf{d}_{\phi}(\mathbf{x}, \theta, \theta')},$$


its intended use case. However, Ref. [41] points out that the classifier's output can also be used as an



(a) Distribution of log-scores for each flow-based posterior estimator. While the mean log-score is presented as a summary statistic in Tab. 3, it does not effectively capture the substructure in the metric distributions; each log-score distribution is bimodal, representing the dichotomy of pseudoexperiments to be generated for signal- and null-like physics parameters. The bump at lower values of log-score corresponds to the flow assigning an approximately uniform probability over parameter space when the pseudoexperiment is thrown from a set of physics parameters to which the global fit is not sensitive (indistinguishable from null). The bump at higher values of log-score correspond to the networks assigning a high posterior probability at the true injected parameters when the global fit can distinguish those parameters meaningfully (a signal-like sample). The high log score for null-like samples in the case of the NPSE posterior estimator is physically unreasonable.



(b) Maximum *a posteriori* (MAP) squared errors on out-of-sample test entries with $\log_{10} U_{e4} > -0.7$, where we expect the global fit to be able to identify a signal.



(c) Empirical cumulative distribution of rank scores for each posterior estimator and the DNRE baseline.

Fig. 1: Comparison of the posterior density estimation methods based on flows and the neural likelihood ratio estimator from Ref. [10] using (a) log-score distributions, (b) MAP error statistics, and (c) rank score distributions.

estimate of the posterior probability density:

$$\log \mathbb{P}(\theta|\mathbf{x}) \approx -\text{LogSumExp}[-\log \hat{r}(\mathbf{x}|\theta, \theta'_i)] + \log M + \log \mathbb{P}(\theta). \quad (5)$$

Given DNRE's relationship to the SBI global fits of sterile neutrino models, we include DNRE as a posterior estimation strategy to test in this work.

We also note that a better-performing posterior estimation strategy (particularly one with a better maximum *a posteriori* error, see Sec. 3.3) than DNRE could be used in place of Eq. 5 in the best fit point determination of the method described in Ref. [10], further motivating its inclusion in this body of work.

Table 2: Performance of posterior estimation strategies presented in Sec. 3, including scores for network calibration and timing considerations. Five-fold cross validation was performed to ensure result consistency among disparate validation sets (reported as mean $\pm$ standard deviation).

Strategy	Log-score	MAP error	Rank score
SNPE-C	-0.803 ± 0.716	0.141 ± 0.152	0.025 ± 0.050
FMPE	-0.701 ± 0.078	0.102 ± 0.038	0.167 ± 0.322
NPSE	2.157 ± 0.186	0.291 ± 0.109	< 0.001
DNRE	-1.33 ± 0.031	0.086 ± 0.014	0.0032 ± 0.004

3.3 Evaluation metrics

Given a posterior density estimator $\hat{\mathbb{P}}(\theta|\mathbf{x})$, certain desirable characteristics motivate performance metrics which can be cross-compared between multiple algorithms, and should quantify the inference strategy’s precision, accuracy, and calibration in an unbiased way.

As such, for an experimental sample $\mathbf{x}_i$ simulated from oscillation parameters θ_i , a good posterior estimation strategy will have a high log-score, $\log \mathbb{P}(\theta_i|\mathbf{x}_i)$, a measure of the probabilistic weight assigned at the true underlying oscillation parameters θ_i . From an estimated posterior distribution, one can extract the maximum *a posteriori* (MAP) $\hat{\theta}_i^{\text{MAP}} = \text{argmax}_{\theta} \hat{\mathbb{P}}(\theta|\mathbf{x}_i)$.¹ The accuracy of the MAP estimate (computed on the same grid in $U_{e4}, \Delta m_{41}^2$ as the data was generated) can be quantified by the MAP mean squared error $\|\theta_i - \hat{\theta}_i^{\text{MAP}}\|_2^2$ in regions where one expects the injected signal parameters to be well-reconstructed ($\log_{10} U_{e4} > -0.7$ in the case of the global fit).

The calibration of the posterior estimator is determined by the rank score. Recall that in a Bayesian analysis, model parameters θ are treated as random variables, whose probability distribution is determined by the posterior. As such, given $\mathbf{x}_i$, the corresponding model parameters θ_i should be thought of as a random draw according to the posterior: $\theta_i \sim \mathbb{P}(\theta|\mathbf{x}_i)$. For the true posterior distribution, if one were to draw N samples of θ from the posterior, one would expect that when these N samples and the true model parameters θ_i are ranked according to their log-scores, the rank of

θ_i ,

$$\text{rank}(\theta_i) = 1 + \sum_{j=1}^N \mathbb{1}(\log \mathbb{P}(\theta_j|\mathbf{x}_i) > \log \mathbb{P}(\theta_i|\mathbf{x}_i))$$

should be uniformly distributed. A well-calibrated posterior estimator $\hat{\mathbb{P}}(\theta|\mathbf{x}_i)$ will share this property. The measure of the rank uniformity of the true oscillation parameters is referred to as the rank score. To determine rank uniformity, we use a Kolmogorov-Smirnov test to verify that $\text{rank}(\theta_i) \sim \text{DUnif}(1, N + 1)$. The p -value of the test is presented in Tab. 3.

These three metrics are the basis of comparison for the posterior density estimation strategies presented in Sec. 3.2. Cross-validated scores across many MC realizations in the test set are summarized in Tab. 3 and Fig. 1.

3.4 Discussion of Flow Choice

Based on the results of the network evaluation tasks presented in tables 2 and 3 and Fig. 1, the clear front-runner for a precise and accurate posterior density estimation model for the task considered in this work is FMPE. Across each of the five folds of the validation data, FMPE most consistently provides the strongest (and physically reasonable) log-scores, MAP errors, and rank scores, indicating a well-calibrated model. For the task of a standalone likelihood-free Bayesian inference of global electron-neutrino disappearance data, we recommend and proceed with FMPE as the chosen strategy for posterior estimation.

However, we point out that the discrete normalizing flow-based SNPE-C method is especially impressive in its faster training and MAP evaluation time, as well as its competitive MAP errors. As such, it is our preferred choice for upgrading

¹The MAP estimate also takes an important role as the best fit for SBI frequentist trials-based fitting strategies, like that presented in Ref. [10]

Table 3: Performance of posterior estimation strategies presented in Sec. 3, including scores for network calibration, and timing considerations for both training and MAP computation (a standardized task for each network, especially relevant for the best fit point determination and test statistic construction necessary for the method presented in Ref. [10]). Equivalents computed using five-fold cross validation are given in Tab. 2.

Strategy	Calibration metric			Time [s]	
	Log-score	MAP error	Rank score	Network training	MAP computation (avg.)
SNPE-C	−0.567	0.035	0.030	2,864	0.186
FMPE	−0.702	0.056	0.067	46,639	23.442
NPSE	2.160	0.302	< 0.001	170,014	631.371
DNRE	−1.335	0.132	< 0.001	986,375	503.403
See also	Fig. 1a	Fig. 1b	Fig. 1c	-	-




Fig. 2: Using simulated data, posterior density estimation performed by FMPE on out-of-sample null-like (left) and signal-like (right) pseudoexperimental throws. The fit to the signal sample includes an inset demonstrating excellent agreement between the sharply peaked posterior estimate and the true injected sterile oscillation parameters. These fits are to all of the electron neutrino disappearance experiments presented in Tab. 1.

the frequentist trials-based global fitting framework presented in Ref. [10]; namely, for use in the MAP estimation which satisfies the ranking criterion of the Feldman-Cousins test statistic (see Sec. 5).

4 The ν_e Disappearance Global Fit with Flow Matching Posterior Density Estimation

We will now apply our methods to the electron flavor oscillation data to highlight the value of this approach.




Fig. 3: Posterior density estimated with FMPE for the published $\nu_e/\bar{\nu}_e$ disappearance data in Tab. 1. White contours show highest-posterior-density credibility regions. The MAP estimate lies at $U_{e4} = 0.28$ ($\sin^2 2\theta_{ee} = 0.29$), $\Delta m_{41}^2 = 48.69 \text{ eV}^2$. Edge effects cause the normalization seen in the source experiments to taper at large Δm_{41}^2 (see discussion). This posterior approximates the MCMC result in Fig. 7 and is qualitatively consistent with the frequentist allowed region in Fig. 6.

4.1 The Expectation for Classes of Experiments

For the classes of experiments considered in this work, previously published results provide clues as what reasonable global fits to ν_e -disappearance experiments should find. Before considering the results of our fits, let us consider the expectations.

First, consider the reactor experiments. As discussed in Sec. 2.1, we have renormalized the reactor data sets to unity to remove the potential confusion between a high Δm_{41}^2 signal causing rapid oscillations and theoretical uncertainty in the HM flux. We have also followed the steps recommended by each experiment to remove sensitivity to the 5 MeV bump. With these corrections, the reactor experiments report results that are generally consistent with the Standard Model. We

point out, however, that while the newest publication from DANSS sees a 2.3σ preference for 3+1, with a best fit point at $(\sin^2 2\theta_{ee}, \Delta m_{41}^2) = (0.07, 0.34 \text{ eV}^2)$ [42], the dataset implemented in `sb1mc` is significantly smaller and consistent with the Standard Model [15]. The RENO+NEOS search excludes even the best fit point of the larger set at 95% confidence [16]. Therefore, taken as a whole, we expect our treatment of the reactor experiments to overall exclude the 3+1 model over the parameter space where they are sensitive.

On the other hand, the source experiments are expected to collectively prefer 3+1 over the Standard Model, with null rejection at $\gtrsim 5\sigma$ significance when taken together [43]. This is a strikingly large signal. No SM source for the deficit seen in these experiments has been identified. The

ν_e fluxes from monoenergetic electron-capture-decaying isotopes like argon-37 and chromium-51 are well-modeled, and counting of germanium-71 atoms in liquid gallium is reliable.

Recall that the 3+1 model we are testing requires the same oscillation parameters for ν_e and $\bar{\nu}_e$ fluxes. Thus, given this disagreement in source versus reactor results, it is interesting to look to the accelerator-based data set, which is also from a ν_e flux. While those results reported are not statistically very strong, they are consistent with null.

Overall, the source experiments are expected to dominate the global fit, but the significance will be reduced due to exclusions from the reactor and accelerator experiments.

Given these expectations, which were guided by the published frequentist results, let us proceed to consider what the SBI-based Bayesian global fit reports based on publicly available experimental data presented the appendix (in Fig. 8).

4.2 The Simulation-Based Inference Global Fit for Simulated and Published Results

Sample fits using FMPE on Monte Carlo pseudo-data are given by Fig. 2. The fit of the trained FMPE posterior density estimator to published ν_e and $\bar{\nu}_e$ disappearance data is given by Fig. 3. One observes that in the global fit, the signal from the source experiments more than off-sets the exclusion from the reactor experiments leaving an allowed region. We discuss how each class of experiment contributed below.

4.3 Fits to Each Experimental Class

Fits to each of the experiment classes (reactor, source and accelerator) are shown in Fig. 4. Based on the overview presented in Sec. 4.1, likelihood-free inference of the three classes are in good agreement with historical precedent. In particular, fits to the reactors and the accelerators show no significant preference for a 3+1 signal, while the source experiments show a strong preference for a sterile neutrino. Moreover, the allowed region preferred by the source experiments remains in the global picture (Fig. 3), though the permissible model parameter space is limited to $\Delta m_{41}^2 \gtrsim$

10 eV², driven by the exclusion set by the reactor experiments. Moreover, the Bayesian fits are in good agreement with their frequentist counterparts, using the method of Ref. [10]; see Fig. 5.

5 Upgrading the Trials-Based Frequentist Global Fit

In a frequentist global fit, the most powerful test-statistic is the likelihood ratio $\hat{r}(x, \theta, \hat{\theta}) = \frac{\mathbb{P}(x|\theta)}{\mathbb{P}(x|\hat{\theta})}$, where $\hat{\theta}$ denotes the parameter values that maximize the likelihood. In our frequentist muon neutrino disappearance analysis [10], finding the maximum likelihood estimator (MLE) is challenging within the SBI pipeline, and we instead take $\hat{\theta}$ to maximize the posterior distribution which we obtain from DNRE. Eq. 5 provides a prescription for posterior estimation from DNRE likelihood ratios. This is crucially dependent on hyper-parameter M , which is the number of samples drawn from the prior over which likelihood ratios are averaged. Specifically,

$$\frac{\mathbb{P}(x|\theta)}{\mathbb{P}(x)} = \frac{1}{M} \sum_i^M \hat{r}(x|\theta, \theta'_i). \quad (6)$$

Increasing M improves coverage accuracy [41], but comes at the cost of significant reduction in computational efficiency. To balance these tradeoffs, in our frequentist study we fixed $M = 600$. With this choice, evaluating likelihood ratios over the grid of parameter space for a single experiment sample took $\mathcal{O}(10)$ s.

While this is a significant improvement from computing the test-statistic using traditional Feldman-Cousins, applying a normalizing flow yields an additional speed-up (Table 3). Although the network training time for the normalizing flows is notably greater than for DNRE, the largest bottleneck in the frequentist study is MAP generation. The strategy with the fastest MAP computation average is SNPE-C by a factor of around 200. Given the higher log and rank score and lower MAP error (Table 3), we anticipate substituting SNPE-C for the DNRE posterior estimation into the frequentist pipeline will yield both more accurate MAP estimates and greater computational efficiency.




Fig. 4: Fits to published data from each of the three experiment classes (reactors, including STEREO, PROSPECT, DANSS, and NEOS; sources, including SAGE/GALLEX and BEST; accelerators, including the Karmen-LSND cross section analysis) using FMPE. These agree qualitatively with the frequentist allowed regions in Fig. 5.




Fig. 5: Acceptance regions for each experiment class within the frequentist framework, computed using the DNRE posterior evaluation.

To find the MAP using SNPE-C, we train the flow on the same data used to train the DNRE, which is still needed for likelihood ratio estimation. For each pseudo-experiment, we sample 50,000 points from the posterior. The MAP is the parameter point with the highest log-probability. The frequentist confidence levels using SNPE-C and DNRE to compute the MAP are shown in Fig. 6. Both methods yield roughly the same allowed region. The SNPE-C method finds a MAP located at $U_{e4} = 0.31$, $\Delta m_{41}^2 = 13.36 \text{eV}^2$, while DNRE finds a MAP located at $U_{e4} = 0.28$, $\Delta m_{41}^2 = 18.42 \text{eV}^2$.

On the same grid and for the same data sample, evaluating likelihood ratios over the parameter space using SNPE-C for posterior estimation takes $\sim 0.1\text{s}$, corresponding to a $\mathcal{O}(100)$ -fold speed up. The use of SNPE-C also removes the dependence on the hyper-parameter M , which no longer needs to be optimized.

A full 3+1 neutrino global fit involves scanning over three independent parameters, greatly expanding the parameter space. As it stands, performing a three-parameter fit using DNRE-based posterior estimation is computationally infeasible. The adoption of SNPE-C overcomes this limitation, making such analyses possible.




Fig. 6: Frequentist allowed regions for the global $\nu_e/\bar{\nu}_e$ disappearance fit, using the SNPE-C posterior density to estimate the MAP. Red and green denote the 90% and 95% confidence regions. The DNRE-derived 95% confidence level is also shown, following the procedure in [10]. These regions can be qualitatively compared with the Bayesian credibility regions from FMPE (Fig. 3) and the high fidelity MCMC global fit in Fig. 7.




Fig. 7: Posterior density estimated with `sblmc` for the published $\nu_e/\bar{\nu}_e$ disappearance data in Tab. 1. Colored regions indicate highest-posterior-density credibility intervals. The MAP estimate occurs at $U_{e4} = 0.31$ ($\sin^2 2\theta_{ee} = 0.35$), $\Delta m_{41}^2 = 13.7 \text{ eV}^2$. See Ref. [1] for details of the method. This is the exact posterior approximated in Fig. 3 and is qualitatively consistent with the features in Fig. 6.

6 Discussion

6.1 Fit Comparisons

In the case of the global fit to all available electron-flavor disappearance data, excellent topological agreement can be seen between the directly comparable true Markov Chain Monte Carlo fit (Fig. 7) and FMPE (Fig. 3). Moreover, amortized inference with flow matching is significantly faster (~ 10 s) than the posterior sampling performed by `sblmc` ($\sim 10^4$ s for convergence). The identified maximum *a posteriori* from the high-fidelity Bayesian fit is in good agreement with that of the flow matching, though we note a slight discrepancy in the found value of Δm_{41}^2 . This is unsurprising – in frequentist fits, the shallow minimum of the $\Delta\chi^2$ landscape due to the influence of the gallium anomaly on the fit is a known artifact. Poor detector segmentation of the source experiments means poor sensitivity to the mass splitting. Topological agreement can also be observed between the SBI-based global estimated posterior in Fig. 3 and the SBI-powered trials-based method (developed in Ref. [10]) in Fig. 6. There is almost exact agreement between the MAP found by the true Markov Chain Monte Carlo fit and that from SNPE-C (Sec. 5). With one method a Bayesian approach and the other a frequentist approach, the fit results are not directly comparable but do indicate similar localized signal preference driven by the source experiments.

Excellent topological agreement in preferred model parameter space for experiments split by class (reactors, sources, and accelerators) can be seen between the Bayesian (Fig. 4) and frequentist (Fig. 5) SBI approaches. Both the reactor experiments and accelerator experiments agree in their regions of exclusion, and the source experiments agree in their preference for a normalization effect consistent with the gallium anomaly. Comparisons of both SBI approaches for single experiments are also shown in Fig. 9. While STEREO, PROSPECT, DANSS, and BEST appear to be in good agreement with one another, the frequentist fit for NEOS draws a closed 95% CL, while the Bayesian fit seems to prefer a result more consistent with null. Regions of high estimated posterior probability assigned by FMPE seem to agree with the region preferred by the frequentist fit. It is worth noting that at a 99% confidence level, the

NEOS frequentist fit draws an exclusion in agreement with the Bayesian counterpart. A similar story emerges in the comparisons of SAGE and GALLEX; despite the weak exclusions drawn by FMPE, regions of parameter space preferred by the fit are in good agreement with the confidence level drawn by the method from Ref. [10].

6.2 Limitations

Given that field expectations are usually presented in a frequentist framework, care must be taken when conducting Bayesian analyses of particle physics data. Credibility regions are not subject to the same coverage guarantees as frequentist confidence levels [44, 45]. The choice of prior (defined implicitly by the model parameter space covered by the training data) can also strongly bias the posterior, credibility regions, and maximum *a posteriori* in low-statistics or degenerate regions. Hard parameter boundaries (introduced via the uniform prior in this study) can give rise to edge effects that a density estimator may amplify. In higher parameter space dimensions, posterior mass can accumulate on thin, highly non-intuitive manifolds due to the concentration-of-measure phenomenon; low-dimensional marginal plots can therefore be misleading, hiding multimodality, degeneracy ridges, and the fact that MAP estimates are often unrepresentative of typical posterior mass [46]. Communicating the geometry of the posterior thus becomes inherently challenging.

While conducting inference with flows is generally faster than a high-fidelity MCMC framework, significant up-front costs in data generation and network training must be considered. Before any analysis can be performed, a sufficient amount of training data must be simulated, inefficient for one-off analyses. There is also an amortization mismatch: once trained, the network is fixed, and changes to the prior, experimental configuration, or simulator require retraining of the flow. As demonstrated in this study, posterior estimation quality depends strongly on network architecture.

Finally, there are challenges of parameter inference using simulation-based inference generally. Biases or mismodeling in the simulator propagates errors into the posterior estimator; anecdotally, however, we point out that unexpected behaviors of the flows’ helped us to narrow down previously undiscovered bugs in the simulation. When

observed data is fed to a flow, and that data may fall outside the training distribution by way of, for instance, unmodeled systematic effects, possibly overconfident (or even nonsensical) posteriors can arise [10]. Inference failures, especially in tails or low-likelihood regions, can be subtle and difficult to diagnose, requiring specialized tests.

In spite of these limitations, amortized posterior density estimation still remains a powerful tool in the analysis pipeline. When carefully validated and used alongside frequentist methods, they can illuminate complimentary structure in the data and provide rapid exploratory insights impractical with sampling-based methods alone. Methods presented in this work are not aimed to replace existing high-fidelity techniques; instead, they are additional tool in the arsenal of global fitters to serve as an additional diagnostic and interpretive tool, help to identify modeling issues, and rapidly iterate on complex particle physics models.

7 Conclusion

In this work, we performed a comprehensive evaluation of modern neural network-based posterior estimation techniques for sterile neutrino oscillation analyses via electron-flavor disappearance, comparing sequential neural posterior estimation, flow matching, and diffusion-based neural posterior score estimation against a direct amortized neural likelihood-ratio baseline. Balancing considerations to accuracy, precision, and computational efficiency, we recommend flow matching for the purposes of sterile neutrino global fits, and recommend sequential neural posterior estimation to improve a previously developed frequentist global fit pipeline. Together, techniques considered in this work enable rapid and flexible global fits retaining compatibility with established methods.

Using these posterior estimators, we performed a global analysis of electron-flavor disappearance data collected from reactor, source, and accelerator experiments, representative of the complexities of a full-fledged 3+1 global fit. By performing fits stratified by experiment and experiment class, we illustrate how neural posterior estimators (with an amortized data generation cost) can help to disentangle the relative contributions of each dataset revealing how the global 3+1 picture may fit

together. Importantly, we showed that the resulting posterior surfaces closely match both high-fidelity Bayesian MCMC sampling and frequentist simulation-based inference approaches. This agreement provides evidence that neural networks can conduct both accurate and interpretable inferences at a fraction of the computational cost.

Taken together, our results indicate that flow-based posterior estimation offers a promising and practical complement to existing analysis paradigms in neutrino oscillation physics, groundwork for an end-to-end simulation-based inference 3+1 global fit, and clues for analysis methods of more complex beyond Standard Model scenarios. While care must be taken in prior specification, calibration, and validation, these methods provide a powerful means of accelerating global fits, enriching the interplay between Bayesian and frequentist perspectives. As sterile neutrino searches continue to evolve and datasets grow in complexity, neural posterior estimation offers a valuable path toward more flexible, scalable, and transparent inference pipelines capable of supporting the next generation of precision oscillation analyses.

Acknowledgements. We are grateful to Phil Harris at MIT and Gabriel Collin at the University of Adelaide for helpful discussions at the birth of this project. JV, JW, JH, and JC thank MIT for support on this project. This material is based upon work supported in part by the National Science Foundation Graduate Research Fellowship under Grant No. 2141064.

A Frequentist Acceptance Regions by Experiment Class


Allowed and exclusion regions for each class of experiment at 95% confidence, shown in Fig. 5.

B Published Data Distributions


Published data is made available in Fig. 8 for convenience.

C Single-Experiment Fits


Comparisons of the frequentist trials-based SBI approach (from Ref. [10]) to the posterior estimates from flow matching are shown in Fig. 9.




(a) Ratio of data to the no-oscillation model for STEREO Phase I and II, distributed among 6 cells [13].



(b) PROSPECT predicted spectrum at 10 different baselines, weighted by the ratio of total observed and predicted event rate in each energy bin [14].




(c) Ratio of positron energy at the bottom and top detector positions in DANSS [42].




(d) Ratio of NEOS and predicted RENO prompt energy spectra. The RENO $\bar{\nu}_e$ flux is used to minimize NEOS dependence on reactor flux [16].


Fig. 8: Published data (black) and averaged `sb1mc` realizations generated at the published best fit (blue) for each experiment. For some experiments, the prediction under the null hypothesis (no sterile oscillation) is shown in gray.



(e) Ratio of measured to expected ^{71}Ge event rates in the Gallium experiments [21].



(f) Production rate in ^{71}Ge atoms/day normalized by the reference time for the inner (left) and outer (right) volumes in BEST [22].



(g) KARMEN (left) and LSND (right) cross section measurements at various neutrino energies [24].

Fig. 8: (continued)




Fig. 9: Approximate posterior distributions and highest probability regions using FMPE for each of the single experiments, as well as the combined KARMEN/LSND cross section measurement. Results from the trials-based DNRE construction of frequentist confidence intervals, from Ref. [10], are also overlaid.

References

- [1] Diaz, A., Argüelles, C.A., Collin, G.H., Conrad, J.M., Shaevitz, M.H.: Where Are We With Light Sterile Neutrinos? *Phys. Rept.* **884**, 1–59 (2020) <https://doi.org/10.1016/j.physrep.2020.08.005> [arXiv:1906.00045](https://arxiv.org/abs/1906.00045) [hep-ex]
- [2] Hardin, J.M., Martinez-Soler, I., Diaz, A., Jin, M., Kamp, N.W., Argüelles, C.A., Conrad, J.M., Shaevitz, M.H.: New Clues about light sterile neutrinos: preference for models with damping effects in global fits. *JHEP* **09**, 058 (2023) [https://doi.org/10.1007/JHEP09\(2023\)058](https://doi.org/10.1007/JHEP09(2023)058) [arXiv:2211.02610](https://arxiv.org/abs/2211.02610) [hep-ph]
- [3] Dasgupta, B., Kopp, J.: Sterile Neutrinos. *Phys. Rept.* **928**, 1–63 (2021) <https://doi.org/10.1016/j.physrep.2021.06.002> [arXiv:2106.05913](https://arxiv.org/abs/2106.05913) [hep-ph]
- [4] Acero, M.A., *et al.*: White paper on light sterile neutrino searches and related phenomenology. *J. Phys. G* **51**(12), 120501 (2024) <https://doi.org/10.1088/1361-6471/ad307f> [arXiv:2203.07323](https://arxiv.org/abs/2203.07323) [hep-ex]
- [5] Giunti, C., Li, Y.F., Ternes, C.A., Tyagi, O., Xin, Z.: Gallium Anomaly: critical view from the global picture of ν_e and $\bar{\nu}_e$ disappearance. *JHEP* **10**, 164 (2022) [https://doi.org/10.1007/JHEP10\(2022\)164](https://doi.org/10.1007/JHEP10(2022)164) [arXiv:2209.00916](https://arxiv.org/abs/2209.00916) [hep-ph]
- [6] Wilks, S.S.: The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Stat.* **9**(1), 60–62 (1938)
- [7] Algeri, S., Aalbers, J., Morã, K.D., Conrad, J.: Searching for new phenomena with profile likelihood ratio tests. *Nat. Rev. Phys.* **2**(5), 245–252 (2020)
- [8] Hardin, J.M.: Wilks’s theorem, global fits, and neutrino oscillations. *Eur. J. Phys.* **45**(2), 025806 (2024) <https://doi.org/10.1088/1361-6404/ad261d> [arXiv:2211.06347](https://arxiv.org/abs/2211.06347) [hep-ph]
- [9] Feldman, G.J., Cousins, R.D.: Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D* **57**, 3873–3889 (1998) <https://doi.org/10.1103/PhysRevD.57.3873>
- [10] Villarreal, J., Woodward, J., Hardin, J., Conrad, J.: A frequentist simulation-based inference treatment of sterile neutrino global fits. *Machine Learning: Science and Technology* (2025)
- [11] Abbasi, R., *et al.*: Search for an ev-scale sterile neutrino using improved high-energy ν_μ event reconstruction in icecube. *Phys. Rev. Lett.* **133**, 201804 (2024) <https://doi.org/10.1103/PhysRevLett.133.201804>
- [12] Abbasi, R., *et al.*: Methods and stability tests associated with the sterile neutrino search using improved high-energy ν_μ event reconstruction in icecube. *Phys. Rev. D* **110**, 092009 (2024) <https://doi.org/10.1103/PhysRevD.110.092009>
- [13] Almazán, H., Bernard, L., Blanchet, A., Bonhomme, A., Buck, C., Amo Sanchez, P., El Atmani, I., Haser, J., Kandzia, F., Kox, S., Labit, L., Lamblin, J., Letourneau, A., Lhuillier, D., Licciardi, M., Lindner, M., Materna, T., Minotti, A., Pessard, H., Réal, J.-S., Roca, C., Rogly, R., Salagnac, T., Savu, V., Schoppmann, S., Sergeeva, V., Soldner, T., Stutz, A., Vialat, M.: Improved sterile neutrino constraints from the stereo experiment with 179 days of reactor-on data. *Phys. Rev. D* **102**, 052002 (2020) <https://doi.org/10.1103/PhysRevD.102.052002>
- [14] Andriamirado, M., Balantekin, A.B., Band, H.R., Bass, C.D., Bergeron, D.E., Berish, D., Bowden, N.S., Brodsky, J.P., Bryan, C.D., Classen, T., Conant, A.J., Deichert, G., Diwan, M.V., Dolinski, M.J., Erickson, A., Foust, B.T., Gaison, J.K., Galindo-Uribarri, A., Gilbert, C.E., Goddard, B.W., Hackett, B.T., Hans, S., Hansell, A.B., Heeger, K.M., Jaffe, D.E., Ji, X., Jones, D.C., Kyzlyova, O., Lane, C.E., Langford, T.J., LaRosa, J., Littlejohn, B.R., Lu, X., Maricic, J., Mendenhall, M.P., Meyer, A.M., Milincic, R., Mitchell, I., Mueller, P.E., Mumm, H.P.,

- Napolitano, J., Nave, C., Neilson, R., Nikkel, J.A., Norcini, D., Nour, S., Palomino, J.L., Pushin, D.A., Qian, X., Romero-Romero, E., Rosero, R., Surukuchi, P.T., Tyra, M.A., Varner, R.L., Venegas-Vargas, D., Weatherly, P.B., White, C., Wilhelmi, J., Woolverton, A., Yeh, M., Zhang, A., Zhang, C., Zhang, X.: Improved short-baseline neutrino oscillation search and energy spectrum measurement with the prospect experiment at hfir. *Phys. Rev. D* **103**, 032001 (2021) <https://doi.org/10.1103/PhysRevD.103.032001>
- [15] Skrobova, N.A.: New results of the danss experiment taking into account the absolute antineutrino counting depending on distance. *Bulletin of the Lebedev Physics Institute* **50**(12), 566–572 (2023) <https://doi.org/10.3103/S1068335623120151>
- [16] Atif, Z., Choi, J.H., Han, B.Y., Jang, C.H., Jang, H.I., Jang, J.S., Jeon, E.J., Jeon, S.H., Joo, K.K., Ju, K., Jung, D.E., Kim, B.R., Kim, H.J., Kim, H.S., Kim, J.G., Kim, J.H., Kim, J.Y., Kim, J.Y., Kim, S.B., Kim, S.Y., Kim, W., Kim, Y.D., Ko, Y.J., Kwon, E., Lee, D.H., Lee, H.G., Lee, J., Lee, J.Y., Lee, M.H., Lim, I.T., Moon, D.H., Oh, Y.M., Pac, M.Y., Park, H.K., Park, H.S., Park, J.S., Park, K.S., Park, R.G., Seo, H., Seo, J.W., Seo, K.M., Shin, C.D., Siyeon, K., Sun, G.M., Yang, B.S., Yeo, I.S., Yoo, J., Yoon, S.G., Yu, I.: Search for sterile neutrino oscillations using reno and neos data. *Phys. Rev. D* **105**, 111101 (2022) <https://doi.org/10.1103/PhysRevD.105.L111101>
- [17] Gavrin, V.N.: Contribution of gallium experiments to the understanding of solar physics and neutrino physics. *Physics of Atomic Nuclei* **76**(10), 1238–1243 (2013) <https://doi.org/10.1134/S106377881309007X>
- [18] Kostensalo, J., Suhonen, J., Giunti, C., Srivastava, P.C.: The gallium anomaly revisited. *Physics Letters B* **795**, 542–547 (2019) <https://doi.org/10.1016/j.physletb.2019.06.057>
- [19] Giunti, C., Laveder, M.: Statistical Significance of the Gallium Anomaly. *Phys. Rev. C* **83**, 065504 (2011) <https://doi.org/10.1103/PhysRevC.83.065504> [arXiv:1006.3244](https://arxiv.org/abs/1006.3244) [hep-ph]
- [20] Abdurashitov, J.N., *et al.*: Measurement of the solar neutrino capture rate with gallium metal. III: Results for the 2002–2007 data-taking period. *Phys. Rev. C* **80**, 015807 (2009) <https://doi.org/10.1103/PhysRevC.80.015807> [arXiv:0901.2200](https://arxiv.org/abs/0901.2200) [nucl-ex]
- [21] Abdurashitov, J.N., Gavrin, V.N., Girin, S.V., Gorbachev, V.V., Gurkina, P.P., Ibragimova, T.V., Kalikhov, A.V., Khairnasov, N.G., Knodel, T.V., Matveev, V.A., Mirmov, I.N., Shikhin, A.A., Veretenkin, E.P., Vermul, V.M., Yants, V.E., Zatsepin, G.T., Bowles, T.J., Elliott, S.R., Teasdale, W.A., Cleveland, B.T., Haxton, W.C., Wilkerson, J.F., Nico, J.S., Suzuki, A., Lande, K., Khomyakov, Y.S., Poplavsky, V.M., Popov, V.V., Mishin, O.V., Petrov, A.N., Vasiliev, B.A., Voronov, S.A., Karpenko, A.I., Maltsev, V.V., Oshkanov, N.N., Tuchkov, A.M., Barsanov, V.I., Janelidze, A.A., Korenkova, A.V., Kotelnikov, N.A., Markov, S.Y., Selin, V.V., Shakirov, Z.N., Zamyatina, A.A., Zlokazov, S.B.: Measurement of the response of a ga solar neutrino experiment to neutrinos from a ^{37}Ar source. *Phys. Rev. C* **73**, 045805 (2006) <https://doi.org/10.1103/PhysRevC.73.045805>
- [22] Barinov, V.V., Cleveland, B.T., Danshin, S.N., Ejiri, H., Elliott, S.R., Frekers, D., Gavrin, V.N., Gorbachev, V.V., Gorbunov, D.S., Haxton, W.C., Ibragimova, T.V., Kim, I., Kozlova, Y.P., Kravchuk, L.V., Kuzminov, V.V., Lubsandorzhiev, B.K., Malyshekin, Y.M., Massarczyk, R., Matveev, V.A., Mirmov, I.N., Nico, J.S., Petelin, A.L., Robertson, R.G.H., Sinclair, D., Shikhin, A.A., Tarasov, V.A., Trubnikov, G.V., Veretenkin, E.P., Wilkerson, J.F., Zvir, A.I.: Results from the baksan experiment on sterile transitions (best). *Phys. Rev. Lett.* **128**, 232501 (2022) <https://doi.org/10.1103/PhysRevLett.128.232501>
- [23] Barinov, V., Gorbunov, D.: Best impact on sterile neutrino hypothesis. *Phys. Rev. D* **105**, 051703 (2022) <https://doi.org/10.1103/PhysRevD.105.051703>

- [24] Conrad, J.M., Shaevitz, M.H.: Limits on Electron Neutrino Disappearance from the KARMEN and LSND ν_e - Carbon Cross Section Data. Phys. Rev. D **85**, 013017 (2012) <https://doi.org/10.1103/PhysRevD.85.013017> arXiv:1106.5552 [hep-ex]
- [25] Huber, P.: On the determination of anti-neutrino spectra from nuclear reactors. Phys. Rev. C **84**, 024617 (2011) <https://doi.org/10.1103/PhysRevC.85.029901> arXiv:1106.0687 [hep-ph]. [Erratum: Phys.Rev.C 85, 029901 (2012)]
- [26] Mention, G., Fechner, M., Lasserre, T., Mueller, T.A., Lhuillier, D., Cribier, M., Letourneau, A.: The Reactor Antineutrino Anomaly. Phys. Rev. D **83**, 073006 (2011) <https://doi.org/10.1103/PhysRevD.83.073006> arXiv:1101.2755 [hep-ex]
- [27] Giunti, C., Li, Y.F., Ternes, C.A., Xin, Z.: Reactor antineutrino anomaly in light of recent flux model refinements. Phys. Lett. B **829**, 137054 (2022) <https://doi.org/10.1016/j.physletb.2022.137054> arXiv:2110.06820 [hep-ph]
- [28] An, F.P., *et al.*: Comprehensive Measurement of the Reactor Antineutrino Spectrum and Flux at Daya Bay. Phys. Rev. Lett. **134**(20), 201802 (2025) <https://doi.org/10.1103/PhysRevLett.134.201802> arXiv:2501.00746 [nucl-ex]
- [29] Elliott, S.R., Gavrinn, V.N., Haxton, W.C., Ibragimova, T.V., Rule, E.J.: Gallium neutrino absorption cross section and its uncertainty. Phys. Rev. C **108**(3), 035502 (2023) <https://doi.org/10.1103/PhysRevC.108.035502> arXiv:2303.13623 [nucl-th]
- [30] Zeitnitz, B., Drexlin, G.: The karmen neutrino experiment. Interdiscip. Sci. Rev. **18**(3), 313–320 (1993)
- [31] Athanassopoulos, C., Auerbach, L.B., Bauer, D., Bolton, R.D., Burman, R.L., Cohen, I., Caldwell, D.O., Dieterle, B.D., Donahue, J.B., Eisner, A.M., Fazely, A., Federspiel, F.J., Garvey, G.T., Gray, M., Gunasingha, R.M., Highland, V., Imlay, R., Johnston, K., Kim, H.J., Louis, W.C., Lu, A., Margulies, J., Mills, G.B., McIlhenny, K., Metcalf, W., Reeder, R.A., Sandberg, V., Schillaci, M., Smith, D., Stancu, I., Strossman, W., Taylor, R., VanDalen, G.J., Vernon, W., Wang, Y.-X., White, D.H., Whitehouse, D., Works, D., Xiao, Y., Yellin, S.: The liquid scintillator neutrino detector and lampf neutrino source. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **388**(1), 149–172 (1997) [https://doi.org/10.1016/S0168-9002\(96\)01155-2](https://doi.org/10.1016/S0168-9002(96)01155-2)
- [32] Cranmer, K., Brehmer, J., Louppe, G.: The frontier of simulation-based inference. Proceedings of the National Academy of Sciences **117**(48), 30055–30062 (2020) <https://doi.org/10.1073/pnas.1912789117>
- [33] Boelts, J., Deistler, M., Gloeckler, M., Tejero-Cantero, Lueckmann, J.-M., Moss, G., Steinbach, P., Moreau, T., Muratore, F., Linhart, J., Durkan, C., Vetter, J., Miller, B.K., Herold, M., Ziaemehr, A., Pals, M., Gruner, T., Bischoff, S., Krouglova, N., Gao, R., Lappalainen, J.K., Mucsányi, B., Pei, F., Schulz, A., Stefanidi, Z., Rodrigues, P., Schröder, C., Zaid, F.A., Beck, J., Kapoor, J., Greenberg, D.S., Gonçalves, P.J., Macke, J.H.: sbi reloaded: a toolkit for simulation-based inference workflows. Journal of Open Source Software **10**(108), 7754 (2025) <https://doi.org/10.21105/joss.07754>
- [34] Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., Bao, B., Bell, P., Berard, D., Burovski, E., Chauhan, G., Chourdia, A., Constable, W., Desmaison, A., DeVito, Z., Ellison, E., Feng, W., Gong, J., Gschwind, M., Hirsh, B., Huang, S., Kalambarak, K., Kirsch, L., Lazos, M., Lezcano, M., Liang, Y., Liang, J., Lu, Y., Luk, C., Maher, B., Pan, Y., Puhersch, C., Reso, M., Saroufim, M., Siraichi, M.Y., Suk, H., Suo, M., Tillet, P., Wang, E., Wang, X., Wen, W., Zhang, S., Zhao, X., Zhou, K., Zou, R., Mathews, A., Chanan, G., Wu, P., Chintala, S.: PyTorch 2: Faster Machine Learning

- Through Dynamic Python Bytecode Transformation and Graph Compilation. In: 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24). ACM, ??? (2024). <https://doi.org/10.1145/3620665.3640366> . <https://docs.pytorch.org/assets/pytorch2-2.pdf>
- [35] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org (2015). <https://www.tensorflow.org/>
- [36] Greenberg, D., Nonnenmacher, M., Macke, J.: Automatic Posterior Transformation for Likelihood-Free Inference (2019). <https://doi.org/10.48550/arXiv.1905.07488>
- [37] Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow Matching for Generative Modeling (2023). <https://arxiv.org/abs/2210.02747>
- [38] Wildberger, J., Dax, M., Buchholz, S., Green, S., Macke, J.H., Schölkopf, B.: Flow matching for scalable simulation-based inference. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems, vol. 36, pp. 16837–16864. Curran Associates, Inc., ??? (2023). https://proceedings.neurips.cc/paper_files/paper/2023/file/3663ae53ec078860bb0b9c6606e092a0-Paper-Conference.pdf
- [39] Geffner, T., Papamakarios, G., Mnih, A.: Compositional score modeling for simulation-based inference. In: Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., Scarlett, J. (eds.) Proceedings of the 40th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 11098–11116. PMLR, ??? (2023). <https://proceedings.mlr.press/v202/geffner23a.html>
- [40] Sharrock, L., Simons, J., Liu, S., Beaumont, M.: Sequential Neural Score Estimation: Likelihood-Free Inference with Conditional Score Based Diffusion Models (2024). <https://arxiv.org/abs/2210.04872>
- [41] Cobb, A.D., Matejek, B., Elenius, D., Roy, A., Jha, S.: Direct amortized likelihood ratio estimation. In: Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence. AAAI'24/IAAI'24/EAAI'24. AAAI Press, ??? (2024). <https://doi.org/10.1609/aaai.v38i18.30018> . <https://doi.org/10.1609/aaai.v38i18.30018>
- [42] Alekseev, I., *et al.*: “Search for sterile neutrinos at the DANSS experiment”. Phys. Lett. **B787**, 56–63 (2018) <https://doi.org/10.1016/j.physletb.2018.10.038> arXiv:1804.04046 [hep-ex]
- [43] Cadeddu, M., Cargioli, N., Carotenuto, G., Dordei, F., Ferro, L., Giunti, C.: Reassessing the gallium anomaly using exact electron wave functions (2025). <https://arxiv.org/abs/2507.13103>
- [44] Cook, S.R., Gelman, A., Rubin, D.B.: Validation of software for bayesian models using posterior quantiles. Journal of Computational and Graphical Statistics **15**(3), 675–692 (2006). Accessed 2025-06-24
- [45] Talts, S., Betancourt, M., Simpson, D., Vehtari, A., Gelman, A.: Validating Bayesian Inference Algorithms with Simulation-Based Calibration (2018) arXiv:1804.06788 [stat.ME]

- [46] Giannopoulos, A.A., Milman, V.D.: Concentration property on probability spaces. *Adv. Math. (N. Y.)* **156**(1), 77–106 (2000)