

MIT Open Access Articles

A Search for Neutrino Point-source Populations in 7 yr of IceCube Data with Neutrino-count Statistics

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: 10.3847/1538-4357/ab7af9

Publisher: American Astronomical Society

Persistent URL: <https://hdl.handle.net/1721.1/132207>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.





A Search for Neutrino Point-source Populations in 7 yr of IceCube Data with Neutrino-count Statistics

M. G. Aartsen¹, M. Ackermann², J. Adams¹, J. A. Aguilar³, M. Ahlers⁴ , M. Ahrens⁵, C. Alispach⁶, K. Andeen⁷, T. Anderson⁸, I. Ansseau³, G. Anton⁹, C. Argüelles¹⁰, J. Auffenberg¹¹, S. Axani¹⁰, P. Backes¹¹, H. Bagherpour¹, X. Bai¹², A. Balagopal V.¹³, A. Barbano⁶, S. W. Barwick¹⁴, B. Bastian², V. Baum¹⁵, S. Baur³, R. Bay¹⁶, J. J. Beatty^{17,18}, K.-H. Becker¹⁹, J. Becker Tjus²⁰, S. BenZvi²¹, D. Berley²², E. Bernardini^{2,56} , D. Z. Besson^{23,57}, G. Binder^{16,24}, D. Bindig¹⁹, E. Blaufuss²², S. Blot², C. Bohm⁵, M. Börner²⁵, S. Böser¹⁵, O. Botner²⁶, J. Böttcher¹¹, E. Bourbeau⁴, J. Bourbeau²⁷, F. Bradascio², J. Braun²⁷, S. Bron⁶, J. Brostean-Kaiser², A. Burgman²⁶, J. Buscher¹¹, R. S. Busse²⁸, T. Carver⁶, C. Chen²⁹, E. Cheung²², D. Chirkin²⁷, S. Choi³⁰, K. Clark³¹, L. Classen²⁸, A. Coleman³², G. H. Collin^{10,33}, J. M. Conrad¹⁰, P. Coppin³⁴, P. Correa³⁴, D. F. Cowen^{8,35}, R. Cross²¹, P. Dave²⁹ , C. De Clercq³⁴, J. J. DeLaunay⁸ , H. Dembinski³², K. Deoskar⁵, S. De Ridder³⁶, P. Desiati²⁷, K. D. de Vries³⁴, G. de Wasseige³⁴, M. de With³⁷, T. DeYoung³⁸, A. Diaz¹⁰, J. C. Díaz-Vélez²⁷, H. Dujmovic³⁰, M. Dunkman⁸, E. Dvorak¹², B. Eberhardt²⁷, T. Ehrhardt¹⁵, P. Eller⁸, R. Engel¹³, P. A. Evenson³² , S. Fahey²⁷, A. R. Fazely³⁹, J. Felde²², K. Filimonov¹⁶, C. Finley⁵, A. Franckowiak² , E. Friedman²², A. Fritz¹⁵, T. K. Gaisser³², J. Gallagher⁴⁰ , E. Ganster¹¹, S. Garrappa² , L. Gerhardt²⁴, K. Ghorbani²⁷, T. Glauch⁴¹, T. Glüsenkamp⁹, A. Goldschmidt²⁴, J. G. Gonzalez³², D. Grant³⁸, Z. Griffith²⁷, S. Griswold²¹, M. Günder¹¹, M. Gündüz²⁰, C. Haack¹¹, A. Hallgren²⁶, L. Halve¹¹, F. Halzen²⁷, K. Hanson²⁷, A. Haungs¹³, D. Hebecker³⁷, D. Heereman³, P. Heix¹¹, K. Helbing¹⁹, R. Hellauer²², F. Henningsen⁴¹, S. Hickford¹⁹, J. Hignight⁴², G. C. Hill⁴³, K. D. Hoffman²², R. Hoffmann¹⁹, T. Hoinka²⁵, B. Hokanson-Fasig²⁷, K. Hoshina^{27,58}, F. Huang⁸, M. Huber⁴¹, T. Huber^{2,13}, K. Hultqvist⁵, M. Hünnefeld²⁵, R. Hussain²⁷, S. In³⁰, N. Iovine³, A. Ishihara⁴⁴, G. S. Japaridze⁴⁵, M. Jeong³⁰, K. Jero²⁷, B. J. P. Jones⁴⁶, F. Jonske¹¹, R. Joppe¹¹, D. Kang¹³, W. Kang³⁰, A. Kappes²⁸, D. Kappesser¹⁵, T. Karg², M. Karl⁴¹, A. Karle²⁷, U. Katz⁹ , M. Kauer²⁷, J. L. Kelley²⁷, A. Kheirandish²⁷ , J. Kim³⁰, T. Kintscher², J. Kiryluk⁴⁷, T. Kittler⁹, S. R. Klein^{16,24}, R. Koirala³², H. Kolanoski³⁷, L. Köpke¹⁵, C. Kopper³⁸, S. Kopper⁴⁸, D. J. Koskinen⁴, M. Kowalski^{2,37}, K. Krings⁴¹, G. Krückl¹⁵, N. Kulacz⁴², N. Kurahashi⁴⁹, A. Kyriacou⁴³, M. Labare³⁶, J. L. Lanfranchi⁸, M. J. Larson²², F. Lauber¹⁹, J. P. Lazar²⁷, K. Leonard²⁷, A. Leszczyńska¹³, M. Leuermann¹¹, Q. R. Liu²⁷, E. Lohfink¹⁵, C. J. Lozano Mariscal²⁸, L. Lu⁴⁴, F. Lucarelli⁶, J. Lünemann³⁴, W. Luszczak²⁷, Y. Lyu^{16,24}, W. Y. Ma², J. Madsen⁵⁰, G. Maggi³⁴, K. B. M. Mahn³⁸, Y. Makino⁴⁴, P. Mallik¹¹, K. Mallot²⁷, S. Mancina²⁷, I. C. Mariş³, R. Maruyama⁵¹, K. Mase⁴⁴, R. Maunu²², F. McNally⁵², K. Meagher²⁷, M. Medici⁴, A. Medina¹⁷, M. Meier²⁵, S. Meighen-Berger⁴¹, T. Menne²⁵, G. Merino²⁷, T. Meures³, J. Micallef³⁸, D. Mockler³, G. Momenté¹⁵, T. Montaruli⁶, R. W. Moore⁴², R. Morse²⁷, M. Moulai¹⁰, P. Muth¹¹, R. Nagai⁴⁴, U. Naumann¹⁹, G. Neer³⁸, H. Niederhausen⁴¹, S. C. Nowicki³⁸, D. R. Nygren²⁴, A. Obertacke Pollmann¹⁹, M. Oehler¹⁵, A. Olivas²², A. O'Murchadha³, E. O'Sullivan⁵, T. Palczewski^{16,24}, H. Pandya³² , D. V. Pankova⁸, N. Park²⁷, P. Peiffer¹⁵, C. Pérez de los Heros²⁶, S. Philippen¹¹, D. Pieloth²⁵, E. Pinat³, A. Pizzuto²⁷, M. Plum⁷, A. Porcelli³⁶, P. B. Price¹⁶, G. T. Przybylski²⁴, C. Raab³, A. Raissi¹, M. Rameez⁴, L. Rauch², K. Rawlins⁵³, I. C. Rea⁴¹, R. Reimann¹¹, B. Relethford⁴⁹, M. Renschler¹³, G. Renzi³, E. Resconi⁴¹, W. Rhode²⁵, M. Richman⁴⁹, S. Robertson²⁴, N. L. Rodd^{10,59}, M. Rongen¹¹, C. Rott³⁰, T. Ruhe²⁵, D. Ryckbosch³⁶, D. Rysewyk³⁸, I. Safa²⁷, B. R. Safdi^{10,60} , S. E. Sanchez Herrera³⁸, A. Sandrock²⁵, J. Sandroos¹⁵, M. Santander⁴⁸ , S. Sarkar⁵⁴, S. Sarkar⁴², K. Satalecka², M. Schaufel¹¹, H. Schieler¹³, P. Schlunder²⁵, T. Schmidt²², A. Schneider²⁷, J. Schneider⁹, F. G. Schröder^{13,32}, L. Schumacher¹¹, S. Sclafani⁴⁹, D. Seckel³², S. Seunarine⁵⁰, S. Shefali¹¹, M. Silva²⁷, R. Snihur²⁷, J. Soedingrekso²⁵, D. Soldin³², M. Song²², G. M. Spiczak⁵⁰, C. Spiering², J. Stachurska², M. Stamatikos¹⁷, T. Stanev³², R. Stein², P. Steinmüller¹³, J. Stettner¹¹, A. Steuer¹⁵, T. Stezelberger²⁴, R. G. Stokstad²⁴, A. Stöbl⁴⁴, N. L. Strotjohann² , T. Stürwald¹¹, T. Stuttard⁴, G. W. Sullivan²², I. Taboada²⁹, F. Tenholt²⁰, S. Ter-Antonyan³⁹, A. Terliuk², S. Tilav³², K. Tollefson³⁸, L. Tomankova²⁰, C. Tönnis³⁰, S. Toscano³, D. Tosi²⁷, A. Trettin², M. Tselengidou⁹, C. F. Tung²⁹, A. Turcati⁴¹, R. Turcotte¹³, C. F. Turley⁸, B. Ty²⁷, E. Unger²⁶, M. A. Unland Elorrieta²⁸, M. Usner², J. Vandenbroucke²⁷ , W. Van Driessche³⁶, D. van Eijk²⁷, N. van Eijndhoven³⁴, S. Vanheule³⁶, J. van Santen², M. Vraeghe³⁶, C. Walck⁵, A. Wallace⁴³, M. Wallraff¹¹, N. Wandkowsky²⁷, T. B. Watson⁴⁶, C. Weaver⁴², A. Weindl¹³, M. J. Weiss⁸, J. Weldert¹⁵, C. Wendt²⁷, J. Werthebach²⁷, B. J. Whelan⁴³, N. Whitehorn⁵⁵, K. Wiebe¹⁵, C. H. Wiebusch¹¹, L. Wille²⁷, D. R. Williams⁴⁸, L. Wills⁴⁹, M. Wolf⁴¹, J. Wood²⁷, T. R. Wood⁴², K. Woschnagg¹⁶, G. Wrede⁹, D. L. Xu²⁷, X. W. Xu³⁹, Y. Xu⁴⁷, J. P. Yanez⁴², G. Yodh^{14,61}, S. Yoshida⁴⁴, T. Yuan²⁷, and M. Zöcklein¹¹

IceCube Collaboration

¹ Dept. of Physics and Astronomy, University of Canterbury, Private Bag 4800, Christchurch, New Zealand; analysis@icecube.wisc.edu² DESY, D-15738 Zeuthen, Germany³ Université Libre de Bruxelles, Science Faculty CP230, B-1050 Brussels, Belgium⁴ Niels Bohr Institute, University of Copenhagen, DK-2100 Copenhagen, Denmark⁵ Oskar Klein Centre and Dept. of Physics, Stockholm University, SE-10691 Stockholm, Sweden⁶ Département de physique nucléaire et corpusculaire, Université de Genève, CH-1211 Genève, Switzerland⁷ Department of Physics, Marquette University, Milwaukee, WI 53201, USA⁸ Dept. of Physics, Pennsylvania State University, University Park, PA 16802, USA⁹ Erlangen Centre for Astroparticle Physics, Friedrich-Alexander-Universität Erlangen-Nürnberg, D-91058 Erlangen, Germany¹⁰ Dept. of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA¹¹ III. Physikalisches Institut, RWTH Aachen University, D-52056 Aachen, Germany

- ¹² Physics Department, South Dakota School of Mines and Technology, Rapid City, SD 57701, USA
¹³ Karlsruhe Institute of Technology, Institut für Kernphysik, D-76021 Karlsruhe, Germany
¹⁴ Dept. of Physics and Astronomy, University of California, Irvine, CA 92697, USA
¹⁵ Institute of Physics, University of Mainz, Staudinger Weg 7, D-55099 Mainz, Germany
¹⁶ Dept. of Physics, University of California, Berkeley, CA 94720, USA
¹⁷ Dept. of Physics and Center for Cosmology and Astro-Particle Physics, Ohio State University, Columbus, OH 43210, USA
¹⁸ Dept. of Astronomy, Ohio State University, Columbus, OH 43210, USA
¹⁹ Dept. of Physics, University of Wuppertal, D-42119 Wuppertal, Germany
²⁰ Fakultät für Physik & Astronomie, Ruhr-Universität Bochum, D-44780 Bochum, Germany
²¹ Dept. of Physics and Astronomy, University of Rochester, Rochester, NY 14627, USA
²² Dept. of Physics, University of Maryland, College Park, MD 20742, USA
²³ Dept. of Physics and Astronomy, University of Kansas, Lawrence, KS 66045, USA
²⁴ Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
²⁵ Dept. of Physics, TU Dortmund University, D-44221 Dortmund, Germany
²⁶ Dept. of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden
²⁷ Dept. of Physics and Wisconsin IceCube Particle Astrophysics Center, University of Wisconsin, Madison, WI 53706, USA
²⁸ Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, D-48149 Münster, Germany
²⁹ School of Physics and Center for Relativistic Astrophysics, Georgia Institute of Technology, Atlanta, GA 30332, USA
³⁰ Dept. of Physics, Sungkyunkwan University, Suwon 16419, Republic of Korea
³¹ SNOLAB, 1039 Regional Road 24, Creighton Mine 9, Lively, ON, P3Y 1N2, Canada
³² Bartol Research Institute and Dept. of Physics and Astronomy, University of Delaware, Newark, DE 19716, USA
³³ Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
³⁴ Vrije Universiteit Brussel (VUB), Dienst ELEM, B-1050 Brussels, Belgium
³⁵ Dept. of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, USA
³⁶ Dept. of Physics and Astronomy, University of Gent, B-9000 Gent, Belgium
³⁷ Institut für Physik, Humboldt-Universität zu Berlin, D-12489 Berlin, Germany
³⁸ Dept. of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA
³⁹ Dept. of Physics, Southern University, Baton Rouge, LA 70813, USA
⁴⁰ Dept. of Astronomy, University of Wisconsin, Madison, WI 53706, USA
⁴¹ Physik-department, Technische Universität München, D-85748 Garching, Germany
⁴² Dept. of Physics, University of Alberta, Edmonton, Alberta, T6G 2E1, Canada
⁴³ Department of Physics, University of Adelaide, Adelaide, 5005, Australia
⁴⁴ Dept. of Physics and Institute for Global Prominent Research, Chiba University, Chiba 263-8522, Japan
⁴⁵ CTSPS, Clark-Atlanta University, Atlanta, GA 30314, USA
⁴⁶ Dept. of Physics, University of Texas at Arlington, 502 Yates St., Science Hall Rm 108, Box 19059, Arlington, TX 76019, USA
⁴⁷ Dept. of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA
⁴⁸ Dept. of Physics and Astronomy, University of Alabama, Tuscaloosa, AL 35487, USA
⁴⁹ Dept. of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA
⁵⁰ Dept. of Physics, University of Wisconsin, River Falls, WI 54022, USA
⁵¹ Dept. of Physics, Yale University, New Haven, CT 06520, USA
⁵² Department of Physics, Mercer University, Macon, GA 31207-0001, USA
⁵³ Dept. of Physics and Astronomy, University of Alaska Anchorage, 3211 Providence Dr., Anchorage, AK 99508, USA
⁵⁴ Dept. of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK
⁵⁵ Department of Physics and Astronomy, UCLA, Los Angeles, CA 90095, USA

Received 2019 November 6; revised 2020 February 25; accepted 2020 February 26; published 2020 April 20

Abstract

The presence of a population of point sources in a data set modifies the underlying neutrino-count statistics from the Poisson distribution. This deviation can be exactly quantified using the non-Poissonian template fitting technique, and in this work we present the first application of this approach to the IceCube high-energy neutrino data set. Using this method, we search in 7 yr of IceCube data for point-source populations correlated with the disk of the Milky Way, the Fermi bubbles, the Schlegel, Finkbeiner, and Davis dust map, or with the isotropic extragalactic sky. No evidence for such a population is found in the data using this technique, and in the absence of a signal, we establish constraints on population models with source-count distribution functions that can be described by a power law with a single break. The derived limits can be interpreted in the context of many possible source classes. In order to enhance the flexibility of the results, we publish the full posterior from our analysis, which can be used to establish limits on specific population models that would contribute to the observed IceCube neutrino flux.

⁵⁶ Also at Università di Padova, I-35131 Padova, Italy.

⁵⁷ Also at National Research Nuclear University, Moscow Engineering Physics Institute (MEPhI), Moscow 115409, Russia.

⁵⁸ Earthquake Research Institute, University of Tokyo, Bunkyo, Tokyo 113-0032, Japan.

⁵⁹ Now at Berkeley Center for Theoretical Physics, University of California, Berkeley, CA 94720, USA and Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.

⁶⁰ Now at Leinweber Center for Theoretical Physics, Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA.

⁶¹ Deceased.

Unified Astronomy Thesaurus concepts: [Astrostatistics \(1882\)](#); [Astrostatistics distributions \(1884\)](#); [Bayesian statistics \(1900\)](#); [Neutrino astronomy \(1100\)](#); [High energy astrophysics \(739\)](#); [Neutrino telescopes \(1105\)](#); [Posterior distribution \(1926\)](#)

1. Introduction

The conclusive discovery of an astrophysical neutrino flux at IceCube (Aartsen et al. 2013a, 2013b, 2014a, 2015a, 2015b, 2015c) presents a new window through which we can view the universe. With energies in the PeV range, these neutrinos free stream to Earth over scales where extragalactic photons of the same energy are attenuated. This makes the IceCube neutrino window novel in terms of being a messenger as well as providing insight into extreme energy phenomena. While evidence was recently found for a point source of neutrinos (Aartsen et al. 2018), at present, the origin of a large fraction of astrophysical flux remains unknown. For a review, see Spiering (2018).

With an observed flux close to the Waxman–Bahcall bound (Waxman & Bahcall 1999), the leading hypothesis posits that the IceCube neutrinos are produced through extragalactic hadronic processes, where high-energy proton interactions produce charged pions, which in turn decay to produce neutrinos.⁶² There are a number of viable models for the origin of these hadronic collisions involving conventional astrophysical sources; see Murase & Waxman (2016) for a comprehensive discussion. Nevertheless, some of the promising source classes already appear to be disfavored as the sole origin of the observed cosmic neutrinos. The current limits on the presence of point sources in the data, which we will soon discuss, place under tension a pure blazar origin for the observed flux, although they may still contribute (Padovani & Resconi 2014; Padovani et al. 2015; Petropoulou et al. 2015; Murase & Waxman 2016)—especially in light of the recent discovery (Aartsen et al. 2018). Further, there are claims that a dominant starburst galaxy origin is in tension with existing gamma-ray measurements (Tamborra et al. 2014; Bechtol et al. 2017; Sudoh et al. 2018). Viable source classes remain, however, such as radio galaxies (Murase & Waxman 2016; Hooper 2016; Blanco & Hooper 2017). Regardless of their origin, due to the hadronic origin of the neutrinos in these scenarios, a definitive identification of neutrinos from a particular source class would provide a deep insight into the wider problem of high-energy cosmic-ray acceleration (Learned & Mannheim 2000; Halzen & Hooper 2002; Anchordoqui & Montaruli 2010; Anchordoqui et al. 2014). As such, the implications of the IceCube data set for various source classes, even those coming from null results, are leading to important insights into high-energy astrophysics as a whole.

As long as the question of their origin remains unanswered, however, there will be room for speculation as to a potentially more exotic origin for these neutrinos. A possibility that has received significant attention in the literature is that these neutrinos could be produced in the decay of \sim PeV scale dark matter (see, for example, Feldstein et al. 2013; Esmaili & Serpico 2013; Ema et al. 2014; Bhattacharya et al. 2014; Zavala 2014; Chianese & Merle 2017; Chianese et al. 2017; Schmaltz & Weiner 2017). While extragalactic dark-matter decays would be distributed isotropically, decays within the Milky Way would imprint a directional anisotropy within the

data in such a scenario. Although the present data appear to be isotropic (Aartsen et al. 2015b), the measurements have not yet resulted in enough significance to disfavor the dark-matter scenarios (Esmaili et al. 2014). More stringent constraints arise as, generically, such models produce photons or charged cosmic-rays in addition to the neutrinos, and limits on dark matter from such final states tend to disfavor many scenarios (see Kalashev & Kuznetsov 2016, 2017; Cohen et al. 2017; Kuznetsov 2017; Abeysekara et al. 2018). Nevertheless, the possibility remains that there could be hints as to the origin of dark matter within the IceCube data set, a possibility that highlights the importance of a definitive determination of the neutrino origins.

Due to their lack of electric charge, neutrinos remain unattenuated as they travel through the magnetic fields that permeate the universe. The neutrinos point back to their source of origin, raising the possibility that if the incident neutrino direction can be measured accurately enough, the source class could be identified. A single event is unlikely to be determinative, but statistical analyses applied to larger data sets can search for clustering on the sky, which can indicate the presence of point sources. A number of such searches for point sources have already been performed within IceCube (Abbasi et al. 2011; Aartsen et al. 2013c, 2013d, 2014b, 2015e, 2015f, 2016a, 2017a, 2019; Glauch & Turcati 2018) and ANTARES (Albert et al. 2017) but have not yet found evidence of such clustering. The absence of any statistically significant clustering in these analyses thus far suggests that whatever is the primary contributor to the IceCube flux is not a small number of bright sources but rather a larger population of sources. In the present work, we extend this line of investigation through the application of a novel technique for searching for populations of sources, which should be viewed as complementary to existing and ongoing individual source searches. We accomplish this using a technique known as the Non-Poissonian Template Fit (NPTF), which has found widespread application to the Fermi LAT gamma-ray data but is applied here to the IceCube neutrino data for the first time.

The basic principle underpinning the NPTF is that in the presence of unmodeled point sources, the neutrino-count statistics of a data set is distinct with respect to a Poisson distribution. The requirement that the point sources are unmodeled is central to the method; indeed, if a point source is resolved and has a known location, a model for it can be constructed and the observed data will represent a Poisson draw from that model. The NPTF, however, remains agnostic as to the location of the sources and simply accounts for the fact that a population of point sources will result in larger upward and downward fluctuations than can be produced by the Poisson distribution. As we will review in this work, that deviation can be rigorously quantified into a likelihood, which can test the preference for non-Poissonian statistics in the data, and thereby uncover evidence for the point-source population.

The NPTF has a number of advantages over traditional point source search techniques. The method is naturally couched in the language of populations of sources, as the fundamental object constrained is the source-count distribution dN/dF , that is the distribution of sources with a flux between F and $F + dF$. By way of contrast, the standard techniques search for

⁶² The neutrino flavor ratio observed with IceCube is presently consistent with a pionic origin (Aartsen et al. 2015d).

sources of a given flux F' one at a time. These methods are often calibrated against simulations where N point sources each of flux F' have been injected, which corresponds to the special case $dN/dF = N\delta(F - F')$. Studies using a similar method to the NPTF applied to the public High Energy Starting Events (HESE) data have shown that population-based approaches can, under certain assumptions, probe deeper than traditional methods and search for sources too dim to be resolved individually (Feyereisen et al. 2017, 2018). Such techniques have also been used to extend limits on the brightest possible source within the data (Ando et al. 2017). Furthermore, as the NPTF is a method for template fitting, it can readily incorporate a nontrivial spatial dependence for the sources. Template fitting is a technique of fitting data with models following a predetermined spatial distribution, or templates. We will exploit this to search for sources correlated with the disk of the Milky Way and the Fermi bubbles, in addition to extragalactic sources that are distributed isotropically. Even in the case of isotropically distributed extragalactic sources, due to the spatial variation inherent in the IceCube effective area matrix, the resulting neutrino-count map will be non-isotropic. This complication can be readily handled in the NPTF. On the other hand, there are drawbacks to the method. The NPTF is fundamentally a binned technique, which precludes optimization through the use of event-by-event reconstruction information. Moreover, in this work, the NPTF is restricted to be used in a single energy bin. This is not a fundamental limitation of the method but simply a shortcoming in existing implementations of it. Yet, at present, this implies that we are unable to account for the strong variation of the response to neutrinos in the IceCube detector as a function of energy, other than through optimizing the choice of energy bin. Taken together, these highlight the complementary nature of the NPTF based search presented in this work to alternative approaches to the problem. We emphasize that the NPTF is distinct from other novel searches for point sources that have been applied to the IceCube data set, in that it searches for the specific modification to the neutrino-count statistics imprinted by a point-source population. An example of such techniques is the multipole and two-point autocorrelation approach (Aartsen et al. 2015e; Glauch & Turcati 2018), which search for a statistical increase in the number of events with small angular separation that point sources would induce. Fundamentally, point sources produce additional clustering within the data on the scale of the instrument point-spread function (PSF), which is usually smaller than the scale on which the diffuse backgrounds cluster. In that sense, the NPTF is also looking to exploit similar information, but it approaches the problem by instead considering how a population of point sources modifies the statistics of the neutrino-count map. In addition to its different approach, the NPTF also offers a number of advantages over these techniques. Most importantly, being a template method, it can readily account for the spatial variation expected for galactic sources, or non-isotropic detector response, as discussed above. Further, the NPTF can be formulated in terms of an analytic likelihood, as we will review in this work, and this allows for an efficient practical implementation of the method. The NPTF is also distinct to the search for steady point sources with specific flux-characteristics, as considered in Aartsen et al. (2019), where a test-statistic is estimated for a given population model with parameterized source density and luminosity. For the NPTF, a

broader range of population models are tested, allowing the shape of the source distribution to be additionally parameterized. Constraints on models with specified flux-characteristics can, however, be derived from the NPTF results, and we will demonstrate this explicitly by placing constraints on the space of standard-candle luminosity functions using FIRESONG (Taboada et al. 2018), similar to those considered in Aartsen et al. (2019).

The remainder of the discussion will be structured as follows. In Section 2, we outline the event selection used to distill the data set analyzed in this work. Section 3 is dedicated to a review of the NPTF method, and a detailed description of the challenges in applying the method at IceCube, and the associated solutions. Following this, in Section 4, we determine the expected sensitivity of the method based on Monte Carlo simulations. Then, in Section 5, we show the result of applying the NPTF to the real IceCube data, and in the absence of a signal, we derive constraints. The full posterior from our analysis is made publicly available,⁶³ with the details of the file discussed in Appendix A. Finally, our conclusions are presented in Section 6.

2. Event Selection

2.1. The IceCube Neutrino Observatory

IceCube is a cubic-kilometer Cerenkov detector, which is composed of 5160 digital optical modules (DOMs) embedded in the Antarctic ice at the South Pole Acherberg et al. (2006). These DOMs are attached to 86 strings of cable at depths between 1450 and 2450 m beneath the surface of the ice. Most of the DOMs have a vertical spacing of 17 m along the strings and the average distance between neighboring strings is ~ 125 m. Each module consists of a photomultiplier tube, onboard digitization board, and a separate board with LEDs for calibration (Abbasi et al. 2009, 2010).

Construction of IceCube started in 2004 and was completed in 2010 December. Before the full detector was completed, data were being taken in partial configurations with fewer than 86 strings. In the present work, we make use of 7 yr of IceCube data. The first three years were taken during the 40-string (IC-40), 59-string (IC-59), and 79-string (IC-79) configurations, as described in Aartsen et al. (2013c) and Abbasi et al. (2011). The subsequent four years of data exploited the full 86-string (IC-86) configuration, as outlined in Aartsen et al. (2014b, 2017a).

2.2. Neutrino Detection at IceCube

Neutrinos are notoriously difficult to detect, and just because they point back to their origin does not mean we can necessarily extract that direction. As we cannot detect the neutrinos directly, the challenge is in inferring the direction from visible products left behind from a neutral or charged-current interaction the neutrino undergoes within or in the vicinity of IceCube. In order to enhance our sensitivity, we choose to focus on events where the flight direction of the neutrino can be accurately reconstructed, as explained below.

In detail, there are three different event topologies within IceCube produced by neutrino interactions: tracks, cascades, and double-bangs. Tracks result from muons traversing the detector, while cascades result from the charged-current

⁶³ https://icecube.wisc.edu/science/data/NPTF_7yr_posterior

interactions of electron or tau neutrinos or neutral-current interactions of any neutrino. The interactions in cascade events produce an almost spherical light emission making directional reconstruction difficult. Tracks from muons of \sim TeV or greater energies can travel several kilometers while constantly emitting Cerenkov light, making them ideal candidates for an accurate directional reconstruction. Double-bangs result from charged-current interactions of tau neutrinos at very high energies where the tau lepton decays to hadrons far enough from the initial interaction to create two distinct cascades. The first candidates for such events have now been identified (Stachurska 2018). Since tracks are the optimal topology for directional accuracy, in this paper, we will only use tracks.

The only neutrino interactions that can create tracks at TeV energies are charged-current interactions of muon neutrinos (and muon antineutrinos). Track events can be further divided into two subclasses: starting tracks occur when a muon neutrino has its charged-current interaction inside of the detector volume, while through-going tracks occur when that interaction occurs outside of the detector volume. Through-going tracks can result from any high-energy muon, including muons produced by cosmic-ray interactions in the atmosphere, but they also have a much higher effective area for muon neutrinos than starting tracks since the charged-current interaction can occur in a much larger volume than the detector volume. In this paper, we will consider only events reconstructed as through-going tracks, in order to take advantage of the high number of events resulting from the much larger effective area; however, the lower purity of astrophysical neutrinos in the sample creates a background we will need to account for in the NPTF analysis.

2.3. Data and Simulations

With these motivations, the specific data set used in this analysis was the through-going tracks in IceCube’s 7 yr point-source sample (Aartsen et al. (2017a), and we refer to Section 2.2 of that reference for a detailed account of the data set. The 7 yr point-source sample has been accrued through several different event selections. Events from the year of IC-40 data were selected using fixed selection criteria on several parameters (Abbasi et al. 2011), while events from the remaining years (Aartsen et al. 2013c, 2014b, 2017a) were selected using multivariate boosted decision trees (BDTs) to classify events as signal or background. The BDTs were trained with separate background and signal data sets. In the background case, the BDT was trained on the data themselves, a procedure that is justified as the data are known to be strongly dominated by the background. For the signal, the BDT was instead trained on muon neutrino Monte Carlo simulations. These same Monte Carlo-simulated muon neutrinos are also used to calculate the effective area of the detector and the PSF. More information on the Monte Carlo simulation can be found in Aartsen et al. (2016b). Through these selection processes, the sample is divided into two regions, up-going (with decl. $\delta > -5^\circ$) and down-going ($\delta < -5^\circ$). The down-going region is dominated by atmospheric muons, while the up-going region is shielded from these muons by the Earth. The up-going region is dominated by atmospheric neutrinos. Despite this distinction, in our analysis, we will not distinguish between up- and down-going events, instead performing a full sky analysis. As in Aartsen et al. (2017a), the total livetime is 2431 days, with 422,791 up-going events, and 289,078 down going. The

directions of the events in this sample were reconstructed using a likelihood-based method, which uses information on how the photons scatter and get absorbed with the ice (Ahrens et al. 2004). The reconstructions of events in the IC-86 data used a more advanced description of the ice (Aartsen et al. 2013e) and a better parameterization of how the photons interact with it Aartsen et al. (2014b). The muon energy is estimated by approximating its energy loss along its reconstructed track Aartsen et al. (2014c). In order to fine tune this sample for our analysis, we make a cut on the reconstructed muon energy. The expected energy spectrum for an astrophysical neutrino flux is harder than the energy spectrum for the atmospheric neutrino flux. It should also be noted that the reconstructed muon energy is only an estimate of the muon energy at a point where it enters the detector and, thus, can only provide an estimated lower-bound of the primary neutrino energy, as much of the energy of the event can be deposited outside the detector. Nevertheless, a cut on the reconstructed muon energy can still effectively increase the sample’s purity with respect to the harder-spectrum astrophysical flux. We determine the optimal energy cut to be $10^{0.5}$ TeV \approx 3 TeV. This value is determined by maximizing the sensitivity for the full sky analysis presented in this work. We note that as the energy distribution of the background in the northern and southern hemispheres is quite different, if we performed an analysis restricting to either of these, the optimal energy cut would vary. We also make a spatial cut around the poles of the detector ($|\delta| > 85^\circ$), as the scrambling procedure is less effective, and the reconstructions can be poorly behaved, in these regions. These cuts lead us to our final sample of 309,134 events.

Since the NPTF is a binned analysis, we must spatially bin the data. For this purpose, we use HEALPIX Gorski et al. (2005) skymaps to bin the data into pixels of equal solid angle. There is still freedom in terms of how large to choose these bins, controlled by the `nside` parameter; however we find that a value of 64 maximizes our full-sky sensitivity, which corresponds to 49,152 bins, each of size approximately 0.84 square degrees.

3. The Non-Poissonian Template Fit

The NPTF quantifies the following observation into a rigorous analytic likelihood: compared to a count map following a Poisson distribution, a map determined by a distribution of sources will have more hot and cold pixels—the hot pixels associated with locations where there are sources, and the cold pixels where there are none. In slightly more detail, there are at least two steps involved in going from an underlying point-source distribution to a neutrino-count map. First, we need to determine how many point sources are expected and how they are distributed on the sky. Second, given a dN/dF , we must determine the map of neutrino counts that is expected from this distribution of sources. The NPTF likelihood provides the rigorous answer to these questions, as well as incorporating additional complications arising through detector effects such as the finite PSF of the instrument. This likelihood, when applied to the data, allows for a determination as to whether such a population of sources is preferred, and if not, constraints on dN/dF can be set.

Although the NPTF is a relatively recent method, core aspects have long been employed in astronomy. The fundamental observation that point-source distributions lead to more hot and cold pixels is at the heart of the $P(D)$ method that has

been applied to X-ray data sets (Hasinger et al. 1993; Georgantopoulos et al. 1993; Miyaji & Griffiths 2002; Gendreau et al. 1998; Perri & Giommi 2000). The method was extended to gamma-rays in Malyshev & Hogg (2011), where the analytic likelihood was also first derived. The likelihood was extended to the NPTF, i.e., a full template based method, in Lee et al. (2016), and the NPTF or similar methods have found a number of additional applications in gamma-ray astronomy (Lee et al. 2015; Feyereisen et al. 2015; Zechlin et al. 2016a, 2016b; Linden et al. 2016; Lisanti et al. 2016; Di Mauro et al. 2018). As mentioned above, a related method that is similarly population-based has been considered in the context of publicly available IceCube data (Feyereisen et al. 2017, 2018). The NPTF has now been incorporated into a publicly available code `NPTFit` (Mishra-Sharma et al. 2017),⁶⁴ which we use for the present work.

The remainder of this section will begin with a more quantitative review of the method, outlining the core ideas leading to the NPTF likelihood. After this, we focus on some of the challenges that had to be addressed in order to apply this method to the IceCube data set, which, in particular, required a careful treatment of the instrument effective area matrix and PSF. Third, we will describe how we can combine the NPTF likelihood with additional Poissonian models that we will use in our hypothesis testing, and finally, we outline our inference procedure, which we will use to search for point-source populations in the data.

3.1. Overview of the Method

In this section, we provide a brief, quantitative review of the NPTF likelihood framework, particularly emphasizing aspects that will be relevant to an application at IceCube. A more comprehensive description of the method and a derivation of all quoted expressions can be found in Mishra-Sharma et al. (2017).

Our ultimate goal is to write down a likelihood for a set of model parameters θ , given the data d described in Section 2—i.e., we want a function $\mathcal{L}(d|\theta)$. Let us start with a description of the model parameters. In the case where we only have neutrinos originating from a single point-source population, then the model parameters specify the source-count distribution dN/dF . In principle, it is possible to keep the form of dN/dF very general, but a particularly simple analytic expression for the source-count function that is likely a good approximation to many realistic neutrino source classes is a broken power law⁶⁵

$$\frac{dN_p}{dF}(F; \theta) = A T_p \begin{cases} (F/F_b)^{-n_1} & F \geq F_b, \\ (F/F_b)^{-n_2} & F < F_b. \end{cases} \quad (1)$$

The use of a common functional form to describe the flux of both galactic and extragalactic sources is motivated by the fact that this ansatz appears to be a reasonable description of both populations in gamma-rays observed by Fermi (see, for example, Lee et al. 2016). In Equation (1), we have added the term T_p to the source-count function. T_p is a template, or pixelated spatial map describing the spatial distribution of the sources on the sky. It is the only term with an explicit pixel-by-pixel variation. As an example, for an isotropic or extragalactic

distribution of sources, we have $T_p = 1$ for every value of p . This template could be considered as a model parameter and fitted for, but in our analysis, we will take T_p to be fixed before performing any likelihood analysis. With the template fixed, there are four model parameters for a singly broken power law: the normalization A , location of the break F_b , and power-law indices above and below the break n_1 and n_2 . Thus, formally $\theta = \{A, F_b, n_1, n_2\}$.

To provide some intuition for these parameters, note that we can calculate the total expected number of point sources and also the total expected flux from the population in each pixel from direct integration of the source-count function over all possible fluxes as follows

$$\begin{aligned} N_p^{\text{PS}} &= \int_0^\infty dF \frac{dN_p}{dF} = \frac{A T_p F_b (n_1 - n_2)}{(n_1 - 1)(1 - n_2)}, \\ F_p^{\text{PS}} &= \int_0^\infty dF F \frac{dN_p}{dF} = \frac{A T_p F_b^2 (n_1 - n_2)}{(n_1 - 2)(2 - n_2)}. \end{aligned} \quad (2)$$

In performing these integrals, we require $n_1 > 1$ and $n_2 < 1$ to obtain a finite N_p^{PS} , whereas for a finite F_p^{PS} we need $n_1 > 2$ and $n_2 < 2$. There is an important distinction between the number of sources and flux as given in Equation (2). The total expected flux per pixel, F_p^{PS} , is related to the expected number of neutrinos observed through the IceCube effective area, and in this sense is tied to an observable in the data. Yet the total expected number of sources, N_p^{PS} , is not tied to an observable derivable from a map of neutrino counts. In this sense, a real map could have a best-fit value $n_2 \in [1, 2]$, which corresponds to an infinite number of sources but a finite flux. While in practice, the total number of sources should be finite, the effective number of sources, for realistic source populations such as star-forming galaxies (Tamborra et al. 2014), can appear infinite. This effect occurs because the cutoff in the source-count distribution, dN/dS , that makes the number of sources finite appears at flux values well below the level where the sources contribute, on average, more than one photon or neutrino (Lisanti et al. 2016).

When presenting results, it will be helpful to use a different set of variables instead of $\{A, F_b, n_1, n_2\}$. Specifically, we will replace A and F_b with the expected number of point sources across the whole sky, N^{PS} , and the expected flux per source, \bar{F}^{PS} . The change of variables can be implemented as follows:

$$\begin{aligned} N^{\text{PS}} &= \sum_p N_p^{\text{PS}} = \frac{A F_b (n_1 - n_2)}{(n_1 - 1)(1 - n_2)} \sum_p T_p, \\ \bar{F}^{\text{PS}} &= \frac{\sum_p F_p^{\text{PS}}}{\sum_p N_p^{\text{PS}}} = F_b \frac{(n_1 - 1)(1 - n_2)}{(n_1 - 2)(2 - n_2)}. \end{aligned} \quad (3)$$

Given a template and values of n_1 and n_2 , we can then change variables to these more intuitive quantities, which we will exploit when presenting our results.

At this point, we comment on an important assumption that will be used throughout our analysis. We have repeatedly discussed the flux of sources, F , eschewing the question of what energy this flux is being measured at. To resolve this, we will assume that our point-source population follows the canonical astrophysical expectation of E^{-2} . In detail the flux

⁶⁴ <https://github.com/bsafdi/NPTFit/>

⁶⁵ In passing, we note that `NPTFit` can handle a broken power-law source-count function with an arbitrary number of breaks.

from an individual source is given by

$$\frac{d\Phi(E)}{dE} = F \left(\frac{E}{1 \text{ TeV}} \right)^{-2}, \quad (4)$$

which serves to contextualize the quantity F discussed throughout this section. It should be thought of as the normalization constant for the energy dependent flux, and we will assume an identical E^{-2} scaling for all sources. If the actual spectrum is softer than this, the events will be shifted to lower energies where the backgrounds are higher, and generically, we would expect a reduction of the sensitivities shown here.

Returning to our derivation of the likelihood, the source-count function parameterizes our model prediction for the population of sources, but our goal is to embed this into a likelihood that can be fit to neutrino-count data. As a first step toward this, we need to address the fact that the discussion so far has been couched in the language of fluxes per neutrino source, Φ , commonly quoted with units of [neutrinos $\text{cm}^{-2} \text{s}^{-1}$], whereas what is observed in the instrument is an integer number of neutrino events. The conversion between these two variables is provided by the effective area matrix, which accounts for the fact that two-point sources with equal flux at different locations on the sky will contribute a different number of detected neutrinos within IceCube. The effective area is the amalgamation of the detection efficiency for neutrinos incident on the IceCube detector from different directions, as well as an accounting for the fact the detector has a fixed location at the South Pole.

The conversion from flux, F , in units of [neutrinos $\text{cm}^{-2} \text{s}^{-1}$] to counts, S , is achieved with the combination of the effective area and collection time, usually called an exposure map, which we denote by \mathcal{E}_p —a pixel dependent quantity due to the spatial dependence in the effective area matrix. We defer the discussion of how the appropriate \mathcal{E}_p map for our analysis was derived until the following subsection. Assuming for now that the map is known, then using this, we can then convert to a source-count function in terms of counts rather than flux as follows:

$$\frac{dN_p(S; \theta)}{dS} = \frac{1}{\mathcal{E}_p} \frac{dN_p(F = S/\mathcal{E}_p; \theta)}{dF}. \quad (5)$$

As \mathcal{E}_p can take on a different value in every single pixel, the conversion from flux to counts should in truth be performed in every pixel. Nevertheless, this is in practice often unnecessary. It is usually sufficient to divide the full map into a number of subregions, which each have comparable \mathcal{E}_p values, take the mean \mathcal{E}_p in this region, and perform the conversion once per region. Within the NPTFit framework, the number of subregions is controlled by the keyword `nexp`, and results are commonly convergent for small values of this parameter. As the appropriate \mathcal{E}_p for our data set varied significantly, albeit smoothly, as a function of decl., we chose to use 50 exposure regions in order to ensure the transformation from flux to counts was accurately performed.

From dN_p/dS , we can then move toward the NPTF likelihood by deriving the following useful quantity: the expected number of sources that will contribute m neutrinos in a pixel p , $x_{p,m}(\theta)$. To do so, note that dN_p/dS evaluated at a particular S provides the expected number of sources that contribute an expected number of counts S , where of course S

does not need to be an integer. The probability that one such source provides m neutrinos is then determined by the Poisson distribution, specifically $S^m e^{-S}/m!$. From here, $x_{p,m}(\theta)$ is given by weighting this factor by the source-count distribution and integrating over all S , as each value could Poisson fluctuate to m . In detail,

$$x_{p,m}(\theta) = \int_0^\infty dS \frac{dN_p(S; \theta)}{dS} \frac{S^m e^{-S}}{m!}. \quad (6)$$

This expression, while intuitive, has an inherent assumption that will be invalidated in most real applications: if a point source is located in pixel p , then it deposits all of its observed neutrinos in that same pixel. This neglects the finite PSF at IceCube, but we will hold off on a discussion of how to address this until the next subsection.

Our final goal of this subsection is to move from $x_{p,m}$ to the probability of observing k neutrinos in a pixel p , $p_{p,k}$. Combining $p_{p,k}$ with the observed number of neutrinos in the data, d_p , and then taking the product over all pixels p , exactly gives us the likelihood through which we can constrain dN/dF , or more specifically $\theta = \{A, F_b, n_1, n_2\}$. Being fully explicit, we have:

$$\mathcal{L}(d|\theta) = \prod_p p_{p,k=d_p}(\theta), \quad (7)$$

where the product is taken over all pixels in the data set analyzed. In order to derive an expression for $p_{p,k}$, we use the concept of probability generating functions.⁶⁶ If we have a discrete probability distribution described by a set $\{p_{p,k}\}$ known for all k , then the generating function in a given pixel is defined as

$$P_p(t) \equiv \sum_{k=0}^{\infty} p_{p,k} t^k, \quad (8)$$

where t is an auxiliary variable. The probabilities can be recovered from $P_p(t)$ via

$$p_{p,k} = \left. \frac{1}{k!} \frac{d^k P_p(t)}{dt^k} \right|_{t=0}. \quad (9)$$

In the case of models described by the Poisson distribution with expected counts μ_p , substituting the Poisson distribution into Equation (8) reveals the associated generating function to be $P_p(t) = \exp[\mu_p(t - 1)]$.

Next, the aim is to construct the associated non-Poissonian generating function, starting with the expected number of m -neutrino sources, $x_{p,m}(\theta)$, as given in Equation (6). Of course, m can take on any integer value, but for the moment, let us take it to be fixed, and we will determine the generating function for m -neutrino sources, denoted $P_p^{(m)}(t)$. From the definition in Equation (8), we need to know the probability of seeing k neutrinos in the pixel p , given by $p_{p,k}$, a value that will depend on how many m -neutrino sources there are. Specifically, k must be some integer n_m multiple of m , where n_m drawn from a Poisson distribution with mean $x_{p,m}(\theta)$,

$$p_{p,n_m} = \frac{x_{p,m}^{n_m} e^{-x_{p,m}}}{n_m!}, \quad (10)$$

⁶⁶ For a review, see, for example, Section 3.6 of Hoel et al. (1971).

where we have left the dependence on the parameters θ implicit. In terms of this, the probability for seeing k neutrinos in pixel p is simply p_{p,n_m} for the case that $k = n_m \times m$, or zero otherwise as we are still keeping m fixed. Substituting this information into Equation (8), we obtain

$$P_p^{(m)}(t) = \sum_{n_m} t^{m n_m} \frac{x_{p,m}^{n_m} e^{-x_{p,m}}}{n_m!} = \exp[x_{p,m}(t^m - 1)], \quad (11)$$

where we only included the nonzero values in the sum. Now, this result was obtained for a fixed m , but in order to obtain the full non-Poissonian generating function, we need to account for all possible m . As each source is independent, each value of m is independent also. We can then make use of the fact that the generating function of a sum of independent random variables is given by the product of each variable's generating function. Accordingly, the full non-Poissonian generating function is given by

$$P_p(t) = \exp\left[\sum_{m=1}^{\infty} x_{p,m}(t^m - 1)\right]. \quad (12)$$

From here, using the inversion formula in Equation (9), the probability of observing k neutrinos in a pixel p is

$$p_{p,k} = \frac{1}{k!} \frac{d^k}{dt^k} \exp\left[\sum_{m=1}^{\infty} x_{p,m}(t^m - 1)\right] \Bigg|_{t=0}, \quad (13)$$

or combining this with our earlier results

$$p_{p,k}(\theta) = \frac{1}{k!} \frac{d^k}{dt^k} \exp\left[\sum_{m=1}^{\infty} (t^m - 1) \times \int_0^{\infty} dS \frac{1}{\mathcal{E}_p} \frac{dN_p}{dF}\left(\frac{S}{\mathcal{E}_p}; \theta\right) \frac{S^m e^{-S}}{m!}\right] \Bigg|_{t=0}. \quad (14)$$

This expression, combined with Equation (7), gives us our full NPTF likelihood as a function of the source-count function dN_p/dF , which is exactly what we want to constrain. Although this expression contains an unevaluated integral, in the case of a multiply broken power law, it can be efficiently implemented as the integral can be calculated analytically, and furthermore, the $p_{p,k}$ can be evaluated recursively in k . All of these details are described in Mishra-Sharma et al. (2017).

Below, we will outline how this likelihood can be extended to account for the presence of additional contributions to the data beyond just point sources. But before doing so, we turn to the practicalities of implementing the NPTF in the specific case of IceCube.

3.2. Implementation at IceCube

There are two immediate obstacles to implementing the NPTF at IceCube, beyond the basic outline described in the previous subsection. First, we need to calculate an appropriate effective area matrix in order to determine \mathcal{E}_p , which appears in the conversion between the dN/dF and dN/dS , as per Equation (5). Second, we must incorporate the real IceCube PSF and thereby remove the assumption hidden in Equation (6), that a source deposits all of its neutrinos in the

pixel it is located. For some of the neutrinos recorded at IceCube, there is a small but nonzero probability that their true incident direction is separated from the reconstructed value by a significant (angular) distance, so this effect must be accounted for in essentially any binning of the data. We will address each point in turn.

Consider first the effective area matrix, a quantity that specifies the response of the detectors to an incident neutrino of a given energy, R.A., and decl. This, of course, cannot be calculated analytically. Instead we take simulations of individual events and use these to construct \mathcal{E}_p . We do this by reweighting simulated events within our energy range according to an E^{-2} spectrum, as this is the spectrum that we assume our point-source population follows, see Equation (4) and the surrounding discussion. This then provides the average detector response as a function of R.A. and decl., which we can map to galactic coordinates and provide \mathcal{E}_p . As claimed above, this map can vary significantly, by more than an order of magnitude between locations with highest and lowest effective area. Given IceCube's location at the geographic South Pole, this variation is exclusively in decl., which tracks whether the event is arriving above from through the atmosphere, or below through the Earth. In light of this we used a relatively large number of 50 exposure regions to convert from flux to counts.

Turning next to the PSF, recall as discussed already that the NPTF likelihood derived in the last subsection assumed perfect angular reconstruction of every event. This assumption was invoked in writing down the number of sources contributing m neutrinos in a pixel p , denoted $x_{p,m}$, in Equation (6). By moving directly from the expected number of sources in the pixel to the expected number of neutrino counts, implicit is the assumption that the source deposits all of its flux into that pixel. Yet detector effects will smear the flux of a real source among a number of pixels. As in the NPTF, we do not keep track of which pixels are adjacent, what we want is the distribution for how a given source deposits its flux among the pixels on the map, a quantity denoted $\rho(f)$. Here, $f \in [0, 1]$ is the fraction of the point source's flux; the case of near perfect angular reconstruction, as compared to the pixel size, corresponds to a $\rho(f)$ peaked near $f = 1$, as most of the flux tends to be distributed in one pixel (the pixel where the source is located). Indeed in the limit of exact angular reconstruction, we have $\rho(f) = 2\delta(f-1)$.⁶⁷ More generically, however, imperfect angular reconstruction leads to a distribution peaked nearer $f = 0$ as most often a pixel will only get a small fraction of the flux. As a concrete example, consider a 3×3 grid with a point source at the center. Imagine the source deposits 60% of its flux in the central pixel that it inhabits, and then 5% in each of the eight pixels surrounding it. In this case, many more pixels experience a small amount of flux, and so $\rho(f)$ would still be peaked toward smaller values of f . Further note that $\rho(f)$ itself is not a probability density function, instead as the point source must deposit all of its flux somewhere, the distribution is normalized so that

$$\int_0^1 df f \rho(f) = 1. \quad (15)$$

⁶⁷ Note that in this idealized case, $\rho(f)$ will also have a contribution at $f = 0$ as for perfect angular reconstruction, most of the sky will receive no flux. Nevertheless, when $f = 0$ there is no contribution to the neutrino flux, and so in practice, we will always neglect the zero flux case.

Imagine we have the appropriate $\rho(f)$ —we will describe how to derive this shortly—consider how this modifies $x_{p,m}$. Previously, we used the fact that dN_p/dS provides the number of sources that contribute an expected number of counts S , to then reweight this quantity by the probability of fluctuating from S to m , given by the Poisson distribution $S^m e^{-S}/m!$. Now, however, a source is only expected to deposit a fraction f of that flux in the pixel under consideration, and so instead, the reweighting to obtain m neutrinos is $(fS)^m e^{-fS}/m!$. Furthermore, the probability that a given value of f is chosen is dictated by the distribution $\rho(f)$, and integrating over all possible flux fractions, we arrive at the following modification for $x_{p,m}$:

$$x_{p,m}(\theta) = \int_0^\infty dS \frac{dN_p}{dS}(S; \theta) \times \int_0^1 df \rho(f) \frac{(fS)^m e^{-fS}}{m!}. \quad (16)$$

Another way to understand the above modification is to compare this result to Equation (6). In doing so, the modification induced by the finite PSF is seen to be equivalent to substituting in a modified source-count function,

$$\frac{dN_p}{dS}(S) \rightarrow \frac{d\tilde{N}_p}{dS}(S) = \int_0^1 df \frac{\rho(f)}{f} \frac{dN_p}{dS}(S/f). \quad (17)$$

Propagating the modification in Equation (16) through to the NPTF likelihood then gives a full accounting for the effect of the finite PSF, once we have $\rho(f)$. Note that taking $\rho(f) = 2\delta(f-1)$, i.e., perfect angular reconstruction, the above expression reduces to Equation (6), as it must, and this can also be seen clearly in Equation (17).

All that remains then in order to incorporate the IceCube PSF is an algorithm for determining $\rho(f)$. Commonly, the instrument PSF is stated as a probability distribution for a given event to be located at some radius from the center of a source. The median reconstruction angle for $\nu_\mu + \bar{\nu}_\mu$ events can be seen, for example, in Figure 2 of Aartsen et al. (2017a), where the angular error can vary from half a degree to several degrees depending on the energy and event type. Nevertheless, this is just the median reconstruction angle; the tails of this distribution are considerably non-Gaussian and can extend out to very wide angles in the case of poorly reconstructed events. Modeling the tails correctly is critical for an NPTF analysis. As an example of why this is important, if there is a true population of sources all with identical fluxes, a mismodeled PSF can lead to the fit preferring an additional population of lower flux sources associated with mis-reconstructed neutrinos.

The event-by-event determination of the angular reconstruction can be exploited in an unbinned analysis, as done for example in Aartsen et al. (2017a). However, as the NPTF is fundamentally a binned method, we will instead consider quantities averaged within the data set of interest. There is no simple analytic expression for converting from the known PSF to $\rho(f)$, as, for example—unlike the PSF— $\rho(f)$ depends critically on the binning of the underlying map.⁶⁸ We can, however, determine this distribution using the following algorithmic prescription:

1. Simulate N equal-flux point sources on a blank map with the same pixelization as the NPTF will be applied to;

⁶⁸ To highlight this, note that in the limit where the map contains only one pixel, we must have $\rho(f) = 2\delta(f-1)$, independent of the PSF.

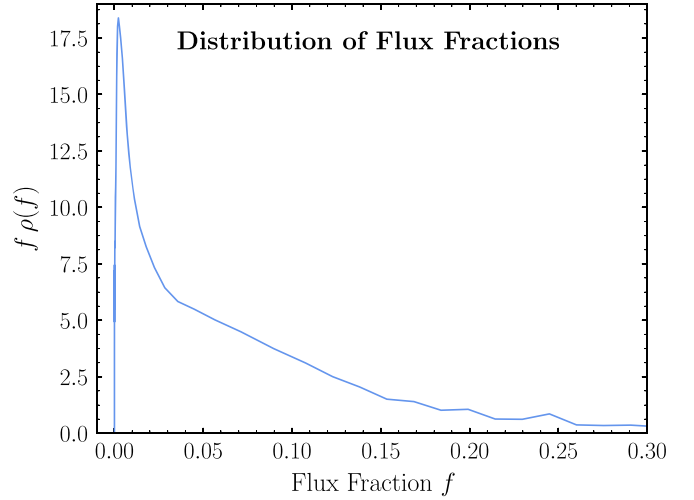


Figure 1. Distribution of the frequency of pixels that contain a fraction f of the flux from a point source, $\rho(f)$, for the full sky template, which is appropriate for modeling isotropic extragalactic sources. This quantity is central to incorporating the PSF of IceCube into the NPTF likelihood, according to Equation (16). We have chosen to show $f\rho(f)$, as following Equation (15) this quantity integrates to one given conservation of flux. See the text for details.

2. Determine the fraction of the total flux in each pixel, f_p , defined such that $\sum_p f_p = 1$;
3. Define a flux binning Δf , and as a function of flux f , define $\Delta n(f)$ the number of pixels that have a flux between f and $f + \Delta f$; and
4. Combine these quantities to define $\rho(f)$ as follows:

$$\rho(f) \approx \frac{\Delta n(f)}{N\Delta f}. \quad (18)$$

The above relation is approximate and only becomes exact in the limit $N \rightarrow \infty$ and $\Delta f \rightarrow 0$.

In order to simulate this in practice, we deposit a large number of sources on the sky, and for each one, we model the neutrino distribution according to the appropriate PSF at each location. To account for the energy dependence inherent in the PSF of IceCube, following Equation (4), we assume each source has an E^{-2} spectrum and draw events for the source according to this distribution. In this way, we can exactly build up $\rho(f)$ as defined in Equation (18), and the result is shown in Figure 1 for our default full-sky analysis. In that figure, we have zoomed in on the small f values where the distribution is peaked. That the flux fraction distribution is peaked at small values is indicative of the fact that the IceCube PSF has tails that extend significantly more broadly than the size of a pixel on the map but is also quite generic of $\rho(f)$ unless the angular reconstruction is significantly better than the pixel size. For comparison, note that the linear size of our pixels is $\sim 0.92^\circ$, which is comparable to the median reconstruction angle of our events, which is $\sim 1^\circ$. As the PSF of IceCube varies across the sky, $\rho(f)$ must be determined for each of our spatial distributions separately.⁶⁹

In light of these results, the NPTF likelihood can now be applied to the IceCube data set. As an important validation of

⁶⁹ In principle, due to the spatial variation of the PSF, $\rho(f)$ also varies spatially. The approximation made in this work is that we will use the mean of $\rho(f)$ across the sky, as weighted by our various spatial templates.

the method, we performed extensive tests based upon Monte Carlo studies where we injected and then recovered point-source populations. An example of this is shown in Figure 6, discussed in detail below. We emphasize again that, in particular, the details discussed in this section were critical to the successful recovery of injected populations. With the method validated, we can now look to calculate the expected sensitivity at IceCube, which we turn to in Section 4. Before doing so, however, in the next two subsections, we will discuss how we incorporate backgrounds that are not associated with point sources into our model and then how these various likelihoods will be combined into an inference framework we can use to test for the presence of point sources.

3.3. Adding Poissonian Models

At the very least, due to the presence of irreducible backgrounds, we know that point sources cannot be the only contribution to the IceCube data set, and in this section, we discuss how to augment the NPTF likelihood to account for these. These additional contributions are generally expected to be described by the Poisson distribution, following an underlying spatial map. To incorporate both the Poissonian statistics and the spatial variation, we will use the language of Poissonian templates, following the language from a recent application of this topic in Lisanti et al. (2018).

To begin with, as in the NPTF case, we imagine that our model follows a spatial distribution that once pixelized can be described by a map T_p . Unlike for an NPTF model, where T_p specified the spatial distribution of point sources, in the Poisson case, we require T_p to be proportional to the expected distribution of counts, not flux. As a concrete example, in the case where our model is for isotropic extragalactic neutrino emission, the appropriate T_p would still have a spatial variation inherited from the effective area matrix described in the previous subsection. As in the non-Poissonian case, we assume for a given model that the Poissonian template, T_p , is specified ahead of time. What we fit for in this case is the overall normalization of this template, in terms of which our Poissonian model prediction is given by

$$\mu_p(\boldsymbol{\theta}) = A T_p, \quad (19)$$

so that $\boldsymbol{\theta} = \{A\}$, where we emphasize that A has no pixel dependence. Given that the sum of two Poisson distributions with means μ_1 and μ_2 is again a Poisson distribution of mean $\mu_1 + \mu_2$, we can readily extend this formalism to account for multiple Poisson models. For example, if we had n of these, described by template T_p^1, \dots, T_p^n , then our combined model prediction in each pixel would be

$$\mu_p(\boldsymbol{\theta}) = \sum_{\ell=1}^n A^\ell T_p^\ell, \quad (20)$$

where now $\boldsymbol{\theta} = \{A^1, \dots, A^n\}$. To provide a concrete example, we may want to model the observed flux using a model that combines three sources: 1. terrestrial backgrounds such as atmospheric neutrinos; 2. diffuse extragalactic emission; and 3. diffuse emission from the Milky Way. In such a scenario, we would have three Poissonian templates, and for each of these, T_p^ℓ would describe the pixel dependence of the flux, and A^ℓ the overall normalization. For the case of extragalactic emission, as

the flux is expected to be isotropic, T_p^ℓ would then be a map of the IceCube detector's response to a uniform incident flux.

In terms of this, if all we had was a set of Poissonian models, then we could write down our likelihood according to the Poisson distribution,

$$P_{p,k}(\boldsymbol{\theta}) = \frac{\mu_p(\boldsymbol{\theta})^k e^{-\mu_p(\boldsymbol{\theta})}}{k!}. \quad (21)$$

Nonetheless, we want to construct a likelihood incorporating both Poissonian and non-Poissonian models. For this purpose, we can use the property of generating functions we exploited earlier, specifically that the generating function that describes the sum of two independent random variables is given by the product of the generating functions for the individual variables. For the Poisson case, given in Equation (21), the associated generating function according to Equation (8), is given by

$$P_p(t; \boldsymbol{\theta}) = \exp[\mu_p(\boldsymbol{\theta})(t - 1)]. \quad (22)$$

Combined with the generating function for the non-Poissonian case given in Equation (12), we arrive at

$$P_p(t; \boldsymbol{\theta}) = \exp\left[\sum_{m=1}^{\infty} x_{p,m}(\boldsymbol{\theta})(t^m - 1) + \mu_p(\boldsymbol{\theta})(t - 1)\right], \quad (23)$$

where now $\boldsymbol{\theta} = \{A, F_b, n_1, n_2, A^1, \dots, A^n\}$. Through the use of Equation (9), this generating function can be used to derive a combined likelihood that includes both Poissonian and non-Poissonian models, accomplishing one of the main aims of this subsection.

The main application for the Poissonian template formalism in our work will be to model the known backgrounds arising from atmospheric neutrinos and muons. For this purpose, we need to derive an appropriate T_p describing the spatial distribution of these contributions. Determining this from first principles is at present out of reach. Fortunately, however, we can estimate the distribution from the data alone. The reason for this is if we assume the data are made up of predominantly background events and a subset of point sources, then we can remove the point sources in the following way. Given the approximate azimuthal symmetry of the effective area of detector, we can take the data collected by IceCube and scramble the events by assigning them a random R.A. value. This process removes any point-source hotspots, as they will be smeared out along bands of constant decl. Furthermore, as the background is only expected to vary with decl., this process does not degrade the spatial information pertaining to the background process. Applying this process once gives a map that is still as noisy as the data. In order to extract a more appropriate map for the mean of a Poisson distribution, we repeated this scrambling process a large number of times and take the average of the resulting maps. Finally, to remove the noise in decl., we convolve this model with a von MisesFisher distribution that has a concentration corresponding to $1^\circ 08$, chosen as this is the median angular resolution at ~ 1 TeV. This last step can be justified as the real data have been scrambled on such a scale due to the PSF.

The map resulting from this procedure is shown in Figure 2, which is the Mollweide projection of the map in galactic coordinates. The most apparent feature in this map is the strong variation away from the poles toward the equator where the




Figure 2. Data-driven background template for the spatial distribution of atmospheric neutrinos and muons, derived using the procedure described in the text. This map is referred to as T_p^{bkg} . The map is a Mollweide projection of an underlying distribution in galactic coordinates, and the overall normalization is arbitrary.

largest flux is observed. Taking this map as template, we can readily incorporate the largest expected background into our likelihood. In subsequent discussions, we will refer to this map as T_p^{bkg} .

3.4. Inference Framework

So far in this section, we have introduced the formalism required to calculate the likelihood for a data set in the presence of a population of sources and additional Poissonian contributions. Our goal now is to put this machinery to work in the form of a test statistic that we can use to test for the presence of point sources in the data. Our test statistic will compare between two hypotheses: that point sources, distributed according to a spatial template T_p^{PS} , are present in the data or they are not. We will refer to these as the non-Poissonian and Poissonian hypotheses, respectively.

In detail, the Poissonian hypothesis is a model consistent of two Poisson templates: one following the dominant background contribution, given by T_p^{bkg} , and the other accounting for the possibility that for a given spatial distribution T_p^{PS} , the data may have a diffuse rather than an unresolved point-source origin. For this second source of emission, as the flux is diffuse, it is better described by Poisson statistics and, thus, follows a template $\mathcal{E}_p T_p^{\text{PS}}$, where the extra factor of \mathcal{E}_p is required to convert to a counts map. In this case, we can write down a likelihood function as described above, and from the data, we can construct the marginal likelihood as follows:

$$\mathcal{L}_0(d) = \int d\theta \mathcal{L}_P(d|\theta)p(\theta), \quad (24)$$

where the subscript 0 indicates that we are using this as our null hypothesis, and the subscript P on the likelihood identifies this as the appropriate form for the Poissonian hypothesis. In detail, $\mathcal{L}_P(d|\theta)$ can be determined directly from Equation (21), as we are only considering Poissonian models. In addition, we have introduced $p(\theta)$, which represents the priors on the parameters. These are given in Table 1, where A^{bkg} and A^P are the normalizations of the templates T_p^{bkg} and $\mathcal{E}_p T_p^{\text{PS}}$, respectively, and so we see this is a two-parameter model.

The non-Poissonian hypothesis is derived from the null hypothesis, except that we append one further model: a non-Poissonian template following the spatial template T_p^{PS} . As we will take a singly broken power law to describe the source-count distribution, this hypothesis is a six-parameter model.

Table 1
List of Priors

Parameter	Prior Range
$\log_{10}(A)$ [TeV cm ² s]	[2.41, 16.41]
$\log_{10}(F_b)$ [TeV ⁻¹ cm ⁻² s ⁻¹]	[-18, -7]
n_1	[2, 10]
n_2	[-6, 2]
A^{bkg}	[0.5, 1.5]
A^P	[-0.5, 1]

Note. All priors are taken to be flat in the given ranges.

Once more, we can form the marginal likelihood using

$$\mathcal{L}_1(d) = \int d\theta \mathcal{L}_{\text{NP}}(d|\theta)p(\theta), \quad (25)$$

where again the priors are given in Table 1. Note that all priors are uniform, except for A and F_b , which are log uniform. Also, the prior for A^P is allowed to float negative, in order to allow the fit to conserve the total amount of flux when evaluating the non-Poissonian hypothesis. To justify this choice, recall that T_p^{bkg} was constructed by scrambling the real data. Accordingly, if there is a detectable point-source population within the data, the flux from these sources would be picked up by the non-Poissonian template but also be present in T_p^{bkg} and, therefore, double counted. A negative A^P can then be loosely thought of as subtracting that flux off, but done in such a way that the combined Poissonian template $A^{\text{bkg}} T_p^{\text{bkg}} + A^P \mathcal{E}_p T_p^{\text{PS}} \geq 0$ in every pixel. We emphasize that the model used for non-Poissonian hypothesis includes both Poissonian and non-Poissonian templates, and thus, the full generating function in Equation (23) is required.

Model selection between these two hypotheses is considered through the use of the Bayes factor

$$\mathcal{B}_{\text{NP/P}}(d) = \frac{\mathcal{L}_1(d)}{\mathcal{L}_0(d)} = \frac{\int d\theta \mathcal{L}_{\text{NP}}(d|\theta)p(\theta)}{\int d\theta \mathcal{L}_P(d|\theta)p(\theta)}. \quad (26)$$

From this definition, it can be seen that the Bayes factor is a summary statistic: it integrates over all possible forms of the source-count function through the integral over parameters. As such, it provides a gross evaluation as to whether a point-source population is preferred by the data, rather than singling out any particular dN/dF . As a summary statistic, it can also serve a secondary role as a test statistic. Our expectation for the value of the Bayes factor can be calibrated through use of frequentist methods such as calculating the p -value of the Bayes factor.

To compare more specific model hypotheses, we introduce the *pointwise likelihood ratio*, defined as

$$\mathcal{M}(d; \phi) = \frac{\tilde{\mathcal{L}}_{\text{NP}}(d|\phi)\pi_{\text{NP}}}{\mathcal{L}_0(d)\pi_P + \mathcal{L}_1(d)\pi_{\text{NP}}}. \quad (27)$$

where π_P and π_{NP} are, respectively, the model priors for the Poissonian and non-Poissonian hypotheses, which have been chosen to be equal for this presentation of the results. Here, $\phi = \{N^{\text{PS}}, \bar{F}^{\text{PS}}\}$ represents the expected number of point sources across the whole sky and the expected flux per source at this location in model space, both of which were defined in Equation (3). In this expression, $\tilde{\mathcal{L}}_{\text{NP}}(d|\phi)$ is similar to $\mathcal{L}_{\text{NP}}(d|\theta)$, except that n_1 and n_2 have been marginalized over.




Figure 3. The three forms of point-source distribution considered in this work, in addition to a purely isotropic distribution. From left to right, these are a model for the Galactic disk, the Fermi bubbles (Su et al. 2010), and the SFD dust map (Schlegel et al. 1998). In effect, these are maps of T_p^{PS} in galactic coordinates, and the normalizations are arbitrary. All maps are on a linear scale except for the SFD dust map, where we use a log color axis to emphasize the more detailed structure within this map. See the text for further description.

Intuitively, $\mathcal{M}(d; \phi)$ should be thought of as the probability for the NPTF model at this particular value of $\{N^{\text{PS}}, \bar{F}^{\text{PS}}\}$, compared to the probability for an equally weighted mixture of the Poissonian and non-Poissonian models. This definition was motivated by the need for a metric that handles both the high and low signal strength regimes. When $\mathcal{B}_{\text{NP/P}} \ll 1$, \mathcal{M} is approximately the model posterior conditioned on ϕ ; while in the $\mathcal{B}_{\text{NP/P}} \gg 1$ regime, \mathcal{M} is approximately the ratio between the posterior and prior: $p(\phi|d)/p(\phi)$.

Before concluding this section, we have said a number of times that we will exploit the power of the NPTF to test several forms for the spatial dependence of the point-source population, specified by T_p^{PS} . In addition to isotropically distributed sources, where $T_p^{\text{PS}} \propto 1$, we will consider three additional templates that consider the possibility of point sources distributed within the Milky Way, all of which are shown in Figure 3. The galactic disk template is used as a generic model for sources distributed according to the disk of the Milky Way. This is the line-of-sight integrated emission from a disk with a source density that scales exponentially in radius and distance from the plane, with a scale height of 0.3 kpc and a scale radius of 5 kpc. Next, we consider sources distributed following the Fermi bubbles (Su et al. 2010), large structures observed in gamma-rays extending perpendicularly from the galactic disk. Although the emission observed from the bubbles so far appears diffuse, neutrino emission from the recently discovered small scale gas clouds within the bubbles (Di Teodoro et al. 2018), could lead to point like sources. Finally, we consider the Schlegel, Finkbeiner, and Davis (SFD) dust map (Schlegel et al. 1998), which provides a two-dimensional distribution of dust within the Milky Way, mapped out using the reddening of starlight. This distribution is interesting to consider, as dust and gas tend to have similar spatial distributions, so the map is a proxy for the distribution of hydrogen in the Milky Way. The hydrogen is a target for cosmic-ray proton to collide with and form pions. The neutral pions then decay to photons, and indeed, the SFD dust map can be seen clearly in the Fermi gamma-ray data. If these same interactions produce higher-energy pions, then the charged variants could produce neutrinos at IceCube, making this an interesting spatial distribution to consider. As the target dust has a diffuse distribution throughout the Milky Way, in order for this template to describe a point-source population, the sources of the cosmic-ray protons would need to be point like, as this would then imprint a point-source-like distribution into the




Figure 4. Distribution of the test statistic, as given in Equation (26), under the background-only hypothesis for four different signal states. These distributions, formed from 1000 trials, are used for establishing sensitivities and p -values. In particular, the vertical lines represent a p -value of 0.5. Here, each trial is formed by scrambling the data in R.A. See the text for details.

neutrino data. Note that we have chosen not to divide our templates between the northern and southern sky, even though this is commonly done for extragalactic point-source searches (see, for example, Aartsen et al. 2017a). Usually, this distinction between the northern and southern sky is imposed due to the different backgrounds that dominate in each hemisphere; in the northern sky, the main background is atmospheric neutrinos, whereas in the southern sky, instead, atmospheric muons dominate. Nonetheless, for the isotropic case, we show both the sensitivity and results for the northern sky in Appendix B. There we will see that the reach for the restricted case is only slightly enhanced, justifying our choice to focus on the full sky.

4. Expected Sensitivity

Using the techniques and statistical framework described in the previous section, we now turn to estimating the expected reach of this technique using Monte Carlo simulations. We will consider both the case of setting limits and quantifying thresholds for a discovery of a point-source population. As we use techniques generally motivated by Bayesian statistics, part of the aim of this section is to help develop intuition for

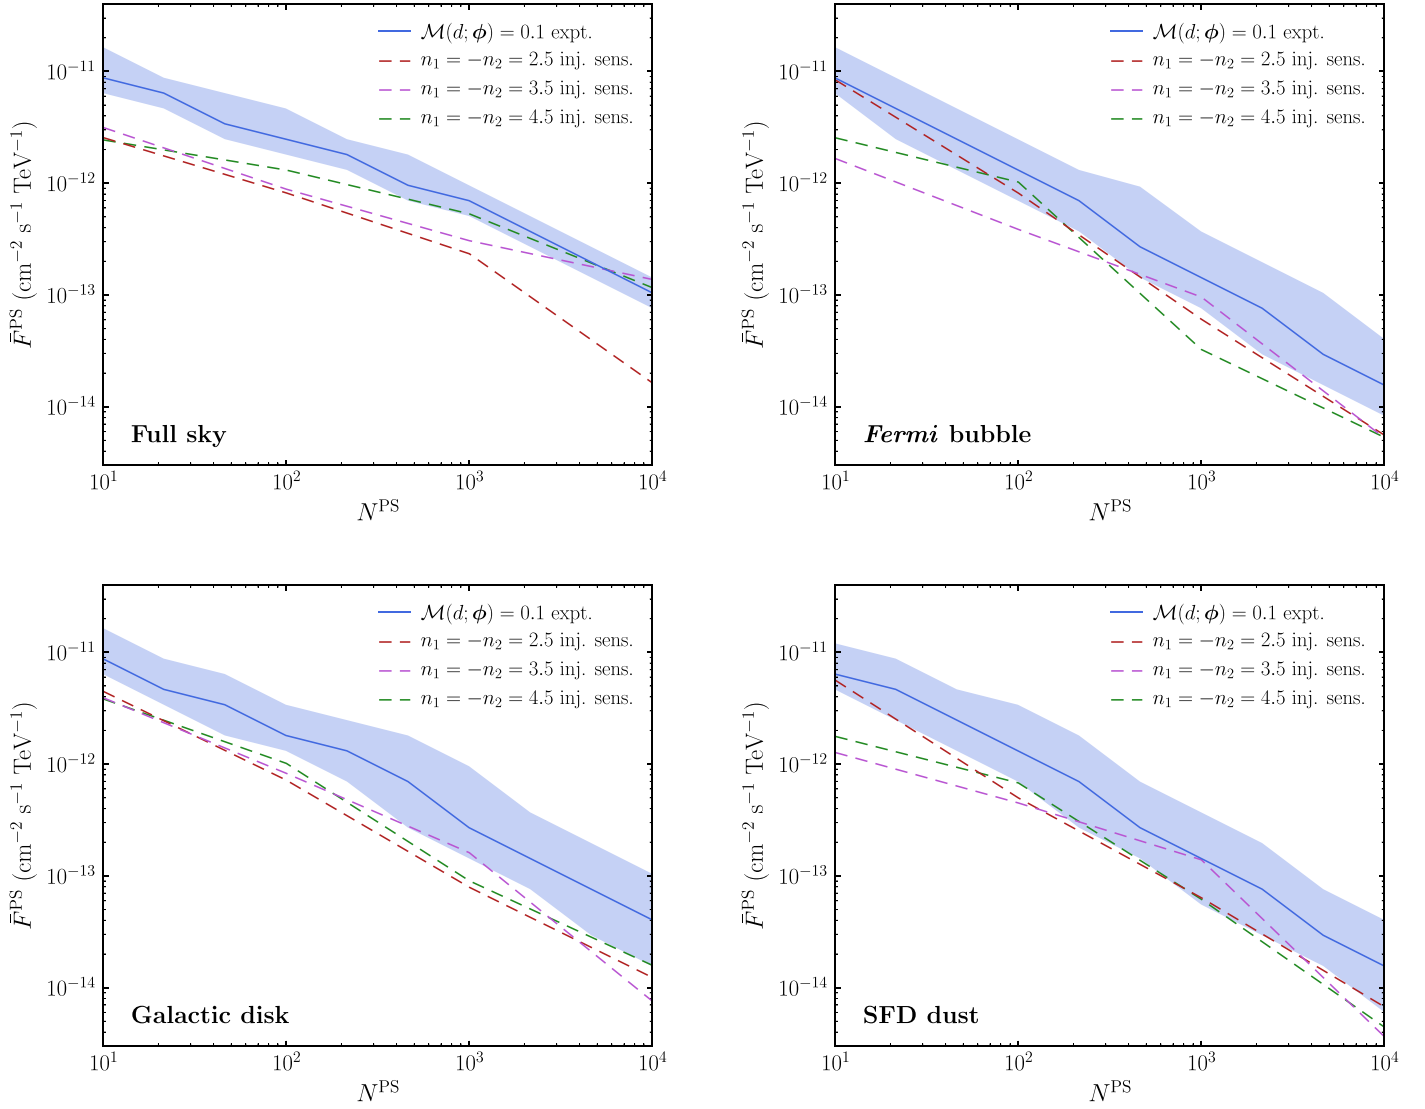


Figure 5. The expected sensitivity and limit for the four different spatial templates considered for the point-source distributions: isotropic extragalactic sources over the full sky (top left panel), Fermi bubbles (top right panel), SFD dust (bottom right panel), and galactic disk (bottom left panel). In each case, the sensitivity is shown as the dashed curves for three different shapes of the source-count function. The median expected limit derived using $\mathcal{M}(d; \phi) = 0.1$ is shown in blue, as well as the associated 10th and 90th percentiles from the distribution. See the text for details.

what discovery and limit setting looks like in our framework. Nevertheless, the primary output of a Bayesian analysis is the posterior, and as mentioned, we make this publicly available,⁷⁰ describing the details in Appendix A.

To begin with, we consider the expected limit sensitivity for the analysis. The sensitivity is determined by comparing how our test statistic, the Bayes factor given in Equation (26), is distributed over many trials for each of the signal plus background and background-only hypothesis. From these distributions, our sensitivity to a given model is defined as when 90% of the signal distribution is above 50% of the background case. For example, the background-only distribution of the natural log of our test statistic, generated from 1000 trials, is shown in Figure 4 for each of our four signal templates. The trials are generated by taking the real data but scrambling the R.A. of each event, with a different scrambling for each trial. This background-only distribution can also be

used to establish p -values, and indeed, the definition of sensitivity is equivalent to requiring the p -value for 90% of the signal distribution to be less than 0.5.⁷¹ Note that as sensitivity is defined in terms of the distributions, there is no statistical variation in its value, as opposed to say a frequentist 90% confidence limit. Where the sensitivity threshold occurs will not be a unique point in the signal-parameter space, established by $\{N^{\text{PS}}, \bar{F}^{\text{PS}}, n_1, n_2\}$. If, however, we fix three of the parameters, for example N^{PS}, n_1 , and n_2 , then we can define the sensitivity as a function of \bar{F}^{PS} . With this in mind, in Figure 5, we show the sensitivity to \bar{F}^{PS} as a function of N^{PS} ,

⁷⁰ https://icecube.wisc.edu/science/data/NPTF_7yr_posterior

⁷¹ A similar procedure can be used to establish the expected discovery sensitivity, which is defined as when 50% of the signal distribution has a p -value less than 2.87×10^{-7} , the threshold usually referred to as a 5σ discovery. Determining the associated test statistic that corresponds to such a small p -value requires generating a large number of background-only trials, and as we find no significant evidence for a signal in the present analysis, we have not quantified the discovery threshold.

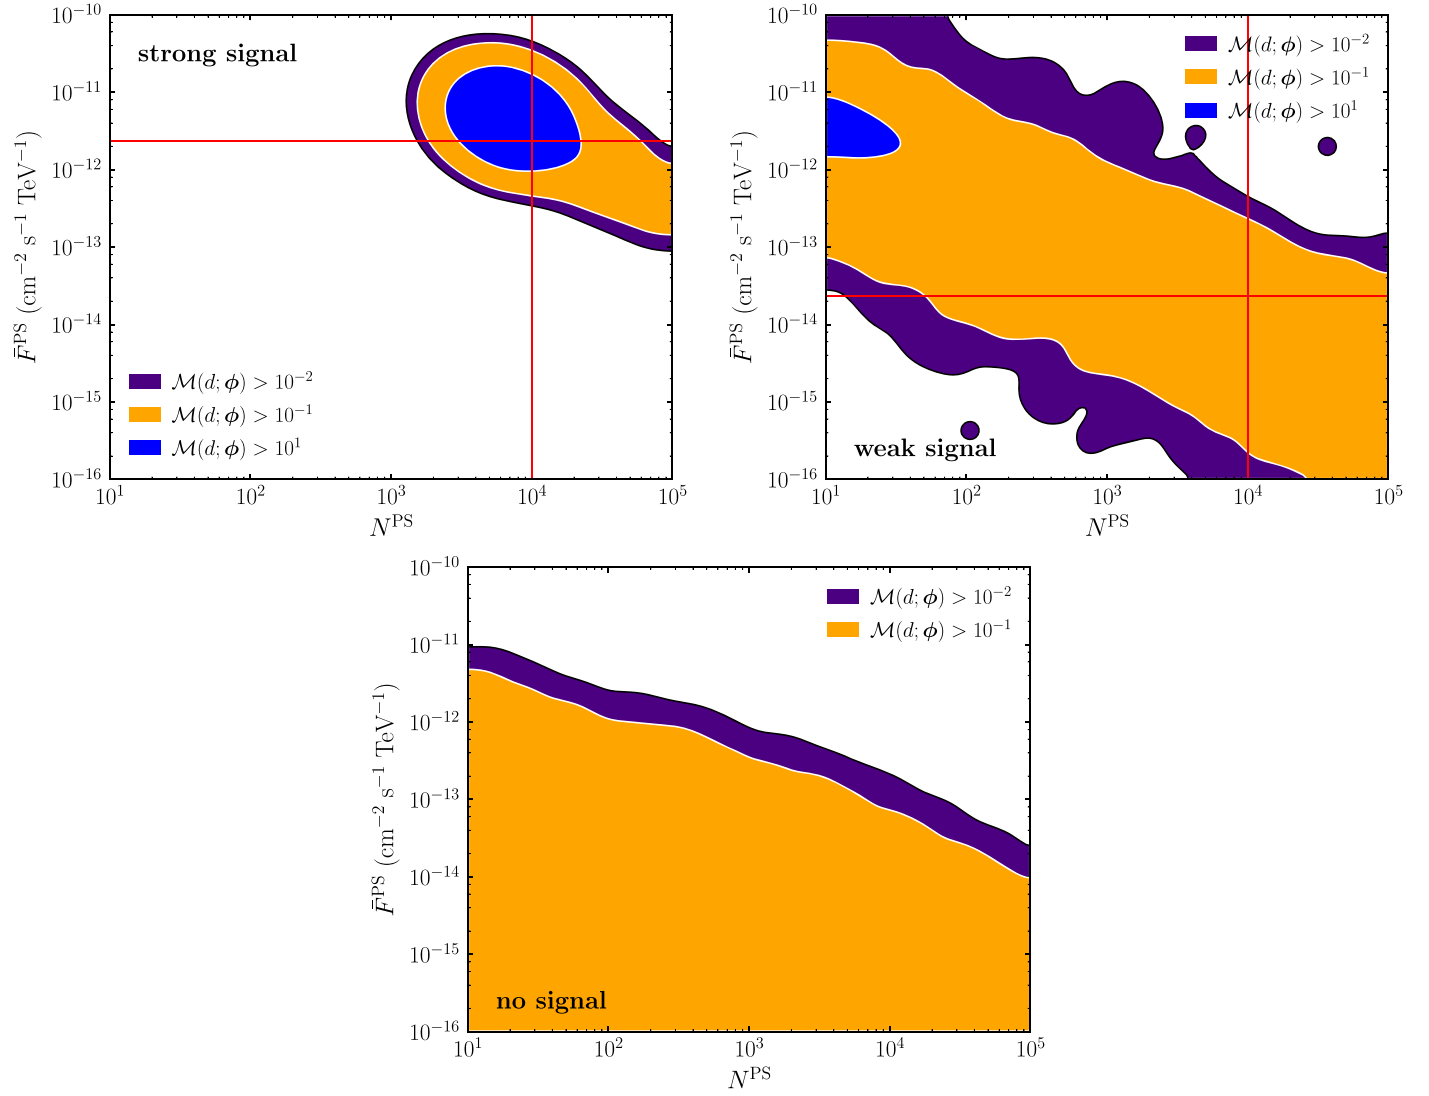


Figure 6. Map of $\mathcal{M}(d; \phi)$ for three different data sets including injected point-source populations. In each of the cases, the parameters of the injected population are shown with the red lines. Only in the case of a strong signal is this method able to identify the specific location in this reduced parameter space of the signal. See the text for details.

for three different values of $n_1 = -n_2$, and for the four signal templates.

In addition to the sensitivity, we define a contour where $\mathcal{M}(d; \phi)$, given in Equation (27), equals 0.1. This contour is the dividing line, above which the odds for a particular point in parameter space are no better than 1 in 10. As $\mathcal{M}(d; \phi) = 0.1$ varies between different data sets, we show the median, 10th, and 90th percentiles on the distribution also in Figure 5. Recall that $\mathcal{M}(d; \phi)$ is defined by marginalizing over n_1 and n_2 . As explained around Equation (3), we can describe an NPTF template in terms of N^{PS} , \bar{F}^{PS} , n_1 , and n_2 . With N^{PS} fixed, and the indices marginalized over, the only remaining degree of freedom is the average flux per source, \bar{F}^{PS} , which, for a fixed number of sources, is equivalent to the total flux associated with the non-Poissonian template. Accordingly, the limit set using this procedure effectively reduces to the weaker constraint obtained by ensuring that such a population does not overproduce the observed neutrino flux, rather than drawing on the full power of the NPTF likelihood. This

explains why, for a large number of sources, the expected limit is weaker for the full sky as compared with the other cases. For the spatially restricted templates, such as the Fermi bubbles, the point sources and the majority of their associated neutrinos are forced within a smaller region on the sky compared with the full sky case, leading to a stronger limit. As discussed earlier, all templates are applied to the full data set; although for the isotropic case, we show the sensitivity and results restricting to the northern hemisphere in Appendix B.

In addition to being useful for setting limits, we can use $\mathcal{M}(d; \phi)$ to map out the signal-parameter space to determine if there are regions that are particularly preferred by the data. In order to calibrate our analysis for this case, we will consider three different scenarios where we look for evidence of an actual point-source signal. In each of these scenarios, we consider a background data set with an additional injected point-source population. The population is distributed isotropically over the full sky, where in all scenarios, the expected number of sources is 10^4 , but each have a source-count




Figure 7. Expected sensitivities for four templates when $n_1 = -n_2 = 20$ is fixed, creating an approximation of an equal-flux population of sources. Northern and southern sky hotspot sensitivities for populations of equal-flux source from Aartsen et al. (2017a) are shown along-side. The hotspot analysis represents a traditional approach to point-source detection, where a “hotspot” is defined as a source that is individually detected at 3σ global significance. See the text for details.

function following a singly broken power law, with $n_1 = -n_2 = 2.5$. The cases are distinguished by the value of the flux break of the source-count function, F_b ; we consider a strong signal with $F_b = 10^{-12}$ neutrinos $\text{cm}^{-2} \text{s}^{-1} \text{TeV}^{-1}$, a weak signal of $F_b = 10^{-14}$ neutrinos $\text{cm}^{-2} \text{s}^{-1} \text{TeV}^{-1}$, and finally, a case almost equivalent to no signal with $F_b = 10^{-15}$ neutrinos $\text{cm}^{-2} \text{s}^{-1} \text{TeV}^{-1}$. Note the strong signal here corresponds to value that is a factor of ~ 100 brighter than the observed diffuse flux, and the sole purpose of such a large value is to validate that the framework has been calibrated correctly. The distribution of $\mathcal{M}(d; \phi)$ for each of these cases is shown in Figure 6, and we see that only in the strong signal case is the actual injected point clearly singled out. Nevertheless, in the weak-signal case, it is clear that a nonzero-point-source population is preferred, but the exact location is not correctly identified. The fit is not able to distinguish between a few bright or many dim sources. Finally, in the no-signal case, the data set is clearly consistent with no point-source population, as the injected population falls below the sensitivity of our analysis. For reference, the $\ln \mathcal{B}_{\text{NP/P}}$ of the strong, weak, and no-signal cases, is approximately 2.3×10^4 , 10.7, and -0.84 , respectively.

It is worth emphasizing that all results shown above, and indeed those we derive on the actual data in the next section, represent slices through the full model space. This is the cost of reducing a four-dimensional parameter space into a two-dimensional limit plot. For a specific model prediction, the full posterior is the more relevant resource.

Finally, we emphasize that a direct comparison of this method to limits set by previous IceCube analyses—such as those in Aartsen et al. (2017a)—cannot be made straightforwardly, as those limits are calculated by assuming a population of equal-flux point sources. This kind of population can be emulated by requiring n_1 and n_2 to be fixed to a large absolute value, thus creating an approximate Dirac delta distribution differential source-count function. In this limited region of

parameter space, a direct comparison is possible, and sensitivities for the four templates under consideration, using $n_1 = -n_2 = 20$, are shown in Figure 7 along side the northern and southern sky sensitivities from Aartsen et al. (2017a).⁷²

5. Results

In this section, we apply the techniques used to estimate the sensitivity discussed in the last section to the actual data. In particular, we plot the distribution of values $\mathcal{M}(d; \phi)$ for each of our signal templates. The results of this are shown in Figure 8. These plots indicate that for each of the investigated cases, there is no indication of a point-source population present in the data, and accordingly, the results are consistent with the expected limits shown in Figure 5. In each case, the significance can be quantified as follows:

1. *Full sky*: $\ln \mathcal{B}_{\text{NP/P}} = -0.79$, p -value = 0.66;
2. *Fermi Bubbles*: $\ln \mathcal{B}_{\text{NP/P}} = -0.94$, p -value = 0.45;
3. *SFD Dust*: $\ln \mathcal{B}_{\text{NP/P}} = -0.92$, p -value = 0.33; and
4. *Galactic Disk*: $\ln \mathcal{B}_{\text{NP/P}} = -0.97$, p -value = 0.74.

The p -values quoted here were determined from the distribution of the background-only hypotheses, which were shown in Figure 4. From the p -values, each signal template is consistent with the Poissonian hypothesis.

We can also consider the full posterior. In Figure 9, we show a triangle plot generated from the posterior for the case of the SFD dust signal template. The signal parameters are clearly consistent with a background-only hypothesis, and we note that the triangle plots for other templates are similar. The posterior for each template is made publicly available, and we refer to Appendix A for details.

The purpose of the public posterior is that they can be used to test any point-source population model where the associated dN/dF can be approximated by a broken power law. There are a wide number of source classes that have been considered as possible contributors to the IceCube neutrino flux. For an overview, see, for example, Murase & Waxman (2016). A fundamental problem, however, is that in many cases, there remains considerable uncertainty in the associated luminosity function. While we often have measurements of the photon luminosity function in the infrared, X-ray, or γ -ray energies, mapping from this to the neutrino luminosity function involves a number of assumptions. For an example in the case of blazars, see Yuan et al. (2019). For these reasons, the model space associated with neutrino sources is significant.

One approach to simplifying this space is to consider standard candles. Under this approach, the luminosity function is chosen to be sharply peaked at a certain value, denoted L_{eff} , and then the problem is reduced to scanning a two-dimension space parameterized by the effective luminosity, and the density of sources, denoted ρ_0 . More quantitatively, following Aartsen et al. (2019), the luminosity function of a standard candle is defined as a log-normal distribution with median L_{eff} and a width of 0.01 in $\log_{10} L_{\text{eff}}$. This model is then converted to an associated dN/dF using FIRESONG Taboada et al. (2018), adopting a density evolution for the source population according to the evolution of the star formation rate in Hopkins & Beacom (2006), and a flat universe with $\Omega_{M,0} = 0.308$,

⁷² The northern sky sensitivity from Aartsen et al. (2017a) has been recalculated to account for an incorrect treatment of signal acceptance in the original publication.

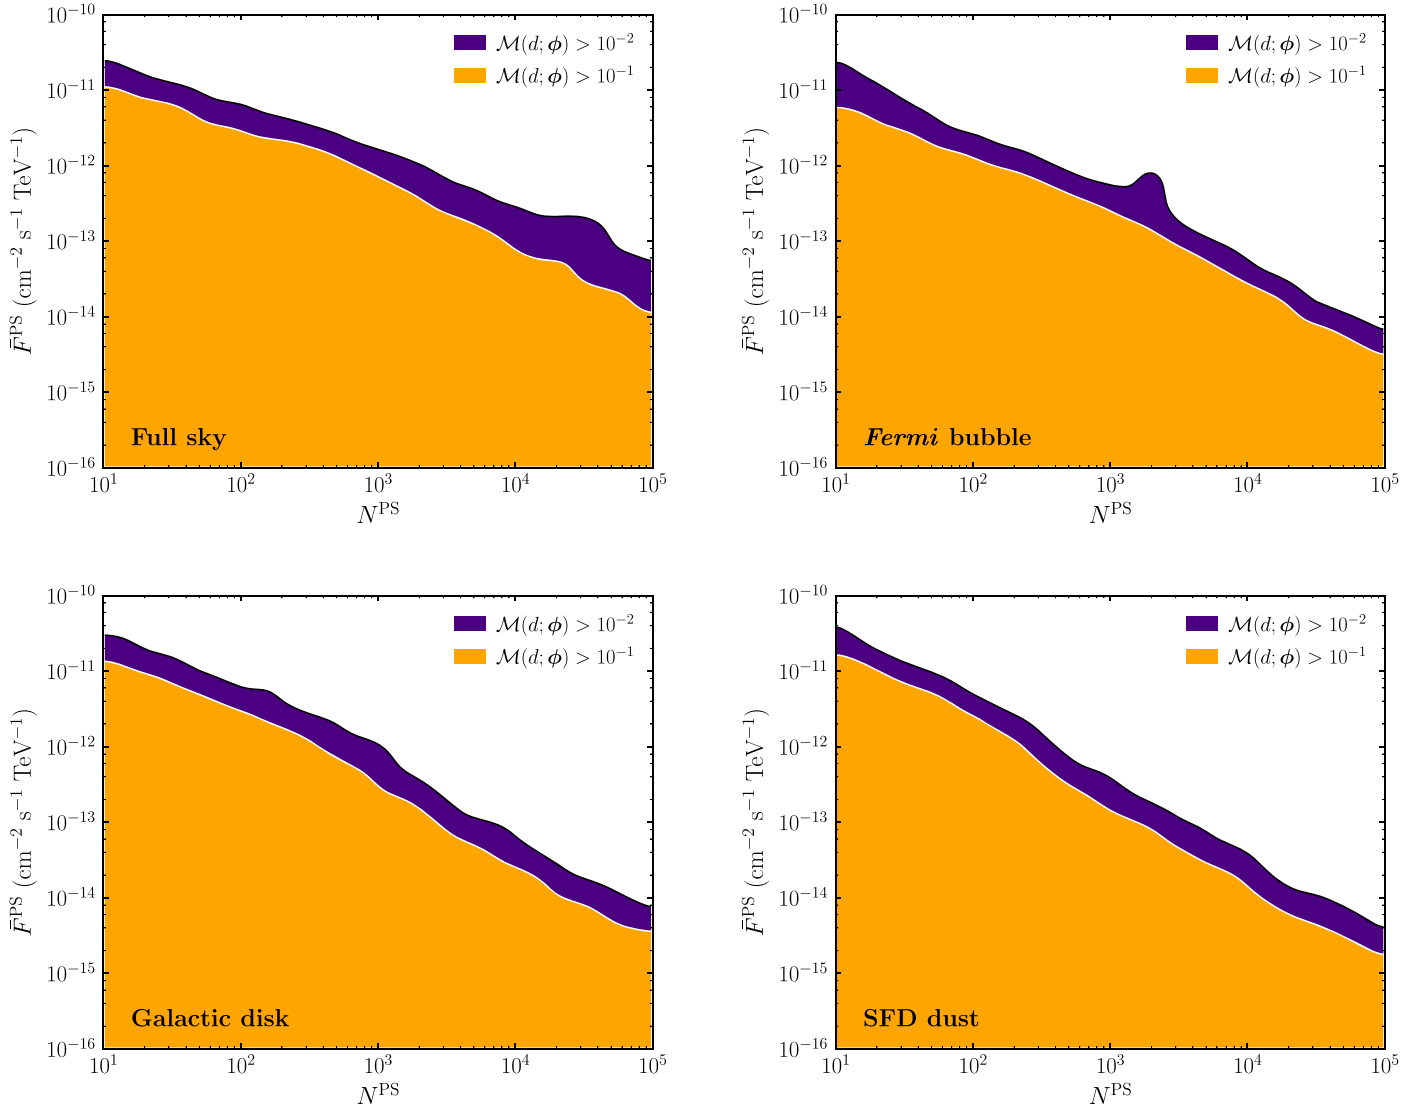


Figure 8. Map of the pointwise likelihood ratio, $\mathcal{M}(d; \phi)$, for the four different point-source spatial distributions considered in this work: isotropic sources over the full sky (top left panel), Fermi bubbles (top right panel), SFD dust (bottom right panel), and galactic disk (bottom left panel). In each case, the results are consistent with the background or Poissonian hypothesis, with the most significant p -value of 0.33 occurring for the SFD dust map.

$\Omega_{\lambda,0} = 0.692$, and $h = 0.678$ Ade et al. (2016). These output source-count distributions were then interfaced with the NPTF posterior, and the value of $\mathcal{M}(d; \phi)$ calculated for each point in parameter space. The A and F_b parameters scale with ρ_0 and L_{eff} , respectively, while n_1 and n_2 are set to 1.9 and -2 , respectively. A lower limit of $L_{\text{eff}} = 10^{52} \text{ erg yr}^{-1}$ was chosen to match the prior on F_b . The result is shown in Figure 10 for the isotropic template, and—consistent with our previous results—we see no evidence for any particular source class. These results, in addition to allowing a comparison with searches for source populations with fixed flux-characteristics (Aartsen et al. 2019), also are representative of the power and generality of the NPTF technique.

6. Conclusion

In this work, we have performed the first application of the non-Poissonian template fitting technique to search within the IceCube data set for neutrino point-source populations. Although IceCube presents novel challenges to the implementation of the NPTF, our work provides an explicit verification

that such difficulties can be addressed and that this technique is a viable method to search for such populations. In addition to being able to search for populations with an, in principle, arbitrary source-count function dN/dF , this method also allows us to search for point sources with peculiar spatial distributions, and here we have considered spatial templates following maps of the isotropic sky, Fermi bubbles, SFD dust map, and galactic disk. In all cases, no significant evidence of a point-source population has been detected, and so we have presented limits in their absence, as shown in Figure 8. Importantly, we have made the full posterior from our analysis publicly available, allowing specific theory predictions for contributions to the IceCube flux to be tested directly. This is exemplified by the application of our results to the space of standard-candle luminosity functions, shown in Figure 10.

There are a number of ways that the analysis presented here can be improved upon. The IceCube data set contains a large amount of information on the reconstruction quality of incident candidate neutrinos on an event-by-event basis. As the NPTF is a fundamentally binned method, much of this information is




Figure 9. Triangle plot for the case of a point-source population following the SFD dust map. The full posterior from which this was created is available for download. A is given in $\text{TeV cm}^2 \text{s}$ and F_b in $\text{TeV}^{-1} \text{cm}^{-2} \text{s}^{-1}$.

lost and is only exploited through the optimization of various high level cuts, such as on the energy range considered. Yet, there is significant scope to incorporate more of this information into the NPTF. For example, there is the potential to incorporate energy binning into the method, and with this, additional event information. Beyond expanding the neutrino data set, such extensions could play an important role in uncovering evidence for a population of astrophysical point sources, and unravelling the mystery surrounding the origin of the IceCube neutrinos.

The IceCube Collaborations acknowledge the significant contributions to this manuscript from Gabriel Collin, Jimmy DeLaunay, and Nicholas Rodd. The authors gratefully acknowledge the support from the following agencies and institutions: USA—U.S. National Science Foundation-Office of Polar Programs, U.S. National Science Foundation-Physics Division, Wisconsin Alumni Research Foundation, Center for High Throughput Computing (CHTC) at the University of

Wisconsin-Madison, Open Science Grid (OSG), Extreme Science and Engineering Discovery Environment (XSEDE), U.S. Department of Energy-National Energy Research Scientific Computing Center, Particle astrophysics research computing center at the University of Maryland, Institute for Cyber-Enabled Research at Michigan State University, and Astroparticle physics computational facility at Marquette University; Belgium—Funds for Scientific Research (FRS-FNRS and FWO), FWO Odysseus and Big Science programmes, and Belgian Federal Science Policy Office (Belspo); Germany—Bundesministerium für Bildung und Forschung (BMBF), Deutsche Forschungsgemeinschaft (DFG), Helmholtz Alliance for Astroparticle Physics (HAP), Initiative and Networking Fund of the Helmholtz Association, Deutsches Elektronen Synchrotron (DESY), and High Performance Computing cluster of the RWTH Aachen; Sweden—Swedish Research Council, Swedish Polar Research Secretariat, Swedish National Infrastructure for Computing (SNIC), and Knut and Alice Wallenberg Foundation; Australia—Australian




Figure 10. Map of the pointwise likelihood ratio for the isotropic template applied to the space of standard-candle luminosity functions. The standard-candle-approach models the luminosity function as sharply peaked in L_{eff} , and then the space of possible models is spanned by this parameter and the density of sources, ρ_0 . This figure also allows us to contrast our results with the 90% upper limit obtained using an analysis for steady point sources with a specified flux-distribution, as derived in Aartsen et al. (2019). Both of these results can be compared to a standard-candle population of sources that is compatible with the observed diffuse flux at $\pm 1\sigma$, as quoted in Aartsen et al. (2017b). To aid interpretation, we have overlaid the electromagnetic luminosities associated with several possible source classes: flat-spectrum radio quasars (FSRQ), BL Lacertae active galactic nuclei (BL LAC), galaxy cluster, and FanaroffRiley Class II radio galaxies (FR-II), following Kowalski (2015) and Taboada et al. (2018). We emphasize that these are not predicted neutrino luminosities, which are unknown, but highlight that current measurements provide information about the relative neutrino to photon luminosities of these sources. We note that the results in this figure were derived using the NPTF posterior, described in Appendix A, and show the power of our result to test specific model hypotheses. See the text for details.

Research Council; Canada—Natural Sciences and Engineering Research Council of Canada, Calcul Québec, Compute Ontario, Canada Foundation for Innovation, WestGrid, and Compute Canada; Denmark—Villum Fonden, Danish National Research Foundation (DNRF), Carlsberg Foundation; New Zealand—Marsden Fund; Japan—Japan Society for Promotion of Science (JSPS) and Institute for Global Prominent Research (IGPR) of Chiba University; Korea—National Research Foundation of Korea (NRF); Switzerland—Swiss National Science Foundation (SNSF); United Kingdom—Department of Physics, University of Oxford. We thank Derek Fox and Kohta Murase for useful discussions. N.L.R. acknowledges support from the Miller Institute for Basic Research in Science at the University of California, Berkeley.

Software: astropy 1.3 (Robitaille et al. 2013), corner 2.0.1 (Foreman-Mackey 2016), FIRESONG 1.5 (Taboada et al. 2018), healpy 1.8.6 (Zonca et al. 2019), HEALPix 3.10

(Gorski et al. 2005), pymultinest 2.1 (Buchner et al. 2014), MultiNest 3.10 (Feroz et al. 2009), and NPTFit 0.2 (Mishra-Sharma et al. 2017).

Appendix A Description of the Public Posterior

The posterior for each of the four templates can be found at https://icecube.wisc.edu/science/data/NPTF_7yr_posterior as an HDF5 file. Within the file, five tables named `Isotropic`, `Galactic_disk`, `Fermi_bubble`, `SFD_dust`, and `Northern_sky` contain the posterior for their respective templates.

Each table describes equally weighted samples using five columns. Four columns, labeled `ln_A`, `ln_Fb`, `n1`, and `n2` contain the coordinates for the sample in natural logarithmic parameter space for the differential source-count function normalization A and break F_b , while the power indices n_1 and n_2 are in linear space. The fifth column—labeled `loglikelihood`—gives the natural logarithm of the likelihood function at the location of the corresponding sample. In addition, each table has two attributes named `P_log_evidence` and `NP_log_evidence` that contain the natural logarithm of the evidence integral for the Poissonian model (\mathcal{L}_0) and non-Poissonian model (\mathcal{L}_1), respectively.

The root node of the HDF5 file also contains a series of attributes named `units_ln_A`, `units_ln_Fb`, `units_n1`, and `units_n2` that specify the units that the posterior sample coordinates are given in. Another series of root attributes named `prior_ln_A`, `prior_ln_Fb`, `prior_n1`, and `prior_n2` give the probability density of the uniform priors for each of the model parameters.

Finally, we emphasize that our analysis and, hence, these posteriors are constructed with the assumption that the astrophysical population produces neutrinos with an E^{-2} spectrum, as given in Equation (4).

Appendix B Isotropic Sources in the Northern Sky

Traditional searches for extragalactic point sources at IceCube are performed restricting to the northern or southern hemispheres. The motivation for this is the northern hemisphere, having a lower background, usually has an enhanced sensitivity. In Figure 11, we show the expected sensitivity and the pointwise likelihood ratio determined from the data for the northern sky, which should be contrasted to the full sky result for both hemispheres shown in Figures 5 and 8. Comparing the two results, it is clear that restricting the NPTF to the lower-background hemisphere only marginally improves the sensitivity. This suggests that the NPTF results are not being degraded by working with the full sky, and given our inclusion of a number of galactic templates, justifies the choice of both hemispheres used in the main text.




Figure 11. Sensitivity and limit (left panel) and the pointwise likelihood ratio (right panel), analogous to the full sky results shown in Figures 5 and 8, respectively, but here restricted to only the northern hemisphere. Interestingly, we see an at most factor of a few improvement in reach, suggesting that the NPTF technique is not being hampered by the increased background in the southern hemisphere. For the top figure, we have explicitly reproduced the median of the full sky $\mathcal{M}(d; \phi) = 0.1$ distribution to allow a direct comparison.

ORCID iDs

M. Ahlers <https://orcid.org/0000-0003-0709-5631>
 E. Bernardini <https://orcid.org/0000-0003-3108-1141>
 P. Dave <https://orcid.org/0000-0002-3879-5115>
 J. J. DeLaunay <https://orcid.org/0000-0001-5229-1995>
 P. A. Evenson <https://orcid.org/0000-0001-7929-810X>
 A. Franckowiak <https://orcid.org/0000-0002-5605-2219>
 J. Gallagher <https://orcid.org/0000-0001-8608-0408>
 S. Garrappa <https://orcid.org/0000-0003-2403-4582>
 U. Katz <https://orcid.org/0000-0002-7063-4418>
 A. Kheirandish <https://orcid.org/0000-0001-7074-0539>
 H. Pandya <https://orcid.org/0000-0002-6138-4808>
 B. R. Safdi <https://orcid.org/0000-0001-9531-1319>
 M. Santander <https://orcid.org/0000-0001-7297-8217>
 N. L. Strotjohann <https://orcid.org/0000-0002-4667-6730>
 J. Vandenbroucke <https://orcid.org/0000-0002-9867-6548>

References

- Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013c, *ApJ*, 779, 132
 Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013d, arXiv:1309.6979
 Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013e, *NIMPA*, A711, 73
 Aartsen, M. G., Abbasi, R., Ackermann, M., et al. 2014c, *JINST*, 9, P03009
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2015b, *ApJ*, 809, 98
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2015c, *PhRvL*, 115, 081102
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2015f, arXiv:1510.05222
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2016a, *ApJL*, 824, L28
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2016b, *ApJ*, 833, 3
 Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2017a, *ApJ*, 835, 151
 Aartsen, M. G., Ackermann, M., Adams, G., et al. 2013a, *PhRvL*, 111, 021103
 Aartsen, M. G., Ackermann, M., Adams, G., et al. 2013b, *Sci*, 342, 1242856
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2015a, *PhRvD*, 91, 022001
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2015e, *Aph*, 66, 39
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2015d, *PhRvL*, 114, 171102
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2017b, arXiv:1710.01191
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2018, *Sci*, 361, eaat1378
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2019, *EPJ C*, 79, 234
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2014a, *PhRvL*, 113, 101101
 Aartsen, M. G., Ackermann, M., Adams, J., et al. 2014b, *ApJ*, 796, 109
 Abbasi, R., Abdou, Y., Abu-Zayyad, T., et al. 2010, *NIMPA*, 618, 139
 Abbasi, R., Abdou, Y., Abu-Zayyad, T., et al. 2011, *ApJ*, 732, 18
 Abbasi, R., Ackermann, M., Adams, J., et al. 2009, *NIMPA*, 601, 294
 Abeyssekara, A. U., Albert, A., Alfaro, R., et al. 2018, *JCAP*, 1802, 049
 Achterberg, A., Ackermann, M., Adams, J., et al. 2006, *Aph*, 26, 155
 Ade, P. A. R., Agnahim, N., Arnaud, M., et al. 2016, *A&A*, 594, A13
 Ahrens, J., Bai, X., Bay, R., et al. 2004, *NIMPA*, 524, 169
 Albert, A., André, M., Anghinolfi, M., et al. 2017, *PhRvD*, 96, 082001
 Anchordoqui, L. A., & Montaruli, T. 2010, *ARNPS*, 60, 129
 Anchordoqui, L. A., Barger, V., Cholis, I., et al. 2014, *JHEAp*, 1, 1
 Ando, S., Feyereisen, M. R., & Fornasa, M. 2017, *PhRvD*, 95, 103003
 Bechtol, K., Ahlers, M., Di Mauro, M., Ajello, M., & Vandenbroucke, J. 2017, *ApJ*, 836, 47
 Bhattacharya, A., Reno, M. H., & Sarcevic, I. 2014, *JHEP*, 06, 110
 Blanco, C., & Hooper, D. 2017, *JCAP*, 1712, 017
 Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, 564, A125
 Chianese, M., & Merle, A. 2017, *JCAP*, 1704, 017
 Chianese, M., Miele, G., & Morisi, S. 2017, *PhLB*, 773, 591
 Cohen, T., Murase, K., Rodd, N. L., Safdi, B. R., & Soreq, Y. 2017, *PhRvL*, 119, 021102
 Di Mauro, M., Manconi, S., Zechlin, H. S., et al. 2018, *ApJ*, 856, 106
 Di Teodoro, E. M., McClure-Griffiths, N. M., Lockman, F. J., et al. 2018, *ApJ*, 855, 33
 Ema, Y., Jinno, R., & Moroi, T. 2014, *PhLB*, 733, 120
 Esmaili, A., Kang, S. K., & Serpico, P. D. 2014, *JCAP*, 1412, 054
 Esmaili, A., & Serpico, P. D. 2013, *JCAP*, 1311, 054
 Feldstein, B., Kusenko, A., Matsumoto, S., & Yanagida, T. T. 2013, *PhRvD*, 88, 015004
 Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601
 Feyereisen, M. R., Ando, S., & Lee, S. K. 2015, *JCAP*, 1509, 027
 Feyereisen, M. R., Gaggero, D., & Ando, S. 2018, *PhRvD*, 97, 103017
 Feyereisen, M. R., Tamborra, I., & Ando, S. 2017, *JCAP*, 1703, 057
 Foreman-Mackey, D. 2016, *JOSS*, 24
 Gendreau, K. C., Barcons, X., & Fabian, A. C. 1998, *MNRAS*, 297, 41
 Georgantopoulos, I., Stewart, G. C., Shanks, T., Griffiths, R. E., & Boyle, B. J. 1993, *MNRAS*, 262, 619
 Glauch, T., & Turcati, A. 2018, ICRC (Busan), 35, 1014
 Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Halzen, F., & Hooper, D. 2002, *RPPH*, 65, 1025
 Hasinger, G., Burg, R., Giacconi, R., et al. 1993, *A&A*, 275, 1
 Hoel, P. G., Port, S. C., & Stone, C. J. 1971, *Introduction to Probability Theory* (Boston, MA: Houghton Mifflin)
 Hooper, D. 2016, *JCAP*, 1609, 002
 Hopkins, A. M., & Beacom, J. F. 2006, *ApJ*, 651, 142
 Kalashev, O. E., & Kuznetsov, M. Y. 2017, *JETPL*, 106, 73
 Kalashev, O. K., & Kuznetsov, M. Yu 2016, *PhRvD*, 94, 063535
 Kowalski, M. 2015, *JPhCS*, 632, 012039
 Kuznetsov, M. Yu 2017, *JETPL*, 105, 561

- Learned, J. G., & Mannheim, K. 2000, [ARNPS](#), **50**, 679
- Lee, S. K., Lisanti, M., & Safdi, B. R. 2015, [JCAP](#), **1505**, 056
- Lee, S. K., Lisanti, M., Safdi, B. R., Slatyer, T. R., & Xue, W. 2016, [PhRvL](#), **116**, 051103
- Linden, T., Rodd, N. L., Safdi, B. R., & Slatyer, T. R. 2016, [PhRvD](#), **94**, 103013
- Lisanti, M., Mishra-Sharma, S., Necib, L., & Safdi, B. R. 2016, [ApJ](#), **832**, 117
- Lisanti, M., Mishra-Sharma, S., Rodd, N. L., Safdi, B. R., & Wechsler, R. H. 2018, [PhRvD](#), **97**, 063005
- Malyshev, D., & Hogg, D. W. 2011, [ApJ](#), **738**, 181
- Mishra-Sharma, S., Rodd, N. L., & Safdi, B. R. 2017, [AJ](#), **153**, 253
- Miyaji, T., & Griffiths, R. E. 2002, [ApJL](#), **564**, L5
- Murase, K., & Waxman, E. 2016, [PhRvD](#), **94**, 103006
- Padovani, P., Petropoulou, M., Giommi, P., & Resconi, E. 2015, [MNRAS](#), **452**, 1877
- Padovani, P., & Resconi, E. 2014, [MNRAS](#), **443**, 474
- Perri, M., & Giommi, P. 2000, [A&A](#), **362**, L57
- Petropoulou, M., Dimitrakoudis, S., Padovani, P., Mastichiadis, A., & Resconi, E. 2015, [MNRAS](#), **448**, 2412
- Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, [A&A](#), **558**, A33
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, [ApJ](#), **500**, 525
- Schmaltz, M., & Weiner, N. 2017, arXiv:1709.09164
- Spiering, C. 2018, [PPN](#), **49**, 497
- Stachurska, J. 2018, in 28th Int. Conf. Neutrino Physics and Astrophysics, **386**
- Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, [ApJ](#), **724**, 1044
- Sudoh, T., Totani, T., & Kawanaka, N. 2018, [PASJ](#), **70**, 49
- Taboada, I., Tung, C. F., & Wood, J. 2018, ICRC (Busan), 35, 663
- Tamborra, I., Ando, S., & Murase, K. 2014, [JCAP](#), **1409**, 043
- Waxman, E., & Bahcall, J. N. 1999, [PhRvD](#), **59**, 023002
- Yuan, C., Murase, K., & Mészáros, P. 2019, arXiv:1904.06371
- Zavala, J. 2014, [PhRvD](#), **89**, 123516
- Zechlin, H.-S., Cuoco, A., Donato, F., Fornengo, N., & Regis, M. 2016a, [ApJL](#), **826**, L31
- Zechlin, H.-S., Cuoco, A., Donato, F., Fornengo, N., & Vittino, A. 2016b, [ApJS](#), **225**, 18
- Zonca, A., Singer, L., Lenz, D., et al. 2019, [JOSS](#), **4**, 1298