PROCEEDINGS OF SCIENCE

Beyond the null hypothesis: Physics-driven searches for astrophysical neutrino sources

Francesca Capel,[a,*] Julian Kuhlmann,[a] Christian Haack,[b] Martin Ha Minh,[c] Hans Niederhausen[d] and Lisa Schumacher[b]

[a] *Max Planck Institute for Physics,*
Föhringer Ring 6, 80805 Munich, Germany
[b] *Erlangen Centre for Astroparticle Physics*
Nikolaus-Fiebiger-Str. 2, 91058 Erlangen, Germany
[c] *Technical University of Munich*
James-Franck-Str. 1, 85748 Garching, Germany
[d] *Michigan State University*
288 Farm Lane, 48824 East Lansing MI, USA

E-mail: capel@mpp.mpg.de

While there have been multiple reports of potential links between astrophysical sources and high-energy neutrinos, definitively identifying these sources remains challenging. No single result has surpassed the significance threshold for discovery, and the theoretical interpretation of these observations has revealed new questions for the bigger picture of neutrino astrophysics. In response to these difficulties, we propose a novel approach that combines state-of-the-art methods with the wealth of information from multi-messenger observations and theoretical models. Our physics-first strategy utilizes advanced techniques such as Hamiltonian Monte Carlo and hierarchical modelling to navigate the high-dimensional space of parameters and uncertainties in this complex setting. Implementing these methods within a Bayesian framework offers a unique perspective on source discovery while enabling us to avoid pitfalls associated with trial factors. Furthermore, considering source populations provides additional constraints from observations and uncovers further implications from a broader perspective, making the most of the existing high-energy neutrino data. We demonstrate the impact of these methods through application to simulated data, considering relevant multi-messenger scenarios for sources and their populations. Our approach complements existing techniques while offering increased sensitivity and interpretability in the context of investigating specific physical models, potentially accelerating decisive discoveries.

38th International Cosmic Ray Conference (ICRC2023)
26 July - 3 August, 2023
Nagoya, Japan

*Speaker

PoS(ICRC2023)1576

1. Introduction

The detection of point sources in astrophysical neutrino data is challenged by low event rates at the highest energies and large backgrounds from both atmospheric and diffuse astrophysical contributions. These difficulties mean that despite the confident detection of an astrophysical component to the total neutrino spectrum, point sources are still just beginning to emerge from the current data. With planned upgrades to existing facilities, as well as the construction of more large-scale neutrino observatories, we expect to move from source discovery to source characterisation in the not-so-distant future (see, e.g. [1] for a recent review).

Current approaches to neutrino point source searches are largely based on null hypothesis significance testing, as described in [2, 3]. The likelihood is formed as a mixture model over source and background components, where the source and background probability density functions (PDFs) are a function of both the reconstructed energies and arrival directions of events produced by incident neutrinos in an instrumented detector volume. A test statistic is then defined as the log-likelihood ratio of the null hypothesis (only background) to the alternate hypothesis (background + source). The observed value of the test statistic for the best-fit alternate hypothesis is calibrated via repeated simulations assuming the null hypothesis to give a p-value, quantifying the significance of the test. Typically, two free parameters are considered in the fit; the expected number of events from the proposed point source, $\bar{n}$ and the spectral index, γ, describing the assumed power-law spectrum of the source. The best-fit values are generally found using quasi-Newton methods implemented in, e.g. `iminuit`[1] or `scipy`[2].

These approaches are limited in the number of free parameters they can handle, in turn limiting the inclusion of more complex models and uncertainties into the statistical analysis. Additionally, the interpretation of p-values is always in relation to the null hypothesis, leading to an indirect approach to asking the question of interest. We are forced to work with "how unlikely is this observation under the assumptions of the background model?" rather than the more direct "what is the probability that these neutrinos come from this source?" [4, 5]. When testing multiple hypotheses, the resulting p-values must be corrected for the incurred "trial factors" to avoid spurious claims of significance. In practice, this means that analysis choices are often driven or limited by the corresponding trial factors, and there is no clear way to keep track of trial factors across independent analyses of the same data. All of the above aspects make it challenging to interpret significant results within the broader context of physical source models and their cosmological populations [6].

We propose a complementary Bayesian approach for searches of astrophysical neutrino point sources that aims to address these challenges, allowing us to exploit the available data further. The method is outlined in Section 2 and then applied to a simple single source example in Section 3. We then summarise the implication of our results and conclude in Section 4.

2. Statistical framework

Our method is based on a hierarchical mixture model framework similar to that discussed in [7] (see also [8] and references therein), but adapted to allow for a joint fit of high-energy

[1] https://iminuit.readthedocs.io/en/stable/
[2] https://scipy.org

PoS(ICRC2023)1576

neutrino reconstructed energies, $\hat{E}$, and directions, $\hat{\omega}$. The likelihood has the form of an inhomogeneous Poisson point process where the rate is a mixture model over the possible point source and background contributions. We have

$$\mathcal{L}(\hat{E}, \hat{\omega}|\Theta) = e^{-\bar{n}} \prod_{i=1}^{n} \sum_{k=1}^{N} \Phi_k P(\hat{E}_i, \hat{\omega}_i|\Theta_k), \tag{1}$$

where Θ is the vector of all source component parameters, and Θ_k are those relevant for individual source components. For point sources, these are the luminosity and power-law spectral index $\Theta_k = \{L_k, \gamma_k\}$. Point sources are also described by fixed directions and distances, ω_k and D_k, respectively. For the diffuse astrophysical component, we have the number flux and power-law spectral index $\Theta_k = \{\Phi_\mathrm{d}, \gamma_\mathrm{d}\}$. Finally, for the atmospheric component, the only free parameter is the number flux $\Theta_k = \{\Phi_\mathrm{a}\}$, as the spectral shape and directional distribution is fixed using `MCEq`[3] [9–11]. The number of detected neutrinos in the sample is given by n, and the total number of proposed source components is denoted N. The mixture is weighted by the source component flux integrated over energy and solid angle to give Φ_k. $P(\hat{E}_i, \hat{\omega}_i|\Theta_k)$ is implemented in a hierarchical way, as the product of conditional statements which connect the observables to the high-level parameters according to their physical relationships. Therefore, each neutrino event denoted by i has latent parameters describing the true energies (both on arrival and at the source), and source labels, λ_i.

In this work, we use the public data release relevant for track-like events observed by the IceCube neutrino observatory for our detector model, as detailed in [12]. This release contains information on the effective area, angular resolution and energy resolution, which are needed to fully describe the reconstructed events and demonstrate the application of our statistical model. The more recent data reported in [13] has also been implemented in our framework and will be applied in future work.

We use the above likelihood in a Bayesian framework and so also specify priors for the highest level parameters, Θ. The choice of these priors depends on the analysis is discussed further in Section 3. We fit the model using Hamiltonian Monte Carlo as implemented in `Stan` [14], which allows for large numbers of free parameters, easing the computational limitations described in Section 1. In this way, our source model can be expanded to study more complex spectral or population models. We ensure convergence of the resulting parameter Markov chains to the target distribution by observing the requirements on the effective sample size, Gelman-Rubin statistic, divergent transitions and other default diagnostics implemented in the `cmdstanpy` Python interface to `Stan` [15]. In taking a Bayesian approach, we also address some of the challenges associated with hypothesis testing. The possibility of drawing false conclusions is mitigated by marginalisation over the priors and model structure, precluding the need for trial factor corrections and avoiding a loss of sensitivity to true differences [16, 17].

3. Application to neutrino alert follow-up

To demonstrate our approach, we consider a simple astrophysical follow-up example motivated by the well-studied association between the blazar TXS 0506+056 and the IceCube high-energy

[3]https://mceq.readthedocs.io/en/latest/

neutrino alert event IC-170922A [18]. We consider TXS 0506+056 as a single point source, described by a luminosity L and a spectral index γ, as well as diffuse astrophysical and atmospheric background components. We assume that the background components are well-understood in relation to the possible source contribution. As such, they are constrained by narrow prior assumptions reflecting the uncertainties of these measurements.

Description	Variables	Implementation
TXS 0506+056	(RA, δ), z	(77.36°, +5.69°), 0.3365
IC-170922A	$\hat{E}$, (RA, δ), σ_ω	23.7 TeV, (77.42°, +5.72°), 0.7°
Energy range	$E_{\rm min}$, $E_{\rm max}$	10 TeV, 100 PeV
Atmospheric flux prior	$\Phi_{\rm a}, \sigma_\Phi^{\rm a}$	9.54×10^{-10} cm^{-2}s^{-1}, 0.1
Diffuse flux prior	$\Phi_{\rm d}, \sigma_\Phi^{\rm d}$	6.25×10^{-11} cm^{-2}s^{-1}, 0.1
Diffuse spectrum prior	$\gamma_{\rm d}, \sigma_\gamma$	2.6, 0.2

Table 1: Key parameters and assumptions considered in this work. σ_ω is the reconstructed angular uncertainty, assuming a von Mises-Fisher distribution description of the angular resolution. $E_{\rm min}$ and $E_{\rm max}$ are defined at Earth and give the energy range used to define the fluxes of all model components. The background flux priors are given by LogNormal(Φ, σ_Φ) and the spectral index priors are Normal(γ, σ_γ).

We first consider the application of our approach to the astrophysical follow-up of a single neutrino event. We start with a dataset containing only simulated background for an observation period of $T_{\rm obs} = 0.5$ yr and a high cut on the reconstructed event energy of $\hat{E}_{\rm th} = 20$ TeV. We then add a single alert event matching the description of the IC-170922A. As we expect the source spectral index to be unconstrained by the single event provided, we consider a narrow prior $\gamma \sim \text{Normal}(2.0, 0.2)$, but allow for a wide prior on L, with the upper limit motivated by previous point source searches at this location on the sky such that $L \sim \text{LogNormal}(10^{38}\ \text{erg s}^{-1}, 6)$ [19]. The results of our fit are shown in Fig. 1. As expected, the posterior for γ is equivalent to the prior due to the lack of information available in the data. However, the single event added leads to a small contribution to higher values of L in the posterior, consistent with the production of a single event under the assumptions of our detector model. We can directly evaluate the probability that the single neutrino event and the proposed source are connected by deriving the marginal posterior of the event source label, $p_{\rm assoc} = P(\lambda_i = \text{src}|\hat{E}_i, \hat{\omega}_i)$, as described in Appendix B of [7]. For the case shown in Fig. 1, we find $p_{\rm assoc} = 0.56$.

The information contained in the IC-170922A-like event alone is insufficient to derive a conclusive association probability, and the results for $p_{\rm assoc}$ are dependent on the prior assumptions for the source properties. We examine the prior dependence by taking the same dataset and assuming $L \sim \text{LogNormal}(\mu_L, 0.5)$ and $\gamma \sim \text{Normal}(\mu_\gamma, 0.2)$, scanning over μ_L and μ_γ as shown in Fig. 2. We see that $p_{\rm assoc}$ transitions from 0 to 1 for μ_L in the range 10^{44} to 10^{46} erg s^{-1}, consistent with what we see in Fig. 1, and is also weakly dependent on μ_γ. In this way, if there are constraints on the source properties from complementary observations, these can be used to design informative priors or explore the implications of different proposed source models.

While following up on individual neutrino alerts can be key to identifying transient sources, it is difficult to draw strong conclusions from a single event without further constraints on, e.g.

PoS(ICRC2023)1576

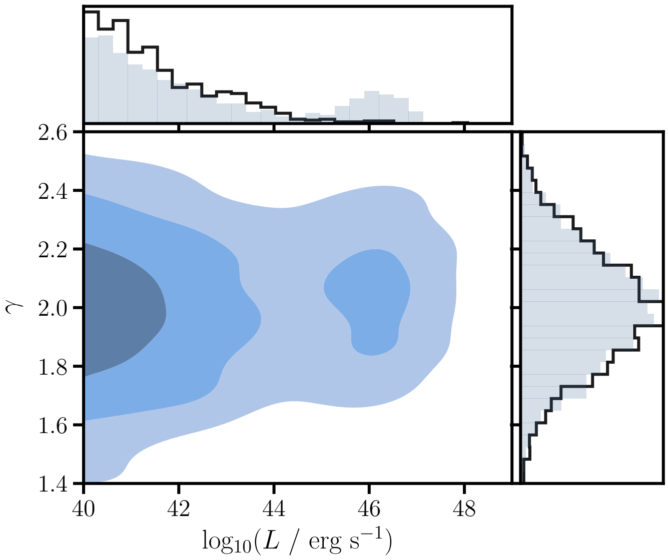

Figure 1: The joint posterior distribution for L and γ is shown with contours corresponding to the 30, 70 and 95% regions of highest posterior density. The upper and right panels show the marginal distributions as shaded histograms for L and γ, respectively, with the solid line showing the relevant prior distribution for comparison.

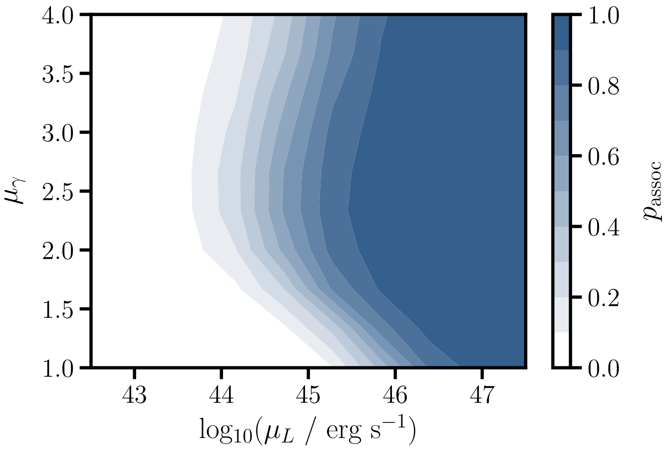

Figure 2: The event–source association probability found for fits of a dataset containing a single IC-170922A-like event with priors centred on L_μ and γ_μ, as described in the text. The colour scale shows the $p_{\rm assoc}$ in steps of 0.1.

the neutrino production time. With this in mind, we also explore the impact of including more information from the neutrino data into the analysis by decreasing $\hat{E}_{\rm th}$ and increasing $T_{\rm obs}$, with the aim of increasing the total number of source events in our sample. To do this, we set up simulations based on the above assumptions for 1 TeV < $\hat{E}_{\rm th}$ < 20 TeV and 0.5 yr < $T_{\rm obs}$ < 2 yr in a rectangular $20^\circ \times 20^\circ$ region of interest centred on the location of TXS 0506+056. We set $L_{\rm true} = 5\times10^{46}$ erg s^{-1} and $\gamma_{\rm true} = 2$, which corresponds to roughly $\bar{n}_{\rm src} \sim 1$ expected source events for our original assumptions on $\hat{E}_{\rm th}$ and $T_{\rm obs}$. Assuming wide priors such that $L \sim \mathrm{LogNormal}(L_{\rm true}, 5)$ and

$\gamma \sim \text{Normal}(\gamma_{\text{true}}, 1)$, we fit these simulations and show the marginal posterior distributions for L and γ in Fig. 3. We see that even including a small number of additional source events in the sample can have a large impact on the accurate reconstruction of the source properties and robustness to prior assumptions.

In [6], we demonstrate via a simulation-based approach that considering the blazar population as a whole helps to place the TXS 0506+056–IC170922A association in context and inform more physical choices for the likelihoods used to study such associations. Our framework can also be extended to include further information from other sources and their related population in a consistent way, allowing for a more direct interpretation of the results under the model assumptions and possible further constraints on source parameters from the total diffuse flux. We plan to investigate these possibilities further in upcoming work.

In Fig. 3, we also introduce the derived parameter $f = \bar{n}_{\text{src}}/\bar{n}_{\text{tot}}$, which represents the expected fraction of detected events which are associated with the proposed point source. f summarises the p_{assoc} values of individual events for the entire sample considered and provides a complementary way to define source association or detection. As we decrease $\hat{E}_{\text{th}}$ and increase T_{obs}, f decreases as we also collect more background events in the sample relative to source events. However, the constraints on f become much stronger, allowing us to state $f > 0$ more confidently. The derived parameters p_{assoc} and f can be used to define sensitivity and discovery within this parameter estimation framework. For example, as proposed in [20], by finding the α_f% highest posterior density region of f and defining a threshold value of p_{assoc} with some level, α_p, we can make decisions based on their compatibility with 0 or 1. The values of α_f and α_p can also be calibrated via simulations to provide a certain level of coverage, as desired.

4. Conclusions

We propose a hierarchical Bayesian approach to searching for point sources of astrophysical neutrinos. Motivated by the emergence of possible sources but the difficulty in their interpretation, we aim to include more information into source searches and ease their interpretation in terms of physical models, thus making the most of the data we have. Our approach can cope with many free parameters and high model complexity as it uses a Hamiltonian Monte Carlo algorithm for efficient computation. We also present a complementary way to quantify discovery and sensitivity from a Bayesian perspective that avoids some of the pitfalls associated with null hypothesis testing. We plan to apply our approach to the publicly available data from the IceCube experiment and release our framework publicly in the near future.

Acknowledgements

F. Capel acknowledges the financial support from the Excellence Cluster ORIGINS during the development of this work, which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094-390783311. J. Kuhlmann acknowledges support from the DFG through the Sonderforschungsbereich SFB 1258 "Neutrinos and Dark Matter in Astro- and Particle Physics" (NDM).

PoS(ICRC2023)1576

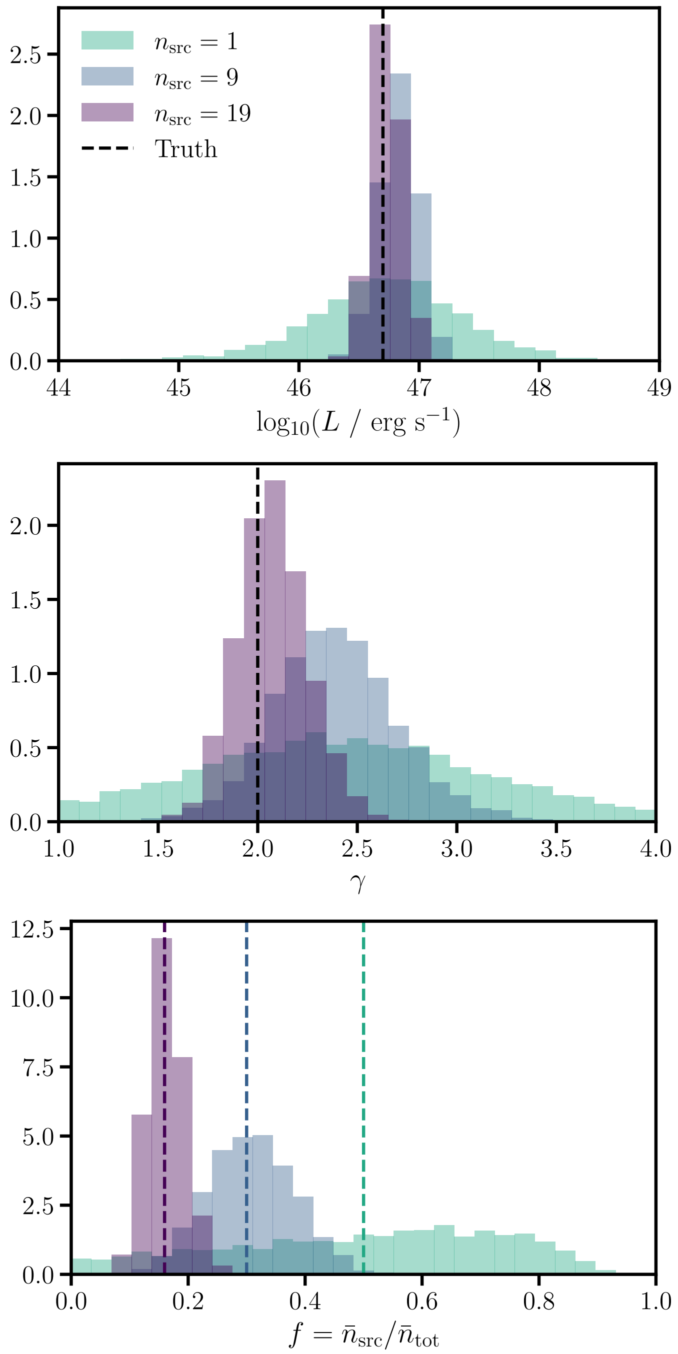

Figure 3: The marginal posterior distributions for L, γ and f are shown in the panels from left to right, respectively. Each panel shows three different distributions corresponding to different values of simulated source events, as shown in the legend. The green, blue and purple distributions correspond to values of $\hat{E}_{\mathrm{th}} = \{20, 5, 1\}$ TeV and $T_{\mathrm{obs}} = \{0.5, 1.25, 2.0\}$ yr. The true simulated values are shown as vertical dashed lines. For f, the value changes for each simulation, so the colour denotes the relevant distribution.

References

[1] N. Kurahashi et al., *Annual Review of Nuclear and Particle Science* **72** (2022) 365 [2203.11936].

[2] J. Braun et al., *Astroparticle Physics* (2008) 299 .

[3] J. Braun et al., *Astroparticle Physics* **33** (2010) 175 [0912.1572].

[4] G. Gigerenzer et al., *The Null Ritual: What You Always Wanted to Know About Significance Testing but Were Afraid to Ask*, in *The SAGE Handbook of Quantitative Methodology for the Social Sciences*, pp. 392 – 409, SAGE Publications, Inc. (2012), DOI.

[5] R.L. Wasserstein et al., *The American Statistician* **70** (2016) 129.

[6] F. Capel et al., *Astronomy & Astrophysics* **668** (2022) 190 [2201.05633].

[7] F. Capel and D.J. Mortlock, *Monthly Notices of the Royal Astronomical Society* **484** (2019) 2324–2340.

[8] T. Budavári and T.J. Loredo, *Annual Review of Statistics and Its Application* **2** (2015) 113 .

[9] D. Heck et al., *CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers*, Tech. Rep. Forschungszentrum Karlsruhe GmbH, Karlsruhe (1998).

[10] T.K. Gaisser, *Astroparticle Physics* **35** (2012) 801.

[11] F. Riehn et al., *The hadronic interaction model Sibyll 2.3c and Feynman scaling*, in *ICRC 2017*, Proceedings of 35th International Cosmic Ray Conference — PoS(ICRC2017), p. 301, 8, 2017, DOI.

[12] M.G. Aartsen et al., *Physical Review Letters* **115** (2015) 96.

[13] IceCube Collaboration, *Icecube data for neutrino point-source searches years 2008-2018*, 2021. 10.21234/CPKQ-K003.

[14] Stan Development Team, *Stan modeling language users guide and reference manual 2.32*, 2023.

[15] Stan Development Team, `cmdstanpy` *(1.1.0)*, 2023.

[16] J.G. Scott and J.O. Berger, *Journal of Statistical Planning and Inference* **136** (2006) 2144.

[17] A. Gelman et al., *Journal of Research on Educational Effectiveness* **5** (2012) 189.

[18] IceCube Collaboration et al., *Science* **361** (2018) eaat1378.

[19] M.G. Aartsen et al., *The European Physical Journal C* **79** (2019) 234 [1811.07979].

[20] A.A. Esfahani et al., *Physical Review C* **103** (2021) 065501 [2012.14341].