

Evaluating cosmic coincidences in the context of astrophysical source populations

Francesca Capel^{a,*}

^a*Technical University of Munich,
Boltzmannstr. 2, 85748 Garching, Germany.
E-mail: f.capel@tum.de*

The multi-messenger era is now well underway, with high-energy neutrinos providing a unique opportunity to study particle acceleration. Recent reports describe possible coincident detections of single IceCube neutrinos with both a flaring blazar and a tidal disruption event. While compelling, these sources cannot be considered in isolation. We present various strategies to put these associations into the context of the relevant astrophysical source populations. Firstly, we can use the non-observation of point sources in IceCube searches to place constraints on the high-level properties of the unknown source population. In particular, current measurements disfavour populations of rare and bright sources. Secondly, multi-messenger simulations of proposed populations and their transient behaviour can be used to evaluate the probability of chance coincident detections in a principled manner. Finally, these simulations can also be harnessed to predict the contribution to the overall neutrino flux that is consistent with an assumed source–neutrino association. We demonstrate the application of these methods, using the proposed blazar–neutrino connection as a case study. The results raise further questions for the bigger picture of neutrino astrophysics.

*37th International Cosmic Ray Conference (ICRC 2021)
July 12th – 23rd, 2021
Online – Berlin, Germany*

*Presenter

1. Introduction

Multi-messenger observations give new insights into the behaviour of astrophysical sources. Strong efforts in the experimental community mean that we now have more data available than ever before, with even larger-scale experiments and surveys planned for the near future. As the amount of data from different messengers increases and we initiate more targeted search efforts, we hope to increase the potential for discoveries. However, it is also important to note that the potential for chance coincidences also grows. We typically search for coincident signals by looking for spatial or temporal correlations. There are cases where it may be evident that two signals are connected if they originate at the same location and time, or disconnected if they arrive from completely different regions of the sky. However, it is also possible to find interesting associations that are uncertain due to poor signal localisation or loose temporal connections. In this case, the statistical methods and information used to assess these connections can significantly impact the conclusions drawn.

Two prominent examples in neutrino astronomy are the detection of single high-energy neutrino IceCube alert events with the blazar TXS 0506+056 [1], and the tidal disruption event (TDE) AT2019dsg [2]. In both cases, the result is reported as a p-value with a $\sim 3\sigma$ significance level, and the sources lie within the 90% confidence region of the reconstructed alert event direction. Additionally, the reconstructed alert event energy is $\gtrsim 200$ TeV, but still has a $\sim 40\%$ chance of being of atmospheric rather than astrophysical origin. The temporal connection is also not so clear as the alerts arrive during an active period of the source with a duration of ~ 6 months. If these associations are confirmed, it would have important implications for modelling particle acceleration in these sources. As the time scale for confirming these with future observations is uncertain and potentially long, it is worthwhile to explore how to use existing data to address these questions in complementary ways.

Here, we explore ways to use information from the relevant source populations to put potential associations into a wider context, using the above blazar–neutrino association as a relevant example. These ideas can also be applied to other multi-messenger associations. A full exposition can be found in [3] and [4], but we present some motivating factors and simplified calculations in this work. In Section 2, we discuss developments in the statistical methods used to evaluate individual source–event coincidences and how more information can be brought into the analysis. Section 3 then continues with general population considerations and constraints. In Section 4 we show how simple calculations can be used to estimate the chance coincidence rates and population implications for the cases of blazars, motivating further study with dedicated simulations. Finally, we conclude in Section 5.

2. Individual associations

The standard approach to searching for point sources in neutrino telescopes uses a frequentist hypothesis testing approach, as described in [5, 6]. Recently, several alternative approaches making use of Bayesian methods have been proposed and applied to multi-messenger searches (see, e.g. [7–10]). Much of this work builds upon ideas reviewed in the work of Budavári and Loredo [11]. In particular, Bayesian hierarchical approaches allow us to close the gap between theory and data by bringing more information about the known physics of a problem directly into the statistical analysis.

The motivation for doing this is outlined in Figure 1 for a simple example. In practice, we also want to include information regarding the detection process, relevant sources of uncertainty, and possible multi-messenger connections. By conditioning on this extra information, it is possible to design more sensitive and principled analyses. Another benefit is that source detection can be assessed independently of standard methods in a complementary way, given the underlying differences in the frequentist and Bayesian paradigms.

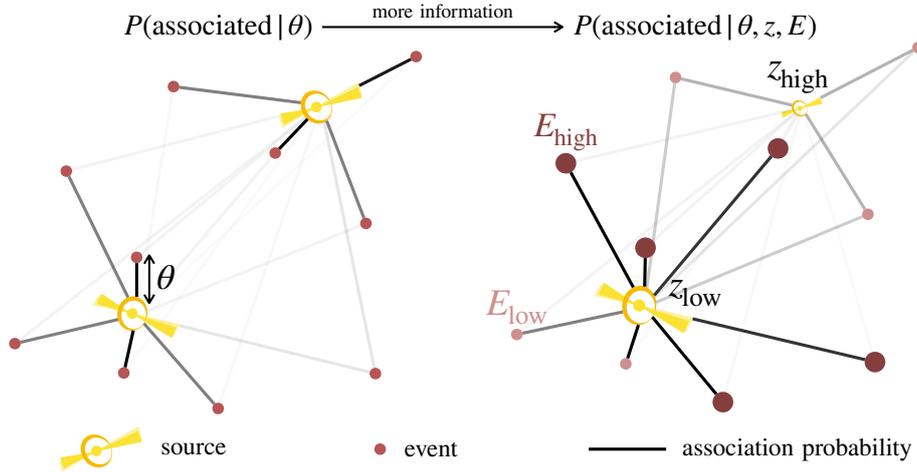


Figure 1: The left panel shows the case of a correlation analysis where the association of sources and events depends purely on their angular separation, θ . The weight of the connecting lines represents the association probability. Consider a case where the upper source is at high redshift, and as such, we do not expect very energetic events to come from it. The right panel shows this case with the more energetic events as larger, darker markers. The association probability then depends on the angular separation, the energy, E , and the redshift, z . Ignoring these factors would lead to inaccurate calculation of the association probabilities. In reality, there are many more factors to be considered than those shown here.

These methods can provide an extendable framework upon which future analyses can be built. However, for complex models, it is not always possible to include all the available and relevant data due to constraints on the tractability of the likelihood and computational restrictions. For this reason, we should also explore other strategies to study and understand potential source–event associations.

3. General population constraints

The sources of astrophysical neutrinos remain elusive, but it is generally thought that they come from point-like, extra-Galactic sources that are capable of accelerating cosmic rays to ~ 100 PeV energies. We expect that any individual detected source will belong to a population of sources with similar properties, and this bigger picture should also comply with the existing neutrino observations. Thanks to electromagnetic surveys, we usually have some knowledge of the luminosity function and density distribution of candidate source populations. This information can be used to study the multi-messenger implications of interesting source–event associations, and to question their validity from a population perspective.

We start by considering constraints on a generic population of extra-Galactic sources. The IceCube neutrino observatory has detected a flux of astrophysical neutrinos [12, 13], but no obvious point sources in the time-integrated data [14]. As neutrinos are weakly interacting, we expect much of the observed flux to appear as diffuse due to the integrated contribution of distant sources that are too faint to be detected individually. If we want to explain the astrophysical flux with a single source population, it must be powerful enough to account for the flux observed by IceCube, whilst not being too sparse or bright, as then we would be able to detect point sources with existing instrument exposures.

Several authors have used this logic to study constraints on the effective density and luminosity of proposed source populations [15–19]. In order to calculate the results, it is necessary to assume some cosmological density evolution model and quantify the detection threshold for point sources seen by IceCube. Typically, this is done by considering a set of discrete evolution models to judge the impact on the final constraints. Source detection is quantified by arguing that at high energies (100 – 200 TeV) if a single source produces two or more events from the same direction (a multiplet) in the detector, it will be detected. For this reason, the results of such analyses are often referred to as “multiplet constraints”. However, point source searches such as that described in [14] make use of larger samples of lower energy ($\gtrsim 100$ GeV) neutrinos to increase their sensitivity, with ~ 5 to 20 signal events required for a 5σ detection, depending on the location of the source in the sky. Additionally, as IceCube continues to gather data, the probability of high-energy multiplets occurring by chance increases. Indeed, we see that there are already several multiplets above 50 TeV in the public IceCube 3-year dataset [20], with the exact number depending on the declination and angular resolution cuts considered.

In our recent work, we go beyond these assumptions by including the important effects of uncertainties in the density evolution of the unknown neutrino sources and a rigorous description of source detection based on the likelihood ratio method described in [5]. To achieve this, we develop a hierarchical statistical formalism that connects the high-level population parameters to observable quantities. We focus on the Northern sky muon track events detected by IceCube as the most constraining dataset due to its low background level and high angular resolution. To place our constraints, we make use of the point source search results reported in [21] together with the diffuse flux characterisation detailed in [22, 23], as these are based on the same underlying event sample. This analysis is detailed in [3], with the code also available as the `nu_pop` GitHub repository¹. Here, we briefly summarise the resulting constraints in Figure 2, for the case of two priors on the source density evolution, encompassing both plausible positive and negative evolutions. In general, rare and bright sources with $n_0 \lesssim 10^{-10}$ Mpc⁻³ and $L \gtrsim 10^{46}$ erg s⁻¹ are incompatible with the neutrino observations. The constraints are compared to the expected values for the two blazar types, BL Lacertae-like objects (BL Lacs) and flat spectrum radio quasars (FSRQs), which are expected to have negative and positive evolutions, respectively. We see that both types of blazars are strongly constrained, implying that they cannot dominate the observed neutrino flux. However, as we consider time-integrated data here, this does not necessarily pose a problem for the transient association with TXS 0506+056 reported in [1].

¹https://github.com/cescalara/nu_pop

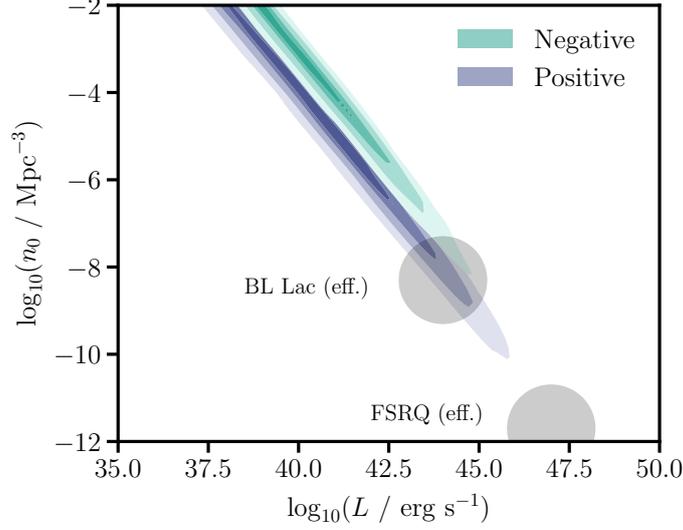


Figure 2: Constraints on the local density, n_0 and luminosity, L , of a generic source population. The contours show the 30%, 60%, 90% and 99% highest posterior density credible regions for the joint distributions of n_0 and L . The results are shown for both positive and negative source density evolutions, and compared with the expected values for BL Lac and FSRQ blazars, based on [19].

4. The blazar–neutrino connection

Section 3 demonstrates how the neutrino observations can be used to constrain a generic source population. To examine a possible blazar–neutrino connection more closely, it is necessary to include further details regarding the proposed association and specific source population. Here, we focus on the case of the IceCube event IceCube-170922A and the blazar TXS 0506+056. As presented in [1], the 3σ result relies on a direct connection between the integrated gamma-ray and neutrino fluxes at the time of the observation. Removing this assumption reduces the significance to $\sim 1.4\sigma$, so the connection between gamma-rays and neutrinos is central to the result. The blazar TXS 0506+056 was initially thought to be a BL Lac, but has been flagged as a “masquerading BL Lac”, an intrinsic FSRQ with hidden broad emission lines [24]. For simplicity, we consider a possible connection with the entire blazar population here, including both BL Lacs and FSRQs.

We can use what we know about blazars from gamma-ray observations and the detection of neutrino alerts to estimate the number of coincidences between high-energy neutrino alerts and flaring blazars in a given observation that occur simply by chance. We start by assuming that there are a fixed number, N_{blazar} , of gamma-ray detected blazars, and that f_{var} of these sources are variable. The variable sources in the detected population are assumed to be flaring a fraction f_{flare} of the time. For an annual rate of neutrino alert events of R_ν , each with a reconstructed localisation of solid angle Ω_ν , the expected number of spatial coincidences between neutrino alerts and blazars is

$$N_{\text{co}}^{\text{spatial}} = N_{\text{blazar}} \frac{R_\nu T_{\text{obs}} \Omega_\nu}{4\pi}, \quad (1)$$

where T_{obs} is the duration of the observation period. The expected number of coincidences between

neutrino alerts and flaring blazars can then be approximated by

$$N_{\text{co}}^{\text{flare}} = N_{\text{co}}^{\text{spatial}} f_{\text{var}} f_{\text{flare}}. \quad (2)$$

For a rough estimate of these values, we take $N_{\text{blazar}} = 3000$ based on the Fermi 4FGL catalog [25], as well as $f_{\text{var}} = f_{\text{flare}} = 0.1$, based on the second Fermi FAVA catalog of variable sources [26]. For the IceCube neutrino alerts we assume $R_{\nu} = 7 \text{ yr}^{-1}$ and $\Omega_{\nu} \sim 10^{-3} \text{ sr}$ based on the public IceCube alert catalog [27]. Using these values along with the relevant observation period of 7.5 yr gives $N_{\text{co}}^{\text{spatial}} \sim 10$ and $N_{\text{co}}^{\text{flare}} \sim 0.1$. While this is a very simplified calculation, we see that we expect to find several spatial coincidences, and a single chance coincidence is not implausible, meriting further investigation. The estimated number of coincidences is also clearly proportional to T_{obs} , showing that we can expect a larger number of chance coincidences as we continue to search for associations.

Similarly, we can also estimate how many neutrino alerts we would expect to see from the whole population, if the connection between IceCube-170922A and TXS 0506+056 is taken to be true. Let us assume a simplified population that is completely characterised by a constant density, n , and a single luminosity, L . For a Euclidean Universe, the gamma-ray flux of a single source is given by

$$\phi_{\gamma} = \frac{L}{4\pi r^2}, \quad (3)$$

where r is the distance to the source and L is defined for the energy range of 0.1 to 100 GeV. The total gamma-ray flux from sources out to some maximum distance, r_{max} , can be expressed as

$$\Phi_{\gamma} = \int_0^{r_{\text{max}}} dr 4\pi r^2 n \phi_{\gamma}(r) = n L r_{\text{max}}. \quad (4)$$

We then assume a simple calorimetric connection between gamma-rays and neutrinos such that

$$\Phi_{\nu} = \epsilon \Phi_{\gamma}, \quad (5)$$

where the relevant energy range for neutrinos is 200 TeV to 7.5 PeV. From this, we can estimate the total number of neutrino alerts in a detector as

$$N_{\nu}^{\text{alert}} = \frac{\Phi_{\nu}}{\langle E \rangle} A_{\text{eff}} T_{\text{obs}}, \quad (6)$$

where $\langle E \rangle$ is the expected energy value, dependent on the spectral distribution, and A_{eff} is the effective area. As the neutrino alerts detected thus far appear to be roughly isotropically distributed on the sky, it can also be interesting to estimate the expected number of alert multiplets. We do this here by considering the number of sources likely to produce more than one neutrino in a detector. This number can be approximated by

$$N_{\text{src}}^{\text{multi}} \approx \sqrt{\pi} N_{\text{src}} (\phi_{\nu} \geq \phi_{\nu}^1) = \sqrt{\pi} n \frac{4}{3} \pi r_1^3, \quad (7)$$

where ϕ_1 is the neutrino flux for a single source that leads to one expected neutrino alert, and r_1 is the distance corresponding to this flux. We have also made use of an approximation detailed in [15].

To estimate N_{ν}^{alert} and $N_{\text{src}}^{\text{multi}}$, we first consider the value for ϵ inferred by assuming the neutrino flux from TXS 0506+056 to equal ϕ_{ν}^1 . This is what one might naively expect, given a true association of a single neutrino. In [1], $L_{\nu}/L_{\gamma} \sim 0.5\text{--}1.7$, depending on the timescale over which neutrino emission takes place. Taking $\epsilon = 0.5$, we consider a simple blazar-like population with $L_{\gamma} = 10^{46} \text{ erg s}^{-1}$, $n = 10^{-9} \text{ Mpc}^{-3}$, and $r_{\text{max}} = 10 \text{ Gpc}$. We also assume $\langle E \rangle = 1 \text{ PeV}$ and $A_{\text{eff}} = 50 \text{ m}^2$ for the high-energy alert events. With these assumptions, we estimate $N_{\nu}^{\text{alert}} \sim 400$ and $N_{\text{src}}^{\text{multi}} \sim 40$ for $T_{\text{obs}} = 7.5 \text{ yr}$, clearly overshooting the ~ 50 observed alerts and lack of observed multiplets. We could also assume that instead the neutrino emission is only due to flaring sources, represented by a lower effective density, $n = 10^{-11} \text{ Mpc}^{-3}$, and higher effective luminosity, $L = 10^{47} \text{ erg s}^{-1}$. In this case we arrive at a more reasonable $N_{\nu}^{\text{alert}} \sim 40$, but still overshoots the expected multiplets with $N_{\text{src}}^{\text{multi}} \sim 10$.

As we realistically expect that other sources in the population may contribute on the same level as TXS 0506+056, the individual neutrino contribution from this blazar could be $\ll 1$, but the integrated contribution from all sources in the population could still be $\mathcal{O}(1)$ (see e.g. [28]). So, allowing lower values of ϵ is still consistent with a real association, and the population can then be used to constrain ϵ to consistent values.

The calculations shown above serve to illustrate the high-level ideas, but miss many important factors that should be included in a realistic assessment. For example, the blazar population has a distribution of luminosities and a source density that evolves over cosmological scales. A realistic cosmological model should also be used, and selection effects accounted for. Additionally, the variability and spectral models for blazars are complex and stochastic, with variations expected from survey to survey. On the neutrino alerts side, the rate and angular resolution are a function of energy, and therefore the spectral models and a more complex detector model is required. In [4], we make use of Monte Carlo simulations with the python packages `popsynth`² and `icecube_tools`³ to overcome these limitations.

5. Conclusions

We highlight ways in which the relevant source populations can be used to assess the validity of proposed multi-messenger observations. The case of the high-energy neutrino alert coincident with the blazar TXS 0506+056 is used as an example, but these methods can also be applied more widely.

Acknowledgements

Francesca Capel acknowledges financial support from the Excellence Cluster ORIGINS, which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094-390783311.

References

- [1] IceCube Collaboration, Fermi-LAT et al., *Science* **361** (2018) eaat1378 [1807.08816].

²<https://github.com/grburgess/popsynth>

³https://github.com/cescalara/icecube_tools

- [2] R. Stein et al., *Nat. Astron.* **5** (2021) 510 [[2005.05340](#)].
- [3] F. Capel, D. J. Mortlock & C. Finley, *Phys. Rev. D* **101**, (2020) 123017 [[2005.02395](#)].
- [4] F. Capel et al., in prep. (2021).
- [5] J. Braun et al., *Astropart. Phys.* **29** (2008) 299 [[0801.1604](#)].
- [6] J. Braun et al., *Astropart. Phys.* **33** (2010) 175 [[0912.1572](#)].
- [7] G. Ashton et al., *ApJ* **860** (2018) 6 [[1712.05392](#)].
- [8] I. Bartos et al., *Phys. Rev. D* **100** (2019) 083017 [[1810.11467](#)].
- [9] F. Capel & D. J. Mortlock, *MNRAS* **484** (2019) 2324 [[1811.06464](#)].
- [10] D. Veske et al., *ApJ* **908** (2021) 216 [[2010.04162](#)].
- [11] T. Budavári & T. J. Loredo, *Annu. Rev. Stat. Appl.* **2** (2015) [[1610.04178](#)].
- [12] IceCube Collaboration, *Science* **342** (2013) 1242856 [[1311.5238](#)].
- [13] IceCube Collaboration, *ApJ* **809** (2015) 98 [[1507.03991](#)].
- [14] IceCube Collaboration, *Phys. Rev. Lett.* **124** (2020) 051103 [[1910.08488](#)].
- [15] P. Lipari, *Phys. Rev. D* **78** (2008) 083011 [[0808.0344](#)].
- [16] A. Silvestri & S. Barwick, *Phys. Rev. D* **81** (2010) 023001 [[0908.4266](#)].
- [17] M. Ahlers & F. Halzen, *Phys. Rev. D* **90** (2014) 043005 [[1406.2160](#)].
- [18] M. Kowalski, *J. Phys.: Conf. Ser.* **632** (2015) 012039 [[1411.4385](#)].
- [19] K. Murase & E. Waxman, *Phys. Rev. D* **94** (2016) 103006 [[1607.01601](#)].
- [20] IceCube Collaboration, dataset (2018) [[10.21234/B4F04V](#)].
- [21] IceCube Collaboration, *Eur. Phys. J. C* **79** (2019) 234 [[1811.07979](#)].
- [22] IceCube Collaboration, *ApJ* **833** (2016) 3 [[1607.08006](#)].
- [23] C. Haack & C. Wiebusch, *Proc. ICRC 2017*, **301** 1005.
- [24] P. Padovani et al., *MNRAS* **484** (2019) L104 [[1901.06998](#)].
- [25] Fermi-LAT Collaboration, *ApJS* **247** (2020) 33 [[1902.10045](#)].
- [26] Fermi-LAT Collaboration, *ApJS* **846** (2017) 34 [[1612.03165](#)].
- [27] IceCube Collaboration, dataset (2018) [[10.21234/B4KS6S](#)].
- [28] Strotjohann et al., *A & A* **622** (2019) L9 [[1809.06865](#)].