PROCEEDINGS OF SCIENCE

# Bayesian neutrino oscillation analysis with 715 kton-yr of KM3NeT/ORCA

**Alfonso Lazo Pedrajas[a] on behalf of The KM3NeT Collaboration**

[a] *IFIC - Instituto de Física Corpuscular (CSIC - Universitat de València), c/Catedrático José Beltrán, 2, 46980 Paterna, Valencia, Spain*

*E-mail:* Alfonso.Lazo@ific.uv.es

KM3NeT/ORCA is a water-Cherenkov neutrino telescope currently under construction in the Mediterranean Sea, aimed at measuring atmospheric neutrino oscillations and determining the neutrino mass ordering. The detector consists of a three-dimensional array of detection units equipped with 18 digital optical modules, each of which houses 31 photomultiplier tubes. The Cherenkov light induced by charged particles produced in neutrino interactions in the 1 – 100 GeV range is used to reconstruct the parent neutrino's direction and energy, allowing for constraints on the oscillation parameters $\Delta m^2_{31}$ and $\theta_{23}$. This work presents the first Bayesian neutrino oscillation analysis performed with KM3NeT, using 715 kton-years of exposure from the ORCA detector. An adaptive Metropolis-Hastings algorithm is employed to build a Markov Chain Monte Carlo, which samples the posterior probability density in the multidimensional space spanned by oscillation and nuisance parameters. The credible region at 90% Confidence Level for the oscillation parameters, $\Delta m^2_{31} \in [-2.62, -1.80] \times 10^{-3}$ eV$^2$ assuming Inverted Ordering ( $\Delta m^2_{31} \in [1.92, 2.72] \times 10^{-3}$ eV$^2$ for Normal Ordering) and $\sin^2 \theta_{23} \in [0.38, 0.62]$, are inferred from the marginal posterior PDF.

39th International Cosmic Ray Conference (ICRC2025)
15–24 July 2025
Geneva, Switzerland





PoS(ICRC2025)1092

## 1. Atmospheric neutrino oscillations

High-energy cosmic ray interactions with atmospheric nuclei produce extensive air showers, primarily composed of mesons such as pions and kaons. These mesons decay as they propagate through the atmosphere, giving rise to a flux of muons and neutrinos emitted in all directions. The latter is dominated by muon and electron (anti)neutrinos, which can traverse the Earth before reaching the detector. As they propagate through the Earth matter, neutrinos undergo coherent forward scattering off fermions, an effect that modifies their oscillation behaviour. In the ultra-relativistic limit, neutrino propagation in matter can be described by introducing a matter potential into the Hamiltonian, altering the vacuum oscillation pattern as follows:

$$\mathcal{H}\text{eff} = \frac{1}{2\text{E}}\mathcal{U}_{\text{PMNS}}\,\text{diag}(0, \Delta m^2_{21}, \Delta m^2_{31})\,\mathcal{U}^{\dagger}_{\text{PMNS}} + \text{diag}(\sqrt{2}G_F n_{\text{e}}(x), 0, 0)\,, \tag{1}$$

where $\mathcal{U}_{\text{PMNS}}$ is the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix that encodes neutrino flavor mixing, including the mixing angle $\theta_{23}$ measured in this analysis. The terms $\Delta m^2_{ij}$ represent the squared mass differences between neutrino mass eigenstates $i$ and $j$, and $n_{\text{e}}(x)$ denotes the electron number density along the neutrino path. Flavor transition probabilities, $P(\nu_\alpha \to \nu_\beta)$ with $\alpha, \beta = e, \mu, \tau$, are obtained by numerically solving the time evolution equation using the effective Hamiltonian in eq. (1). By analysing neutrinos from a wide range of arrival directions in the multi-GeV range, the oscillation parameters $\Delta m^2_{31}$ and $\theta_{23}$ are extracted by comparing Monte Carlo (MC) templates to the observed event distributions in the detector.

## 2. The KM3NeT/ORCA detector

The KM3NeT/ORCA neutrino telescope is currently being built and will consist of 115 vertical Detection Units (DUs) in its final configuration. Each DU holds 18 Digital Optical Modules (DOMs), spaced 9 meters apart vertically, with 20 meters of horizontal separation between the DUs. Each DOM contains 31 3-inch photomultiplier tubes (PMTs), designed to detect the Cherenkov light produced by secondary charged particles resulting from neutrino interactions. The energy and incoming direction of atmospheric neutrinos are inferred by reconstructing the spatial and temporal pattern of detected light. With this setup, ORCA can probe neutrinos travelling along baselines from a few kilometres up to approximately $10^4$ km, covering various layers of the Earth's interior. The detector is sensitive to neutrinos with energies roughly in the range of 1 to 100 GeV. The measurement of $\Delta m^2_{31}$ and $\theta_{23}$ reported here is based on the latest available ORCA data, corresponding to an exposure of 715 kt-y, which includes data from configurations operating with 6, 10, and 11 DUs.

## 3. Markov Chain Monte Carlo

This work employs ORCA data in order to perform a Bayesian neutrino oscillation analysis for the first time in KM3NeT. The goal of the analysis is to perform a random sampling of the posterior probability density function (PDF) in the multidimensional space spanned by the oscillation parameters, $\Delta m^2_{31}$ and $\sin^2\theta_{23}$, here referred to as Parameters of Interest (PoI), and the 13 nuisance parameters modelling different systematic uncertainties. A Markov Chain Monte Carlo

(MCMC) method is used to numerically sample the posterior PDF by means of the Metropolis-Hastings (MH) algorithm [1, 2], which proceeds by iteratively drawing proposals for PoIs and nuisance parameters that identify the model, denoted as $\vec{\omega}^t$ and $\vec{\eta}^t$, respectively. Each proposal then defines a *step* $t$ representing a point in the multidimensional parameter space, denoted as $Y^t \equiv \{\vec{\omega}^t, \vec{\eta}^t\}$, which may be incorporated into the Markov Chain or not. If it is, we denote the point as $X^t$. An acceptance criterion for the subsequent proposals $Y^{t+1}$, based on the last accepted step $X^t$ and the posterior probability of both, is chosen in order to construct a Markov Chain of steps $\{X^0, X^1, ..., X^t, \dots\}$.

In the MH algorithm, the subsequent steps $Y^{t+1}$ are randomly drawn from a proposal PDF, $Q(Y^{t+1}, X^t)$, which depends on the last accepted step of the chain $X^t$, and is symmetric with respect to the interchange of $Y^{t+1}$ and $X^t$. The point represented by $Y^{t+1}$ is then accepted into the Markov Chain with probability $\alpha(Y^{t+1}, X^t)$, computed as

$$\begin{aligned}\alpha(Y^{t+1}, X^t) = \min\left(1, \frac{P(Y^{t+1}|D)Q(Y^{t+1}, X^t)}{P(X^t|D)Q(X^t, Y^{t+1})}\right) = \\ \min\left(1, \frac{\mathcal{L}(D|\vec{\omega}^{t+1}, \vec{\eta}^{t+1})P(\vec{\eta}^{t+1})}{\mathcal{L}(D|\vec{\omega}^t, \vec{\eta}^t)P(\vec{\eta}^t)}\right),\end{aligned} \tag{2}$$

where $P(X^t|D)$ ($P(Y^{t+1}|D)$) represents the posterior probability of $X^t$ ($Y^{t+1}$) given the observed data $D$. Bayes theorem has been invoked in the second line when replacing $P(\vec{\omega}, \vec{\eta}|D) = \mathcal{L}(D|\vec{\omega}, \vec{\eta})P(\vec{\eta})/P(D)$, where $\mathcal{L}(D|\vec{\omega}, \vec{\eta})$ is the likelihood function, $P(\vec{\eta})$ is the prior probability of the constrained nuisance parameters, and the Bayesian evidence $P(D)$ represents the posterior normalization, which simplifies in the quotient at the second line of eq. (2). The symmetry property of $Q$ has been assumed. If $Y^{t+1}$ is accepted into the Markov Chain, the algorithm identifies $X^{t+1} = Y^{t+1}$ and draws the next proposal with probability $Q(Y^{t+2}, X^{t+1})$. Otherwise, the step is rejected and new phase space points are sampled from $Q(Y'^{,\,t+1}, X^t)$ until a proposal is accepted into the chain. The sequence $\{X^0, X^1, ..., X^t\}$ grows with further accepted steps, becoming independent of the first proposal $X^0$ for large enough $t$, in what is said to be the stationary state. Then, the set $\{X^B, X^{B+1}, ..., X^t\}$ is ensured to be a faithful representation of the posterior, provided that the first $B$ steps of the original chain (the *burn-in* period) are discarded, where the chain is still highly correlated to its initial state.

This work employs a multivariate Gaussian centred at the last accepted step as a proposal PDF, $Q \sim \mathcal{G}(X^t, \hat{\Sigma}^{(t)})$, where $\hat{\Sigma}^{(t)}$ represents the Empirical Covariance Matrix (ECM) of model parameters computed using the Markov Chain ensemble up to step $t$. An adaptive MCMC approach is followed [3], in which the ECM is updated at runtime every 50 steps. Furthermore, $\hat{\Sigma}^{(t)}$ is adjusted by means of individual scaling factors which modify the sampling width along each parameter dimension. The scaling factors are tuned in order to increase the acceptance rate of proposed steps while keeping a low autocorrelation time for all sampled parameters.

All nuisance parameters and $\Delta m^2_{31}$ are sampled jointly from $\mathcal{G}(X^t, \hat{\Sigma}^{(t)})$. Due to the true Neutrino Mass Ordering (NMO) being unknown, the sign of $\Delta m^2_{31}$ drawn from the proposal PDF is inverted with a 50% chance. On the other hand, the value of $\theta_{23}$ is drawn independently of all other model parameters from a uniform PDF, $\sin^2\theta_{23} \sim \mathcal{U}[0.25, 0.75]$, which was shown to better reproduce the multimodal posterior density expected from the $\theta_{23}$-octant degeneracy. We adopt

in this work the effective likelihood formalism proposed in ref. [4] which is suited for Bayesian inference.

## 4. Results

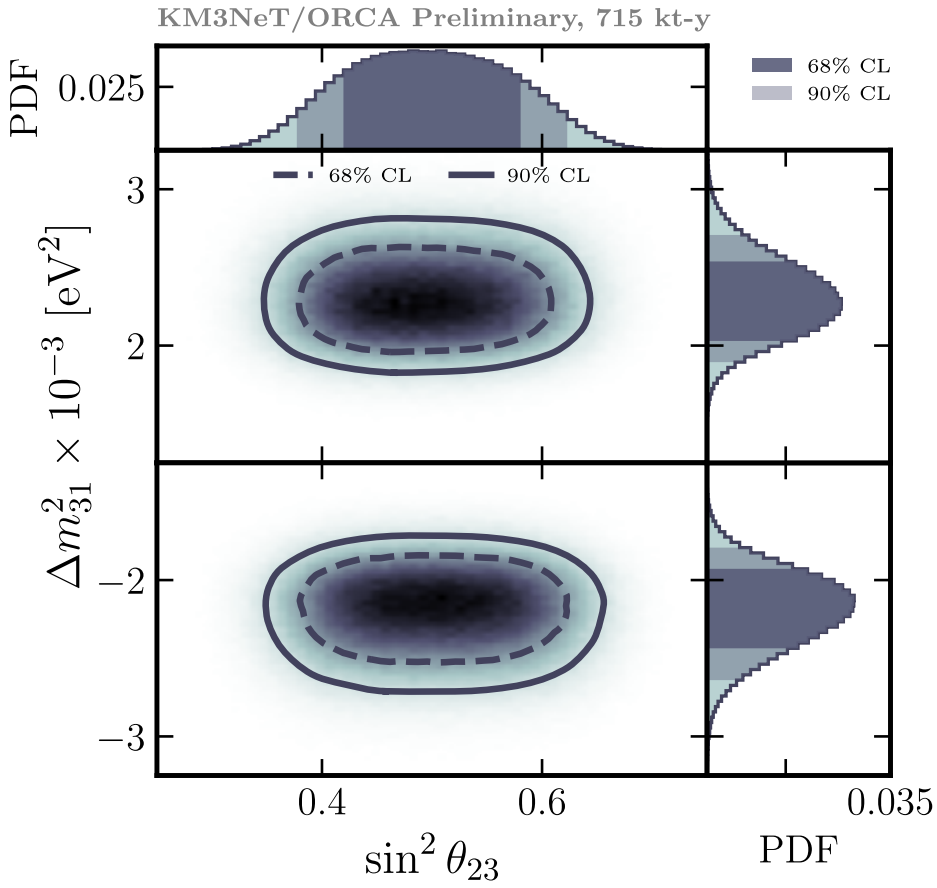


**Figure 1:** Two-dimensional marginal posterior in the space of $\Delta m^2_{31}$ and $\sin^2\theta_{23}$, together with the one-dimensional marginalisation for $\sin^2\theta_{23}$ and every mass ordering possibility for $\Delta m^2_{31}$. The dashed and solid lines represent the limits of the joint credible region for the oscillation parameters at 68% and 90% CL, respectively, while the shaded areas at the top and side projections show the marginal credible regions.

Fifty Markov Chains with independent random starting points were run in parallel and combined after burn-in into a single MCMC set, featuring a total of 1850000 steps. The distribution of the marginal posterior PDF in the joint space of $\Delta m^2_{31}$ and $\sin^2\theta_{23}$, shown in fig. 1, is obtained from the density of MCMC steps per bin in the space of oscillation parameters, represented by the colour scale. The top and side projections show the marginal posterior PDF in the space of each oscillation parameter alone. The credible regions at different confidence levels (CL) are represented by either contour lines in the two-dimensional panels, or shaded areas in the one-dimensional projections; they are computed by integrating the PDF from its maximum in decreasing order of probability

density until the desired CL is reached. The thirteen nuisance parameters listed in ref. [5] have been marginalised over.

The estimated values for the oscillation parameters are extracted from the projected PDFs as the points of highest posterior density (HPD), resulting in

$$\Delta m^2_{31} = -2.13^{+0.21}_{-0.29} \times 10^{-3}\ \text{eV}^2, \quad \sin^2\theta_{23} = 0.48^{+0.09}_{-0.07}. \tag{3}$$

The strongest posterior evidence for either mass ordering favours IO with a Bayes factor of $\mathcal{B}_{\text{IO}} = P(\text{IO}|D)/P(\text{NO}|D) = 1.18$ assuming the upper $\theta_{23}$-octant, corresponding to no significant NMO preference. Fig. 2 presents the joint credible region at 90% CL, shown by the solid black line, in the global context of neutrino oscillation measurements from multiple experiments. The dashed black line represents the confidence region from a frequentist analysis using the same 715 kt-y dataset as in this work [6], which is compatible with the Bayesian result.

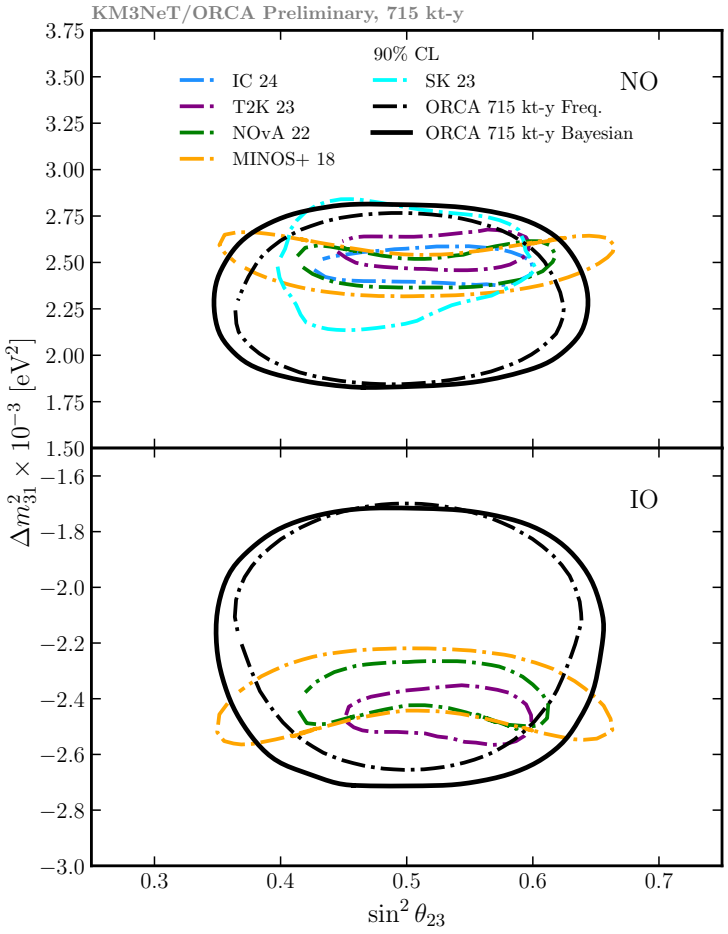


**Figure 2:** Allowed region at 90% CL for the oscillation parameters $\sin^2\theta_{23}$ and $\Delta m^2_{31}$ obtained from the Bayesian analysis of this work, from the frequentist analysis covered in ref.[6] using the 715 kt-y dataset of ORCA, and reported by Super-Kamiokande [7], IceCube [8], T2K [9], NOvA [10] and MINOS+ [11]. Some experiments report their allowed regions exclusively for NO. The results from T2K, NOvA and this work are Bayesian; the remaining ones are frequentist confidence regions.

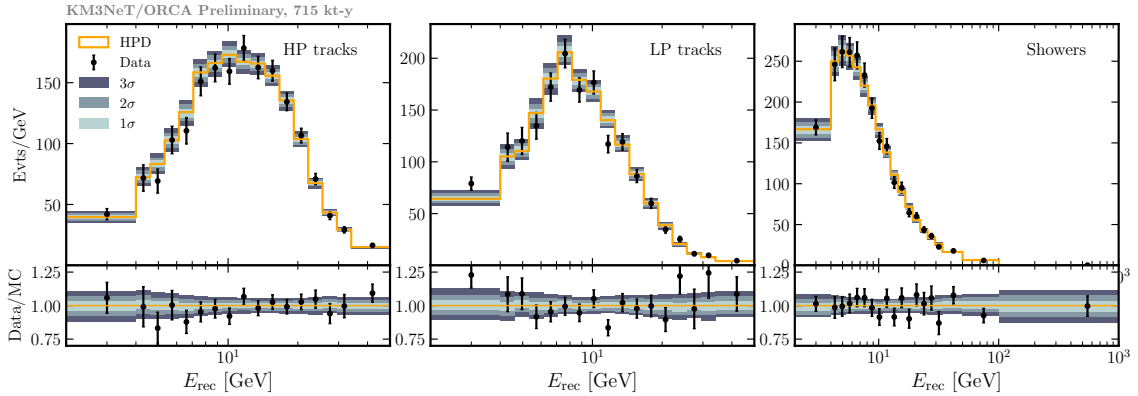


**Figure 3:** Reconstructed energy distribution for the high-purity (HP) tracks (left), low-purity (LP) tracks (centre) and showers (right) classes. The orange solid line represents the HPD prediction for the number of events in every bin, and the coloured boxes show its 1, 2 and 3$\sigma$ credible region. The bottom plots show the data/MC ratio and the posterior relative systematic uncertainty of the MC at the 1, 2 and 3$\sigma$ level.

Fig. 3 shows the reconstructed energy distribution for data and the posterior prediction inferred from it, split into the three event categories defined in the analysis. The posterior predictive distribution is built by reweighting the MC expectation at 10000 values of oscillation and nuisance parameters randomly sampled from the posterior PDF. The coloured bands represent the $1\sigma$, $2\sigma$ and $3\sigma$ variability arising from the posterior systematic uncertainty of the bin-by-bin model prediction. The observed data is to be compared with the orange lines, representing the number of events most frequently expected from the posterior samples (referred to as the HPD of the number of events in that bin). The ratio plots in the bottom row reflect an excellent agreement in reconstructed energy between the data and the expectation from the inferred model.

The compatibility between data and posterior PDF can be quantified by the posterior predictive p-value, serving as an analogue to the frequentist goodness-of-fit p-value. We follow the approach outlined in refs. [9, 10], where the likelihood values from data and the posterior predictive distribution, $\mathcal{L}(\vec{\Omega}_i|D)$ with $\vec{\Omega}_i = \{\vec{\omega}_i, \vec{\eta}_i\}$ for $i = 1, \ldots, 10^4$, are compared to the likelihood computed using pseudodata generated by the Poisson fluctuation of the MC prediction at $\vec{\Omega}_i$, denoted as $\mathcal{L}(\vec{\Omega}_i|D^{\text{fluct}}_{\vec{\Omega}_i})$. The posterior predictive p-value is identified with the proportion of draws satisfying $-2\log\mathcal{L}(\vec{\Omega}_i|D) < -2\log\mathcal{L}(\vec{\Omega}_i|D^{\text{fluct}}_{\vec{\Omega}_i})$, associated with the cases where random fluctuations are less compatible with the posterior predictive distribution than the observed data. The condition happens 43% of the times, shown by the hatched orange region in fig. 4, proving a high compatibility of the data with the posterior model.

## 5. Conclusions

This work reports the results from the first Bayesian neutrino oscillation analysis conducted with ORCA. An adaptive Metropolis-Hastings algorithm was implemented in order to sample the posterior PDF in the multidimensional space spanned by oscillation and nuisance parameters. Fifty Markov Chains were run for that purpose, which were previously tuned to the desired levels of sampling efficiency and autocorrelation time of the parameters. The credible regions inferred from

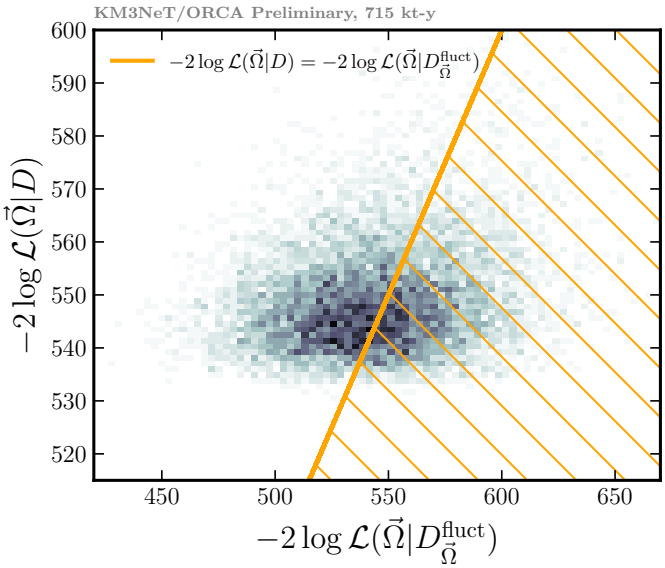


**Figure 4:** Negative log-likelihood of 10000 parameter picks from the posterior PDF, $\vec{\Omega}_i$, evaluated with the observed data $D$ and the model generated at $\vec{\Omega}_i$, $-2\log\mathcal{L}(\vec{\Omega}_i|D)$ , against the log-likelihood computed using a random fluctuation of the model as pseudodata, $-2\log\mathcal{L}(\vec{\Omega}_i|D^{\text{fluct}}_{\vec{\Omega}_i})$. The orange line corresponds to equal values of both likelihoods. The fraction of points $\vec{\Omega}_i$ under the line (hatched region) is the posterior predictive p-value of 43%.

the 715 kt-y dataset of ORCA are $\Delta m^2_{31} \in [-2.62, -1.80] \times 10^{-3}$ eV$^2$ ($[1.92, 2.72] \times 10^{-3}$ eV$^2$ assuming NO) and $\sin^2\theta_{23} \in [0.38, 0.62]$ at 90% CL. The Inverted Ordering was favoured by the posterior model with a Bayes factor of 1.18 assuming upper $\theta_{23}$-octant. The posterior predictive p-value is found to be 43%, showing an excellent agreement between data and MC expectation.

This work supports the frequentist results previously obtained from the same ORCA dataset. Additionally, the use of MCMC methods opens the possibility to explore higher-dimensional oscillation models, such as Beyond the Standard Model scenarios with additional interactions, while naturally incorporating parameter correlations across the entire phase space, and avoiding the computational burden often associated with gradient-descent minimisations in frequentist approaches.

## Acknowledgements

The authors acknowledge the support from MCIN for PID2021-124591NB-C41, -C42, -C43 and PDC2023-145913-I00 funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe", for ASFAE/2022/014 and ASFAE/2022 /023 with funding from the EU NextGenerationEU (PRTR-C17.I01) and Generalitat Valenciana, for Grant AST226.2 with funding from Consejería de Universidad, Investigación e Innovación and Gobierno de España and European Union - NextGenerationEU, for CNS2023-144099, Generalitat Valenciana for CIDEGENT/2020/049, CIDEGENT/2021/23, CIDEIG/2023/20, ESGENT2024/24, CIPROM/2023/51, INNVA1/2024/110 (IVACE+i) and ACIF/2021/233, Spain.

## References


[1] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller, *Equation of State Calculations by Fast Computing Machines*, *The Journal of Chemical Physics* **21** (1953) 1087.

[2] W.K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications*, *Biometrika* **57** (1970) 97.

[3] Jonas Wallin and David Bolin, *Efficient adaptive mcmc through precision estimation*, 1505.03908.

[4] C.A. Argüelles, A. Schneider and T. Yuan, *A binned likelihood for stochastic models*, *JHEP* **06** (2019) 030 [1901.04645].

[5] KM3NeT collaboration, *Measurement of neutrino oscillation parameters with the first six detection units of KM3NeT/ORCA*, 2408.07015.

[6] KM3NeT collaboration, *Measurement of atmospheric neutrino oscillations with KM3NeT/ORCA*, *PoS* **ICHEP2024** (2025) 134.

[7] Super-Kamiokande collaboration, *Atmospheric neutrino oscillation analysis with neutron tagging and an expanded fiducial volume in Super-Kamiokande I–V*, *Phys. Rev. D* **109** (2024) 072014 [2311.05105].

[8] (IceCube Collaboration)‖, IceCube collaboration, *Measurement of Atmospheric Neutrino Oscillation Parameters Using Convolutional Neural Networks with 9.3 Years of Data in IceCube DeepCore*, *Phys. Rev. Lett.* **134** (2025) 091801 [2405.02163].

[9] T2K collaboration, *Measurements of neutrino oscillation parameters from the T2K experiment using* $3.6 \times 10^{21}$ *protons on target*, *Eur. Phys. J. C* **83** (2023) 782 [2303.03222].

[10] NOvA collaboration, *Improved measurement of neutrino oscillation parameters by the NOvA experiment*, *Phys. Rev. D* **106** (2022) 032004 [2108.08219].

[11] MINOS+ collaboration, *Precision Constraints for Three-Flavor Neutrino Oscillations from the Full MINOS+ and MINOS Dataset*, *Phys. Rev. Lett.* **125** (2020) 131802 [2006.15208].


PoS(ICRC2025)1092