

Constraining the contribution of Seyfert galaxies to the diffuse neutrino flux in light of point source observations

Lena Saurenhaus^{1,2,*}, Francesca Capel¹, Foteini Oikonomou³, and Johannes Buchner⁴

¹*Max Planck Institute for Physics, 85748 Garching, Germany*

²*Department of Physics, TUM School of Natural Sciences, Technical University of Munich, 85747 Garching, Germany*

³*Institutt for fysikk, Norwegian University of Science and Technology, 7491 Trondheim, Norway*

⁴*Max Planck Institute for Extraterrestrial Physics, 85748 Garching, Germany*



(Received 8 July 2025; accepted 8 December 2025; published 12 January 2026)

Recently, the IceCube Collaboration reported evidence for TeV neutrino emission from several nearby Seyfert galaxies, with the highest significance found for NGC 1068. Assuming stochastic proton acceleration in magnetized turbulence inside the corona, we model the neutrino emission of Seyfert galaxies as a function of their x-ray luminosity. Applying our model to NGC 1068, we obtain a good fit to the public IceCube data and constrain the coronal radius to $\lesssim 5R_S$ by comparing our MeV γ -ray predictions to Fermi-LAT observations. Extending to the full Seyfert population, we estimate their diffuse neutrino contribution and find that they can explain a significant fraction of the observed flux below 10 TeV. However, scenarios with highly turbulent coronae and high cosmic-ray pressure across the population are ruled out. In particular, if all sources shared the best-fit parameters obtained for NGC 1068, their cumulative neutrino emission would exceed current upper limits at TeV energies by 3.8σ . Our results, informed by both neutrino and γ -ray data, show that those Seyfert galaxies that emerge as neutrino point sources must be exceptionally efficient neutrino emitters and are not representative of the broader population.

DOI: [10.1103/f66p-k6z9](https://doi.org/10.1103/f66p-k6z9)

I. INTRODUCTION

The discovery of a diffuse flux of astrophysical neutrinos by the IceCube Neutrino Observatory has provided unambiguous evidence for the existence of powerful cosmic accelerators capable of producing neutrinos up to $\sim$ PeV energies [1,2]. Yet, more than a decade later, identifying the origins of these neutrinos remains challenging.

The recent IceCube observation of an excess of neutrinos from the direction of the nearby Seyfert galaxy NGC 1068 is a promising development [3]. Evidence is building for neutrino emission from this source class, with further studies by the IceCube Collaboration showing possible signals from other Seyfert galaxies, particularly NGC 4151 and CGCG 420-015 [4,5]. Independently, a search using public IceCube data identified NGC 4151 and NGC 3079 as possible neutrino sources [6], while [7] highlighted the

coincidence of two $\sim$ 100 TeV IceCube alert events with the location of the Seyfert galaxy NGC 7469.

NGC 1068 remains the strongest source in the northern sky, with a significance of 4.2σ . As such, its potential as a neutrino emitter has been extensively studied from both observational and theoretical perspectives (see Ref. [8] and references therein). Located at a distance of $d_L = 11.14$ Mpc [9], NGC 1068 hosts a luminous obscured active nucleus with an intrinsic x-ray luminosity of $L_X = 4.17 \times 10^{43}$ erg s $^{-1}$ in the 2–10 keV range (after rescaling the measurement of [10] to the updated luminosity distance). This makes it the Seyfert galaxy with the highest intrinsic 2–10 keV x-ray flux in the Northern Hemisphere [11].

Notably, no TeV γ -ray emission has been detected from NGC 1068 despite searches by ground-based instruments such as HESS [12], MAGIC [13], or HAWC [14]. This nondetection implies that any γ rays produced in the same hadronic processes responsible for neutrino emission must be efficiently absorbed within the source. NGC 1068, thus, represents a compelling example of a “hidden” neutrino source, a class that has been independently proposed to explain the tension between measurements of the diffuse γ -ray and < 10 TeV neutrino fluxes [15].

Several theoretical models suggest that the observed neutrinos associated with NGC 1068 originate from hadronic interactions in a dense corona located close to the

*Contact author: lsaurenh@mpp.mpg.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

central supermassive black hole (SMBH). Murase *et al.* [16] propose that, inside this corona, protons are accelerated via stochastic acceleration in magnetized turbulence and subsequently interact with ambient gas and coronal x-rays, producing high-energy neutrinos and γ rays. The latter are attenuated in $\gamma\gamma$ interactions with x-rays, making it possible to explain the observed neutrino emission without violating any γ -ray observations. Building a two-zone model of the corona and the circumnuclear starburst region of NGC 1068, Eichmann *et al.* [17] confirm that the neutrinos can be produced only inside the corona, while the starburst region contributes to the γ -ray emission. Based on radio observations of an excess in the millimeter range, which can be interpreted as coronal synchrotron emission, Inoue *et al.* [18,19] assume that both electrons and protons are accelerated via diffusive shock acceleration. In addition, magnetic reconnection inside the corona has been proposed as a possible proton acceleration scenario [20–22].

Alternative scenarios involve outflows driven by the active galactic nucleus (AGN) [23–25], jet interactions with the interstellar medium [26], or stellar-mass black holes embedded in the AGN accretion disk [27], as well as muon pair production [28,29] or beta decays of neutrons [29,30]. However, the AGN corona is particularly compelling as a neutrino production site, since in this scenario the coronal x-ray photons attenuate the high-energy γ rays that would otherwise be expected to accompany the $> \text{TeV}$ neutrinos.

Several studies have also investigated the diffuse neutrino emission expected from the Seyfert population as a whole in light of the NGC 1068 neutrino observations. Padovani *et al.* [31] used a template neutrino spectrum based on the disk-corona model reported in [32] and assumed a linear scaling between x-ray and neutrino luminosities. They found that the Seyfert population can account for the observed neutrino flux at energies of $\sim 1\text{--}10 \text{ TeV}$, but an additional population is needed to explain the flux at higher energies. Other works have explored the diffuse contribution using alternative models for neutrino production in Seyfert galaxies, including a leaky-box model with fixed proton injection [33] and a strong turbulence acceleration scenario [34], finding results that slightly overshoot or are broadly consistent with the IceCube observations, respectively.

In this work, we aim to explore the constraints of the diffuse flux from the Seyfert galaxy population for the key physical parameters of the disk-corona model and quantify how unusual a source such as NGC 1068 is expected to be in this context. Furthermore, we also investigate the variance in the model parameters that is required to explain neutrino observations from the directions of other nearby Seyfert galaxies.

We present the details of our neutrino emission model in Sec. II. This model is then fit to the public IceCube data for the case of NGC 1068 in Sec. III, including an evaluation of the γ -ray flux predictions. Using the result of this fit as a

benchmark, we extrapolate our model to the entire population of Seyfert galaxies discussed in Sec. IV and calculate the resulting diffuse neutrino flux in Sec. V. We discuss our results in Sec. VI and conclude in Sec. VII.

II. NEUTRINO EMISSION MODEL FOR A SINGLE SOURCE

The corona of an AGN is a magnetized plasma of hot electrons located in the vicinity of the central region of the accretion disk [35–41]. It is very luminous in x-rays, which are produced via inverse Compton scattering of optical and UV disk photons (e.g., [42,43]). However, the exact geometry and origin of the corona are not yet fully understood.

Similar to the model proposed by Murase *et al.* [16], we assume that, inside AGN coronae, protons are accelerated via stochastic acceleration in magnetized turbulence up to energies of a few hundred TeV [44–46]. Subsequently, they undergo pp interactions with ambient gas and $p\gamma$ interactions with coronal x-rays, producing high-energy neutrinos and γ rays. The latter are attenuated in $\gamma\gamma$ interactions with x-rays, initiating electromagnetic cascades and escaping the corona at MeV energies. As in [16], we do not consider any primary electron acceleration.

We describe the relevant properties of the corona of a Seyfert galaxy as a function of its intrinsic 2–10 keV x-ray luminosity L_X , following [16]. Making use of the fact that the nonthermal proton spectrum resulting from stochastic acceleration follows a power law with an exponential cutoff, we then calculate the neutrino and the cascaded γ -ray spectrum of the source. To this end, we use the multimessenger simulation code AM^3 [47].

A. Properties of the AGN corona

We model the AGN corona as a sphere with radius $R_c = r_c R_S$, where r_c is the dimensionless coronal radius and $R_S = 2GM_{\text{BH}}/c^2$ is the Schwarzschild radius of the central SMBH. The typical black-hole mass of an AGN with x-ray luminosity L_X is given by

$$M_{\text{BH}} \simeq 2 \times 10^7 M_\odot \left(\frac{L_X}{1.16 \times 10^{43} \text{ erg/s}} \right)^{0.746}, \quad (1)$$

which is an empirical relation derived from x-ray observations of a sample of AGNs [48].

The intrinsic x-ray spectrum of the corona is described by a single power law with an exponential cutoff:

$$\left(\frac{dN_\gamma}{dE_\gamma dt} \right)_X \propto E_\gamma^{-\Gamma_X} \exp(-E_\gamma/E_X^{\text{max}}). \quad (2)$$

The spectral index Γ_X and the cutoff energy E_X^{max} are related to the Eddington ratio $\lambda_{\text{Edd}} = L_{\text{bol}}/L_{\text{Edd}}$ as

$$\Gamma_X = 0.167 \log \lambda_{\text{Edd}} + 2.0 \quad (3)$$

and

$$E_X^{\max} = [-74 \log \lambda_{\text{Edd}} + 150] \text{ keV}, \quad (4)$$

respectively [49,50]. L_{bol} is the bolometric luminosity, which we calculate using the relation derived in Ref. [51], and L_{Edd} is the Eddington luminosity, which is given by $L_{\text{Edd}} \simeq 1.3 \times 10^{45} \text{ erg/s} (M_{\text{BH}}/10^7 M_{\odot})$.

We calculate the Thomson optical depth of the corona, τ_T , as

$$\tau_T \simeq 10^{(2.16 - \Gamma_X)/1.062} \left(\frac{k_B T_e}{1 \text{ keV}} \right)^{-0.3}, \quad (5)$$

where $k_B T_e \simeq E_X^{\max}/2$ is the electron temperature [50]. This relation, which has been derived from simulations of coronal x-ray spectra for different values of τ_T and $k_B T_e$, yields $\tau_T \sim 0.3\text{--}0.6$ for $42 \leq \log L_X \leq 46$ (see Table I); i.e., in our model, the corona is optically thin for all considered x-ray luminosities. However, other recent studies suggest that some AGNs may also have optically thick coronae with $\tau_T > 1$ [52,53]. Assuming that the fraction of electron-positron pairs in the coronal plasma is negligibly small, we set the thermal proton number density n_p to be equal to the thermal electron number density n_e , which results in

$$n_p = n_e = \frac{\tau_T}{\sigma_T R_c} \quad (6)$$

with the Thomson cross section σ_T (e.g., [16,54]).

The plasma beta is defined as $\beta = P_{\text{th}}/P_{\text{mag}}$, where $P_{\text{th}} = n_p k_B T_p$ is the thermal gas pressure at the virial temperature

$$T_p = \frac{GM_{\text{BH}} m_p}{3R_c k_B} = \frac{m_p c^2}{6k_B r_c} \quad (7)$$

and $P_{\text{mag}} = B^2/8\pi$ is the magnetic pressure. In general, the plasma beta in AGN coronae is not very well constrained. Results of various magnetohydrodynamics simulations point toward rather low values of $\beta \lesssim 1$ [36,40,41]. However, Inoue, Takasao, and Khangulyan [55] argue that the absence of powerful jets and the resulting lack of large-scale poloidal magnetic fields in the coronal region of

Seyfert galaxies leads to $\beta \sim 10\text{--}100$. In line with other works studying AGN coronae as possible neutrino sources [16,17,29,33,46], we set $\beta = 1$. The strength of the magnetic field inside the corona can then be calculated as

$$B = \sqrt{\frac{8\pi n_p k_B T_p}{\beta}}. \quad (8)$$

The properties of the AGN corona for different x-ray luminosities and $r_c = 5$ are summarized in Table I.

B. Proton acceleration

The presence of strong and turbulent magnetic fields inside AGN coronae suggests that they are possible sites of stochastic particle acceleration in magnetized turbulence, also referred to as second-order Fermi acceleration. Protons are accelerated via repeated scattering by plasma waves, causing them to gradually gain energy. For this process to be efficient, not only a strong magnetic field, but also a sufficient level of turbulence is required. We quantify the latter in terms of the inverse turbulence strength η , defined as

$$\eta = \frac{B^2}{\delta B^2} = \frac{B_0^2 + \delta B^2}{\delta B^2}, \quad (9)$$

where B_0 and δB are the regular and the turbulent component of the magnetic field, respectively.

Stochastic acceleration can be modeled as diffusion in momentum space described by a momentum diffusion coefficient. For relativistic particles, this diffusion coefficient can be approximated as

$$D_{E_p} \simeq \frac{c}{\eta R_c} \left(\frac{V_A}{c} \right)^2 \left(\frac{E_p}{e B R_c} \right)^{q-2} E_p^2, \quad (10)$$

where $V_A = B/\sqrt{4\pi m_p n_p}$ is the Alfvén velocity and q is the spectral index of the turbulence power spectrum [56]. For proton energies $E_p \ll E_p^{\max}$, where $E_p^{\max}$ is the energy at which cooling and escape losses start to dominate over acceleration, we approximate the resulting nonthermal proton spectrum by a power law with spectral index $1 - q$ [57,58]. Starting from the momentum diffusion equation that describes the evolution of the particle distribution function f , under stochastic acceleration,

$$\frac{\partial f}{\partial t} = \frac{1}{E_p^2} \frac{\partial}{\partial E_p} \left[E_p^2 D_{E_p} \frac{\partial f}{\partial E_p} \right], \quad (11)$$

it can be seen that, in a stationary system ($\partial_t f = 0$), f has to satisfy $E_p^2 D_{E_p} \partial_{E_p} f = \text{const}$. With $D_{E_p} \propto E_p^q$ [cf. Eq. (10)] and $dn_p/dE_p = 4\pi E_p^2 f/c^3$, this yields $dn_p/dE_p \propto E_p^{1-q}$. Thus, for Kolmogorov turbulence with $q = 5/3$, the resulting proton spectrum is a power law with spectral

TABLE I. Properties of the AGN corona for different L_X and $r_c = 5$. L_X is given in units of erg/s and n_p in units of cm^{-3} .

$\log L_X$	$\log M_{\text{BH}}/M_{\odot}$	τ_T	$\log n_p$	B [kG]
42.0	6.51	0.56	11.25	15.0
43.0	7.25	0.49	10.45	5.9
44.0	8.00	0.43	9.64	2.3
45.0	8.74	0.37	8.84	0.93
46.0	9.49	0.34	8.05	0.37

index $-2/3$ and an exponential cutoff at $E_p^{\max}$:

$$\frac{dN_p}{dE_p} \propto E_p^{-2/3} \exp(-E_p/E_p^{\max}). \quad (12)$$

The corresponding acceleration timescale is given by (e.g., [16,56])

$$t_{\text{acc}} \simeq \frac{E_p^2}{D_{E_p}} = \eta \left(\frac{c}{V_A} \right)^2 \frac{R_c}{c} \left(\frac{E_p}{eBR_c} \right)^{2-q}. \quad (13)$$

We determine the normalization of the proton spectrum in Eq. (12) by the total energy injected into proton acceleration, which can be expressed in terms of the ratio of the cosmic-ray (CR) pressure to the thermal gas pressure inside the corona, $P_{\text{CR}}/P_{\text{th}}$, with

$$P_{\text{CR}} = \frac{1}{3} \int dE_p E_p \frac{dN_p}{dE_p dV}. \quad (14)$$

We set $P_{\text{CR}}/P_{\text{th}} \leq 0.5$ as a strict upper limit on this pressure ratio [20]. The reasoning behind this is as follows: In the limit $P_{\text{CR}}/P_{\text{th}} = 0.5$, the total kinetic energy of nonthermal particles is equal to half of the gravitational binding energy of the corona. According to the virial theorem, in this case all the kinetic energy of the system is contained in the nonthermal particle population and there is no room for any kinetic energy in the thermal particle population. This means that, for $P_{\text{CR}}/P_{\text{th}} > 0.5$, the corona is no longer stable. The injected CR proton luminosities that correspond to a pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$ are listed in Table II for different L_X and two benchmark values of η . For all considered x-ray luminosities and turbulence levels, the required proton luminosity is much lower than the Eddington luminosity, with a maximum of $L_p/L_{\text{Edd}} \sim 4\%$ for $\eta = 1$ and $L_X = 10^{42}$ erg/s.

C. Proton cooling and escape

Relevant cooling processes for nonthermal protons inside AGN coronae include pp interactions, $p\gamma$ interactions, Bethe-Heitler pair production, and synchrotron radiation [16,17].

TABLE II. Proton luminosities corresponding to a pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$ and their ratio to the Eddington luminosity for different L_X and η . All luminosities are calculated for a coronal radius of $r_c = 5$.

L_X [erg/s]	$\eta = 1$		$\eta = 100$	
	L_p [erg/s]	L_p/L_{Edd}	L_p [erg/s]	L_p/L_{Edd}
10^{42}	1.6×10^{43}	0.039	4.0×10^{42}	0.010
10^{44}	4.9×10^{44}	0.038	1.0×10^{44}	0.008
10^{46}	1.3×10^{46}	0.031	2.7×10^{45}	0.007

The target photon spectrum for $p\gamma$ interactions and Bethe-Heitler pair production is the sum of the optical and UV spectrum of the accretion disk and the x-ray spectrum of the corona. The disk spectrum is a multi-temperature blackbody spectrum, which we approximate as

$$\left(\frac{dN_\gamma}{dE_\gamma dt} \right)_{\text{disk}} \propto E_\gamma^{-5/3} \exp(-E_\gamma/k_B T_{\text{max}}) \quad (15)$$

for photon energies $E_\gamma \geq k_B T(R_{\text{out}})$ [59]. The maximum temperature of the disk near the SMBH, T_{max} , is given by

$$T_{\text{max}} \simeq \left(\frac{L_{\text{disk}}}{200\pi R_S^2 \sigma_{\text{SB}}} \right)^{1/4}, \quad (16)$$

where $L_{\text{disk}} \simeq 0.5 L_{\text{bol}}$ [16] is the disk luminosity and σ_{SB} is the Stefan-Boltzmann constant. $T(R_{\text{out}})$ denotes the temperature at the outer radius of the accretion disk, which is $k_B T(R_{\text{out}}) \simeq 0.2$ eV for $R_{\text{out}} = 10^3 R_S$ [cf. Eq. (8.5) in Ref. [59]]. As already explained in Sec. II A, we describe the x-ray spectrum of the corona by a single power law with an exponential cutoff. The resulting disk-corona spectrum, which is obtained by summing Eqs. (2) and (15), then depends on only the x-ray luminosity of the source.

The proton energy threshold for photopion production is $E_{p-\pi}^{\text{th}} \simeq 3.4$ PeV $(E_\gamma/10 \text{ eV})^{-1}$ [60]. Thus, the proton energies that we consider are not high enough for photopion production with disk photons, and the main target photons for $p\gamma$ interactions are x-rays emitted by the corona. The optical and UV disk spectrum is mainly relevant for Bethe-Heitler pair production with a threshold energy of $E_{p-\text{BH}}^{\text{th}} \simeq 48$ TeV $(E_\gamma/10 \text{ eV})^{-1}$ [60]; i.e., the required proton energy is approximately 2 orders of magnitude lower than for $p\gamma$ interactions.

Besides the cooling processes mentioned above, we also take into account particle escape via infall onto the SMBH and diffusion. The infall timescale is given by

$$t_{\text{fall}} = \frac{R_c}{\alpha V_K} \quad (17)$$

with the viscosity parameter $\alpha = 0.1$ and the Keplerian velocity $V_K = \sqrt{GM_{\text{BH}}/R_c} = c/\sqrt{2r_c}$ [16]. The timescale for diffusive escape from the corona can be estimated as (e.g., [16,58])

$$t_{\text{diff}} \simeq \frac{9 R_c}{\eta c} \left(\frac{E_p}{eBR_c} \right)^{q-2}. \quad (18)$$

We calculate the cutoff energy of the proton spectrum, $E_p^{\max}$, by solving $t_{\text{acc}}(E_p^{\max}) = t_{\text{loss}}(E_p^{\max})$ with the acceleration timescale given in Eq. (13) and the total loss timescale for protons given by

$$t_{\text{loss}}^{-1} = t_{pp}^{-1} + t_{p\gamma}^{-1} + t_{\text{BH}}^{-1} + t_{\text{syn}}^{-1} + t_{\text{diff}}^{-1} + t_{\text{fall}}^{-1}. \quad (19)$$

The timescales for pp interactions, $p\gamma$ interactions, Bethe-Heitler (BH) pair production, and synchrotron radiation are calculated directly by the multimessenger simulation code AM^3 [47], which we use to compute the resulting neutrino and γ -ray emission. The relative importance of the different cooling and escape processes at different x-ray luminosities is discussed in Appendix A.

D. Neutrino spectra

Inside the AGN corona, high-energy neutrinos can be produced in both pp interactions with ambient gas and $p\gamma$ interactions with x-rays. Taking into account all energy and escape loss processes for protons described in Sec. II C, we calculate the neutrino spectrum of a source with a given x-ray luminosity using AM^3 , an open-source code that allows one to simulate the nonthermal particle emission of an astrophysical source [47]. As mentioned in Sec. II A, the coronal radius is taken to be a multiple of the Schwarzschild radius, $R_c = r_c R_S$. The density of thermal protons, which act as a target for pp interactions, is given by Eq. (6), while the magnetic field strength is set according to Eq. (8). The injected nonthermal proton spectrum as well as the target photon field relevant for $p\gamma$ and Bethe-Heitler interactions are specified in Secs. II B and II C, respectively.

Figure 1 shows the resulting neutrino spectra for different x-ray luminosities and two different values of η . In general, sources with a higher x-ray luminosity also have a higher neutrino luminosity. Given the assumption that the pressure ratio $P_{\text{CR}}/P_{\text{th}}$ is constant across sources of all x-ray luminosities, we find that the total neutrino luminosity L_ν roughly follows a power-law scaling of the form

$$L_\nu \propto L_X^{0.7}. \quad (20)$$

This means that, in our model, the neutrino luminosity increases slightly less than linearly with the x-ray luminosity. The turbulence strength parameter η determines the efficiency of the proton acceleration process and, thus, the cutoff energy of the proton spectrum and also of the resulting neutrino spectrum. Smaller values of η corresponding to very turbulent magnetic fields inside the corona lead to higher cutoff energies, since proton acceleration is more efficient. In this case, the acceleration process is stopped by $p\gamma$ and Bethe-Heitler interactions, and we find lower cutoff energies for sources with higher L_X due to the higher target photon density inside the corona. For larger values of η , the confinement time of protons inside the corona becomes shorter, and the acceleration process is stopped by protons escaping the corona via diffusion. This leads to higher cutoff energies for sources with higher L_X , since more luminous sources have larger coronae.




FIG. 1. Neutrino spectra for different x-ray luminosities computed for $\eta = 10$ (top panel) and $\eta = 150$ (bottom panel). The neutrino spectra are normalized so that $P_{\text{CR}}/P_{\text{th}} = 0.5$, and the dimensionless coronal radius is set to $r_c = 5$. $E_\nu L_{E_\nu} \equiv E_\nu^2 dN_\nu/dE_\nu dt$ denotes the differential all-flavor neutrino luminosity.

Overall, both pp and $p\gamma$ interactions contribute to the production of neutrinos. Only for low x-ray luminosities in combination with large η does the contribution from $p\gamma$ interactions become subdominant compared to that from pp interactions due to the low cutoff energy of the corresponding proton spectrum. A more detailed discussion of the relative importance of the different cooling and escape processes for protons can be found in Appendix A.

E. Cascaded γ -ray spectra

In addition to neutrinos, hadronic interactions also produce high-energy γ rays with an intrinsic luminosity comparable to that of neutrinos. Unlike neutrinos, however, these γ rays do not simply escape from the corona. Instead, they are attenuated in $\gamma\gamma$ interactions with x-rays, which initiate electromagnetic cascades inside the

corona [16,17,29,32]. Secondary electrons and positrons produced in pp , $p\gamma$, and Bethe-Heitler interactions also contribute to these electromagnetic cascades through both synchrotron radiation and inverse Compton scattering.

The interaction timescales for the leptonic processes mentioned above are much shorter than those for the hadronic processes discussed in Sec. II C. A single AM^3 simulation that self-consistently accounts for both hadronic and leptonic interactions would, therefore, require a solver time step small enough to resolve all leptonic processes as well as a sufficiently long total simulation time for the hadronic component to reach a steady state. This would make the computation very time consuming and, thus, unfeasible across a grid of source parameters. To address this, we decouple the purely leptonic part of the simulation, which describes the electromagnetic cascade, from the hadronic part and perform two consecutive simulations. First, we run a purely hadronic AM^3 simulation to compute the neutrino spectrum as well as the intrinsic γ -ray and electron-positron spectra. Then, we inject the latter together with the photon spectrum of the accretion disk and the corona into a second, purely leptonic AM^3 simulation in order to calculate the spectrum of cascaded γ rays escaping from the corona of the source.

As shown in Fig. 2, only γ rays with energies $E_\gamma \lesssim 10$ MeV can escape the corona completely unhindered. The shape and, in particular, the cutoff energy of the cascaded γ -ray spectrum depend strongly on the emission radius. Smaller values of r_c lead to lower cutoff energies and less γ rays at energies above ~ 10 MeV due to the higher density of target photons for $\gamma\gamma$ interactions at smaller emission radii. In addition to η and the total energy injected into proton acceleration, r_c is, therefore, an important model

parameter that shapes the multimessenger spectrum of a source with a given x-ray luminosity.

III. MODEL FIT FOR NGC 1068

In order to constrain the remaining free parameters of the neutrino emission model presented in Sec. II, we fit the neutrino spectrum for NGC 1068 to publicly available neutrino data. Moreover, we compare the resulting cascaded γ -ray emission to existing upper limits in the MeV–GeV range to ensure that our model predictions are consistent with both neutrino and γ -ray observations of NGC 1068.

A. Fit of the neutrino spectrum

Besides the coronal radius r_c , the only free parameters of our model are the inverse turbulence strength of the coronal magnetic field, η , and the total energy injected into proton acceleration, expressed in terms of the pressure ratio $P_{\text{CR}}/P_{\text{th}}$. The neutrino spectrum of a source is primarily determined by the latter two parameters. To constrain them, we fit our model neutrino spectrum for NGC 1068 to the public 10-year IceCube dataset of tracklike events for point-source searches [61], which consists of a list of muon events with reconstructed energies $\hat{E}$ and directions $\hat{\theta}$. We do this by performing a fit with an unbinned likelihood function using the open-source software `skyLLH` [62–64]. The likelihood function

$$\mathcal{L}(n_s, \eta) = \prod_{i=1}^N \left[\frac{n_s}{N} \mathcal{S}(\hat{E}_i, \hat{\theta}_i | \eta) + \left(1 - \frac{n_s}{N} \right) \mathcal{B}(\hat{E}_i, \hat{\theta}_i) \right] \quad (21)$$

is a weighted sum of the signal probability density $\mathcal{S}(\hat{E}_i, \hat{\theta}_i | \eta)$, which gives the probability that the i th event is a signal event associated with the source, and the background probability density $\mathcal{B}(\hat{E}_i, \hat{\theta}_i)$, which gives the probability that the i th event is a background event. Here, n_s denotes the number of signal events, and N is the total number of events in the dataset. To obtain the values of η and n_s for which our model for NGC 1068 fits the data best, we first compute the neutrino spectra for the x-ray luminosity of NGC 1068 for $\eta = 1, 2, 3, \dots, 150$. Then, we determine the values of η and n_s that maximize the likelihood given in Eq. (21). In doing so, we impose bounds on n_s to ensure that $P_{\text{CR}}/P_{\text{th}} \leq 0.5$ is satisfied at all times.

For a dimensionless coronal radius of $r_c = 5$, our fit yields $\hat{n}_s = 46.4$ signal events and an inverse turbulence strength of $\hat{\eta} = 56$, which corresponds to a total proton luminosity of $L_p = 1.43 L_X \simeq 6 \times 10^{43} \text{ erg/s} \simeq 0.009 L_{\text{Edd}}$ and a pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$. This is consistent with Refs. [16,17,20,29,65], all of which report that $L_p \simeq L_X$ and $P_{\text{CR}}/P_{\text{th}} \simeq 0.5$ is necessary to explain the observed neutrino flux of NGC 1068. Figure 3 shows the log-likelihood ratio




FIG. 2. Cascaded γ -ray spectra for different emission radii calculated for the x-ray luminosity of NGC 1068 ($L_X = 4.17 \times 10^{43} \text{ erg/s}$) and $\eta = 20$. Here, the underlying proton spectra are normalized so that $L_p/L_X = 0.5$. $E_\gamma L_{E_\gamma} \equiv E_\gamma^2 dN_\gamma/dE_\gamma dt$ denotes the differential γ -ray luminosity.




FIG. 3. Log-likelihood ratio as a function of n_s and η . The best-fit parameters obtained for NGC 1068 are indicated by a red cross, and the gray contours represent the 1σ and 2σ confidence regions according to Wilks' theorem [66]. The region above the red line corresponds to $P_{\text{CR}}/P_{\text{th}} > 0.5$ and is, thus, excluded.

$$-2 \times \Delta \log \mathcal{L} \equiv -2 \log \frac{\mathcal{L}(n_s, \eta)}{\mathcal{L}(\hat{n}_s, \hat{\eta})} \quad (22)$$

as a function of the parameters n_s and η . By setting an upper limit on $P_{\text{CR}}/P_{\text{th}}$, we restrict the allowed values of n_s and η to only part of the considered parameter space. However, when performing the same fit without imposing a constraint on the pressure ratio, we find that our result obtained with $P_{\text{CR}}/P_{\text{th}} \leq 0.5$ still lies within the 1σ confidence region of the fit result obtained without an upper limit on $P_{\text{CR}}/P_{\text{th}}$. This indicates that requiring $P_{\text{CR}}/P_{\text{th}} \leq 0.5$ does not result in a significantly worse fit to the data.

Our best-fit model neutrino spectrum for NGC 1068 is shown in Fig. 4, together with the 1σ and 2σ uncertainty bands. The latter correspond to the confidence regions marked in Fig. 3 and are calculated assuming Wilks' theorem [66] as implemented in SkyLLH. The asymmetrical shape of the uncertainty bands is due to the fact that the best-fit values for n_s and η lie at the boundary of the parameter space. We find that both pp and $p\gamma$ interactions are important for the production of neutrinos. Around the peak of the spectrum at ~ 1.5 TeV, most of the neutrino emission is produced in $p\gamma$ interactions, while at lower energies the spectrum is dominated by neutrinos from pp interactions. The sensitive energy range of our analysis, defined as the central 68% interval of neutrino energies contributing to the excess over the background [67–69], extends from 350 GeV to 190 TeV. A more detailed discussion of the method used to determine this energy range, as well as a comparison of our result to the best-fit power law spectrum of NGC 1068, can be found in Appendix B.

Similar to Ref. [71], we use posterior predictive checks to evaluate the goodness of fit. Using SkyLLH, we simulate 10000 pseudodatasets assuming our best-fit model neutrino




FIG. 4. Best-fit all-flavor neutrino spectrum of NGC 1068 (red) together with the corresponding cascaded γ -ray spectrum (blue). The shaded bands indicate the 1σ and 2σ uncertainty regions of the flux and, for the neutrino spectrum, the central 68% energy range. The latest Fermi-LAT data of NGC 1068 [70] is shown for comparison.

spectrum and select all events in a circular region with radius $r = 5^\circ$ around the position of NGC 1068. In 24 logarithmically spaced energy bins between 100 GeV and 1 PeV, we quantify the discrepancy between the simulated and the observed data and find that it is smaller than 2σ in all bins. This indicates a good agreement between our model and the data.

B. Cascaded γ -ray spectrum of NGC 1068

To ensure that our model predictions do not violate γ -ray observations, we also calculate the cascaded γ -ray spectrum that corresponds to the best-fit neutrino spectrum of NGC 1068. The resulting spectrum is shown in Fig. 4 together with the 1σ and 2σ uncertainty bands calculated from the uncertainty bands around the best-fit neutrino spectrum. Comparing our result to the latest Fermi-LAT observations of NGC 1068 [70], we find that, for a dimensionless coronal radius of $r_c = 5$, our predicted γ -ray spectrum agrees well with the observational data. Especially the 95% upper limit at ~ 30 MeV, which provides the strongest constraint, lies within the 1σ uncertainty band.

This is in agreement with the results obtained by Das, Zhang, and Murase [29], who find $R_c \lesssim (3-15)R_S$ considering a purely photohadronic scenario. A similar study using a previously reported γ -ray spectrum of NGC 1068 that extends only down to ~ 70 MeV [72] yields somewhat looser constraints on the coronal radius with $R_c \lesssim 30R_S$ for a $p\gamma$ -only scenario and $R_c \lesssim 100R_S$ for a pp -only scenario [32]. In addition, various indirect observational methods that are used to determine the size of AGN coronae also point toward rather small and compact coronae with a typical radius of $R_c \lesssim 5R_S$ [73]. The fact that the observed γ -ray spectrum of NGC 1068 at energies $\gtrsim 500$ MeV cannot be explained by hadronic interactions

within the corona is consistent with the results of multi-messenger studies [17,70,74], which suggest that these γ rays are mainly produced in the starburst region. It is expected that the starburst contributes significantly only at energies above ~ 100 MeV [70]. However, if there was another additional source of γ rays with energies $\lesssim 100$ MeV, e.g., due to primary electron acceleration, an even smaller coronal radius or nonuniform γ -ray absorption in the corona would be necessary in order not to exceed the Fermi-LAT observations [19].

IV. SOURCE POPULATION

We apply the neutrino emission model described in Sec. II to an entire population of AGNs, taking into account the fit result for NGC 1068, with the aim of calculating the resulting diffuse neutrino flux. The cosmological evolution of this population is described by the x-ray luminosity function (XLF), which gives the differential number density of AGNs per comoving volume as a function of x-ray luminosity and redshift. Here, we consider three different 2–10 keV XLFs.

The luminosity function by Ueda *et al.* presented in Ref. [75] (hereafter U14) is derived assuming that the XLF in the local Universe can be modeled by a smoothly broken power law, which is then multiplied by a luminosity-dependent evolution factor to obtain the XLF at a given redshift $z > 0$. The free parameters of this function are determined by a fit to the source sample. Heavily obscured Compton-thick AGNs with column

densities $N_H \geq 1/\sigma_T \simeq 10^{24} \text{ cm}^{-2}$ are assumed to follow the same cosmological evolution as obscured Compton-thin AGNs with $10^{22} \text{ cm}^{-2} \leq N_H \leq 10^{24} \text{ cm}^{-2}$, which is in agreement with observations of the cosmic x-ray background.

In contrast, Buchner *et al.* [76] (hereafter B15) use a nonparametric approach to derive the XLF. They fit a 3D histogram in L_X , z , and N_H to their source sample with the only assumption being that the XLF does not vary rapidly between neighboring bins. To ensure this, two different smoothness priors are considered: the constant-slope prior and the constant-value prior. The former keeps power-law slopes intact, while the latter retains the current value of the XLF unless constraints are imposed by the data. Moreover, B15 constrain the contribution from Compton-thick AGNs in a more precise way. For each source in their sample, they fit a spectral model to determine the obscuring column density and the intrinsic x-ray luminosity consistently taking into account different sources of uncertainties. In this way, they do not have to rely on a fit to the cosmic x-ray background, which is rather insensitive to the exact fraction of Compton-thick AGNs [77]. Although Compton-thick sources do not appear very luminous due to their high level of obscuration, they are just as important as less obscured AGNs, since in our model the neutrino emission properties of a source do not depend on N_H and, unlike x-rays, neutrinos are not absorbed by the dusty torus surrounding the central region of an AGN.

Figure 5 shows the three considered XLFs in different redshift bins between $z = 0$ and 5 together with their




FIG. 5. 2–10 keV x-ray luminosity functions of AGNs from Ref. [75] (U14) and Ref. [76] (B15). The red shaded bands correspond to the 1σ and 2σ uncertainty regions of the U14 XLF derived from the statistical errors of the model parameters assuming that all parameters are uncorrelated; i.e., these bands indicate upper limits on the size of the actual uncertainty regions. The blue and green shaded bands indicate the 68% and 95% credible regions of the B15 XLFs for the constant-slope and constant-value prior, respectively.

TABLE III. Properties of the source populations simulated based on the three different XLFs and the sources selected from the BASS catalog. Given are the number of sources with an intrinsic x-ray flux higher than 10^{-11} erg cm $^{-2}$ s $^{-1}$ (column 2) and higher than that of NGC 1068 (column 3), the total number of simulated sources in the population (column 4), and the total intrinsic x-ray flux from the population after replacing all sources with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ by sources from the BASS catalog (column 5).

	Number of sources with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$	Number of sources with $F_X \geq F_X^{\text{NGC 1068a}}$	Number of all sources	F_X^{tot} [erg cm $^{-2}$ s $^{-1}$ sr $^{-1}$]
U14	363	0	5.1×10^8	7.9×10^{-8}
B15 const. slope	865	0	1.3×10^9	9.1×10^{-8}
B15 const. value	28471	46	3.4×10^8	1.2×10^{-7}
BASS catalog	212	1	...	...

^aHere, we assume an intrinsic x-ray flux of $F_X^{\text{NGC 1068}} = 2.80 \times 10^{-9}$ erg cm $^{-2}$ s $^{-1}$ for NGC 1068, which corresponds to $L_X = 4.17 \times 10^{43}$ erg/s and $d_L = 11.14$ Mpc [9,10].

uncertainties. All three XLFs have roughly the shape of a broken power law with significantly more low-luminosity sources than high-luminosity sources. However, they do not agree within their uncertainties. While the U14 XLF and the B15 XLF for the constant-slope prior have an overall similar shape, especially at intermediate redshifts, the B15 XLF for the constant-value prior is much flatter and, thus, predicts many more high-luminosity sources with $L_X \gtrsim 10^{45}$ erg/s. The discrepancy between the two XLFs by B15, particularly at the high-luminosity end, is due to the fact that the underlying x-ray source sample consists mainly of data from deep x-ray surveys with small survey areas. These surveys are not suited to accurately determine the number density of high-luminosity AGNs, since they are very rare. Thus, the shape of the XLFs at high luminosities is primarily determined by the underlying prior and not by the data itself. However, as we show in Sec. VA, the choice of the XLF ultimately has only a minor impact on the resulting diffuse neutrino flux.

For each XLF, we simulate a population of AGNs together with their intrinsic x-ray luminosities. It is particularly important to take into account that the XLF changes not only its normalization, but also its shape with redshift. Otherwise, the number of high-luminosity sources at higher redshifts will be severely underestimated. To this end, we consider 14 redshift bins between $z = 0$ and 5 and fit broken power laws to the XLFs in each bin. We then use the population synthesis framework POPSYNTH [78] to simulate three source populations based on the three different XLFs. For each simulated source, the redshift and the intrinsic x-ray luminosity are drawn from the respective distributions.

To ensure a correct description of the number density of bright Seyfert galaxies in the local Universe, we compare our simulated source populations to the BAT AGN Spectroscopic Survey (BASS) catalog [11]. This is an x-ray catalog of AGNs detected in the 14–195 keV band, which has also been used in previous works to predict and search for neutrino emission from nearby Seyfert galaxies [4,5,20]. Table III lists the number of bright

sources with an intrinsic 2–10 keV x-ray flux of $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ and the number of sources with an intrinsic x-ray flux as high as or higher than that of NGC 1068 in each simulated population and in the BASS catalog. The U14 XLF and the B15 XLF for the constant-slope prior predict the same order of magnitude of bright, nearby Seyfert galaxies as contained in the BASS catalog. The B15 XLF for the constant-value prior, on the other hand, predicts more than 2 orders of magnitude more sources with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$, including 46 sources that are as bright as or even brighter than NGC 1068. As mentioned above, this discrepancy can be attributed to the fact that the x-ray sample used to construct the B15 XLFs is not sensitive to rare sources. In line with the findings by Ananna *et al.* [79], we conclude that the B15 XLF for the constant-value prior likely overestimates the abundance of high-luminosity sources and, therefore, does not provide an accurate description of the number density of bright Seyfert galaxies in the local Universe.

To correct for this effect and model the contribution of nearby bright sources in a more realistic way, we select all Seyfert galaxies with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ from the BASS catalog and add them to our simulated populations. Simultaneously, we remove all simulated sources above this flux threshold to avoid double counting of sources. Overall, we consider 212 Seyfert galaxies from the BASS catalog, including NGC 1068, NGC 4151, and CGCG 420-015. The total number of sources as well as the total intrinsic x-ray flux from all sources in the resulting populations are listed in Table III. The population based on the B15 XLF for the constant-slope prior comprises the largest number of sources, while the population based on the B15 XLF for the constant-value prior produces the highest intrinsic x-ray flux.

V. DIFFUSE NEUTRINO FLUX

To estimate the contribution of Seyfert galaxies to the diffuse astrophysical neutrino flux, we apply our neutrino emission model to each individual source in the simulated

populations, assuming that the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η are the same for all sources. First, we fix these two parameters to the best-fit values found for NGC 1068 and study the effect of the different luminosity functions on the resulting diffuse neutrino flux. We then explore how the diffuse flux spectrum depends on the model parameters and quantify the agreement between our diffuse flux predictions and the latest IceCube observations as a function of $P_{\text{CR}}/P_{\text{th}}$ and η . This allows us to derive constraints on these parameters and to assess the exceptional nature of NGC 1068 in the context of the entire population of Seyfert galaxies.

A. Diffuse neutrino flux for different XLFs

To compute the diffuse neutrino flux from a population of AGNs, we apply the neutrino emission model presented in Sec. II to each source in the population and calculate the neutrino flux taking into account energy losses due to the adiabatic expansion of the Universe. The neutrino flux at Earth from a source located at redshift z is given by

$$\frac{dN_\nu}{dE_\nu dAdt}(E_\nu) = \frac{(1+z)^2}{4\pi d_L^2(z)} \frac{dN_\nu}{dE_\nu dt}[E_\nu(1+z)], \quad (23)$$

where $dN_\nu/dE_\nu dt$ is the differential neutrino emission at the location of the source given by our neutrino emission model and $d_L(z)$ is the luminosity distance calculated using the approximation from Ref. [80]. To obtain the diffuse neutrino flux per solid angle from the entire population, the neutrino flux spectra of all sources are summed up and divided by 4π assuming that the sources are isotropically distributed across the sky:

$$\Phi_\nu(E_\nu) = \frac{1}{4\pi} \sum_{i=1}^N \left(\frac{dN_\nu}{dE_\nu dAdt} \right)_i(E_\nu). \quad (24)$$

Here, i runs over all sources in the population, and $(dN_\nu/dE_\nu dAdt)_i$ is the neutrino flux from the i th source.

First, we compute the diffuse neutrino flux assuming that all sources have the same values of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η that we found for NGC 1068 in Sec. III and a coronal radius of $r_c = 5$, since this is the most straightforward way to extrapolate our findings for NGC 1068 to the entire source population. The diffuse flux spectra obtained for the three different XLFs with $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ for all sources are shown in Fig. 6. The diffuse neutrino flux measured by IceCube in different analyses using starting tracks and the enhanced starting track event selection (ESTES) [81], cascade events [82], and muon tracks from the Northern Hemisphere [83] are displayed for comparison.

Although the underlying XLFs shown in Fig. 5 appear quite different, the resulting diffuse neutrino flux spectra are remarkably similar in both shape and overall




FIG. 6. Diffuse all-flavor neutrino flux calculated for the three different XLFs assuming $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ for all sources. The black, dark gray, and light gray points correspond to the diffuse neutrino flux measured by IceCube in different analyses assuming an E^{-2} power law in each energy bin [81–83]. We assume a flavor ratio of $\nu_e : \nu_\mu : \nu_\tau = 1 : 1 : 1$ to convert the per-flavor fluxes reported by IceCube to all-flavor fluxes.

normalization. They have a slight peak at ~ 1 TeV and drop off toward higher energies, with the contribution to the total observed neutrino flux becoming negligible at $E_\nu \gtrsim 10$ TeV. The main reason for this similarity is that, after replacing all sources with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ by sources from the BASS catalog, the total intrinsic x-ray fluxes of the three populations listed in Table III differ by less than a factor of 2, and the neutrino luminosity increases with the x-ray luminosity of a source as described by Eq. (20). The source population based on the B15 XLF for the constant-slope prior produces a slightly higher diffuse neutrino flux than the other two populations. This is due to the fact that this population contains the largest number of low-luminosity sources with $L_X \lesssim 10^{43}$ erg/s and, under the assumption that $P_{\text{CR}}/P_{\text{th}}$ is constant across sources of all luminosities, low-luminosity sources have a higher neutrino luminosity relative to their x-ray luminosity than high-luminosity sources. This means that the diffuse neutrino flux is determined not only by the total x-ray flux of the population, but also by the distribution of the source luminosities.

From Fig. 6, it can also be seen that our diffuse flux predictions for all three XLFs exceed both the 68% upper limit at ~ 2 TeV from the IceCube ESTES analysis and the flux in the lowest energy bin of the IceCube cascade analysis at ~ 7 TeV. This indicates that it is very unlikely that all sources share the same model parameters as NGC 1068. Especially the assumption that the pressure ratio inside the coronae of all sources reaches the maximum




FIG. 7. Diffuse all-flavor neutrino flux from sources in different redshift bins between $z = 0$ and 5 for the U14 XLF (left panel), the B15 XLF for the constant-slope prior (middle panel), and the B15 XLF for the constant-value prior (right panel) assuming $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ for all sources. Lighter colors correspond to lower redshifts, and the total diffuse flux from all redshift bins together is shown as a gray line.

physically allowed value of $P_{\text{CR}}/P_{\text{th}} = 0.5$ (see discussion in Sec. II B) is quite extreme. Thus, the spectra shown in Fig. 6 are rather to be understood as upper limits on the possible neutrino emission from the population under the assumption that all sources share the same value of η as NGC 1068. A more detailed comparison of our diffuse flux predictions for different values of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η to the IceCube observations is performed in Sec. V B.

Another interesting aspect that we investigate is the extent to which sources in different redshift and luminosity bins contribute to the total diffuse neutrino flux from the population. Figure 7 shows the relative contribution from sources in different redshift bins to the diffuse flux. For all three XLFs, the bulk of the diffuse flux originates from sources at redshifts $z \sim 0.1$ – 2.0 , with sources at higher redshifts contributing to the diffuse flux at lower energies since their neutrino spectra are shifted to lower energies due to the adiabatic expansion of the Universe. At 1 TeV, nearby sources with redshifts $z < 0.1$ account for only $\sim 5\%$ of

the diffuse flux due to their limited number. The contribution from the BASS catalog sources, by which we replace all simulated sources with $F_X \geq 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$, is even smaller with just $\sim 1\%$ at 1 TeV. Because of their large distances and the fact that the x-ray luminosity density of the AGN population peaks at $z \sim 2$ and decreases toward higher redshifts, sources with $z \gtrsim 2$ do not contribute significantly to the diffuse neutrino flux either.

Figure 8 shows the relative contribution from sources in different x-ray luminosity bins to the diffuse neutrino flux. Considering the populations simulated based on the U14 XLF and the B15 XLF for the constant-slope prior, most of the diffuse flux can be attributed to sources with luminosities in the range $10^{42} \text{ erg/s} \lesssim L_X \lesssim 10^{44} \text{ erg/s}$. Looking at the population based on the B15 XLF for the constant-value prior, also brighter sources with luminosities up to $L_X \sim 10^{45} \text{ erg/s}$ provide a significant contribution, because they are so abundant in the population. Under the assumption that all sources are described by the same model parameters as NGC 1068, we conclude that the bulk of the diffuse




FIG. 8. Diffuse all-flavor neutrino flux from sources in different x-ray luminosity bins between $\log L_X = 42$ and 46 for the U14 XLF (left panel), the B15 XLF for the constant-slope prior (middle panel), and the B15 XLF for the constant-value prior (right panel) assuming $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ for all sources. Lighter colors correspond to lower luminosities, and the total diffuse flux from all luminosity bins together is shown as a gray line.

neutrino flux originates from a large population of very faint neutrino sources and the contribution from nearby Seyfert galaxies is subdominant.

B. Comparison to IceCube observations

To constrain our neutrino emission model also from the population side, we quantify the agreement between the diffuse flux predictions and the IceCube measurements as a function of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η , still assuming that all sources in the population share the same parameter values. The dependence of the shape and the normalization of the diffuse neutrino spectrum on $P_{\text{CR}}/P_{\text{th}}$ and η is very similar to that of the single-source neutrino spectra shown in Fig. 1. The parameter η determines the shape of the spectrum, in particular, the cutoff energy, with larger values of η leading to lower cutoff energies, while the pressure ratio determines the normalization. We compare our results to the segmented E^{-2} power-law spectra obtained by IceCube in the analyses using ESTES events [81] and cascade events [82]. The sensitive energy ranges of both analyses extend down to TeV energies, making them well suited to constrain our model. We perform the comparisons to the two IceCube spectra separately, as they differ from each other at energies below ~ 10 TeV (see also Fig. 6). For the following part of our analysis, we consider only the diffuse flux spectra calculated for the source population based on the U14 XLF, since the results for the three different XLFs do not differ significantly anyway.

We compute the diffuse neutrino flux for a total of 100 values of η equally spaced between 1 and 100. In doing so, we consider only sources with $L_X \leq 10^{45.5}$ erg/s after confirming that the contribution from sources with higher x-ray luminosities is negligible ($< 0.25\%$) over the entire energy range.¹ The reason for this luminosity cut is that calculating the neutrino spectra of sources with high x-ray luminosities is very computationally expensive due to the large size of the corona and the time resolution required for the simulation to capture all relevant interaction processes.

In order to compare our diffuse flux spectra to the segmented E^{-2} power-law fluxes reported by IceCube, we calculate the average all-flavor neutrino flux in the i th energy bin as

$$\Phi_i = \frac{1}{E_{i+1} - E_i} \int_{E_i}^{E_{i+1}} dE_\nu E_\nu^2 \Phi_\nu(E_\nu), \quad (25)$$

where $\Phi_\nu(E_\nu)$ is the neutrino flux from the population given by Eq. (24) and E_i and E_{i+1} are the lower and the upper bound of the energy bin, respectively. For each bin, we define a likelihood of the form

$$\mathcal{L}(\Phi_i | \Phi_{\text{obs},i}, \sigma_i) \propto \begin{cases} \frac{1}{\sigma_i \sqrt{2\pi}} & \text{if } \Phi_i < \Phi_{\text{obs},i} \\ \mathcal{N}(\Phi_i | \Phi_{\text{obs},i}, \sigma_i) & \text{if } \Phi_i \geq \Phi_{\text{obs},i} \end{cases} \quad (26)$$

that is constant for $\Phi_i < \Phi_{\text{obs},i}$ and has the shape of a Gaussian with

$$\mathcal{N}(x | \mu, \sigma) \equiv \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (27)$$

for $\Phi_i \geq \Phi_{\text{obs},i}$. The flux and the corresponding 68% uncertainty in the i th energy bin reported by IceCube are denoted by $\Phi_{\text{obs},i}$ and σ_i , respectively, where we assume a flavor ratio of $\nu_e : \nu_\mu : \nu_\tau = 1 : 1 : 1$ to convert per-flavor fluxes to all-flavor fluxes. The likelihood for a total of N energy bins is then given by

$$\mathcal{L}(\Phi | \Phi_{\text{obs}}, \sigma) = \prod_{i=1}^N \mathcal{L}(\Phi_i | \Phi_{\text{obs},i}, \sigma_i) \quad (28)$$

with $\Phi \equiv (\Phi_1, \dots, \Phi_N)$, $\Phi_{\text{obs}} \equiv (\Phi_{\text{obs},1}, \dots, \Phi_{\text{obs},N})$, and $\sigma \equiv (\sigma_1, \dots, \sigma_N)$. Upper limits are treated as flux measurements with $\Phi_{\text{obs},i} = 0$ and σ_i as reported in the respective publications.

Figure 9 shows the log-likelihood ratio

$$-2 \times \Delta \log \mathcal{L} \equiv -2 \log \frac{\mathcal{L}(\Phi | \Phi_{\text{obs}}, \sigma)}{\mathcal{L}(\Phi = 0 | \Phi_{\text{obs}}, \sigma)}, \quad (29)$$

where $\Phi \equiv \Phi(P_{\text{CR}}/P_{\text{th}}, \eta)$ is a function of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η . The regions of the parameter space that lie above the gray and black contours are excluded by the IceCube measurements with a significance of 1σ and 3σ , respectively. These exclusion contours correspond to $-2 \times \Delta \log \mathcal{L} = 1$ and 9 according to Wilks' theorem [66], the use of which is justified since we assume that the uncertainties on the segmented power-law fluxes reported by IceCube are Gaussian. The red solid and dashed contours correspond to the 1σ and 2σ contours, respectively, shown in Fig. 3 with the signal event number n_s converted to $P_{\text{CR}}/P_{\text{th}}$. For the visualization in Fig. 9, we use a polynomial fit to smooth the contours that are calculated on a finite grid of values for n_s and η . The sharp cutoff at $P_{\text{CR}}/P_{\text{th}} = 0.5$ is due to the fact that, for higher pressure ratios, the AGN corona is no longer stable, as explained in Sec. II B.

¹To verify this, we have calculated the diffuse neutrino flux from the full population, including sources with $L_X > 10^{45.5}$ erg/s, on a sparser grid of η values ($\eta = 1, 20, 40, 60, 80, 100$), and compared the result to the flux obtained when considering only sources with $L_X \leq 10^{45.5}$ erg/s. Changing the pressure ratio does not affect the relative contribution of sources with different x-ray luminosities, but only the overall normalization of the diffuse flux spectrum, since we assume a common value of $P_{\text{CR}}/P_{\text{th}}$ for all sources in the population.




FIG. 9. Log-likelihood ratio as a function of $P_{\text{CR}}/P_{\text{th}}$ and η for the comparison of our model predictions for the diffuse neutrino flux to the IceCube measurements using ESTES (top panel) and cascade (bottom panel) events. The gray and black contours indicate the 1σ and 3σ exclusion regions, respectively. The red solid and dashed contours represent the 1σ and 2σ confidence regions, respectively, from the fit of our model neutrino spectrum for NGC 1068 to the public IceCube data, and the red cross corresponds to the best-fit result of $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$.

Comparing our diffuse flux predictions based on the neutrino emission model presented in Sec. II to the IceCube observations, we can exclude scenarios in which all sources in the population have a high pressure ratio in combination with a high level of turbulence inside the corona, i.e., a high value of $P_{\text{CR}}/P_{\text{th}}$ together with a small value of η . Because of the upper limit at ~ 2 TeV and the shape of the diffuse neutrino spectra predicted by our model, the IceCube ESTES measurement is more constraining than the one based on cascade events. As can be seen in the upper panel in Fig. 9, the best-fit result for NGC 1068 indicated by a red cross lies outside the 3σ contour. If all sources in the population had the same model parameters as NGC 1068, the resulting diffuse neutrino flux would exceed the diffuse flux from the ESTES analysis by 3.8σ , echoing our findings




FIG. 10. Diffuse all-flavor neutrino flux for $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ (thick dark red line) together with discrete shaded spectra indicating the 1σ and 2σ confidence regions of the fit for NGC 1068. Shown is a selection of spectra for parameter values that lie within the red solid and dashed contours plotted in Fig. 9. The black and dark gray points correspond to the diffuse flux measured by IceCube using starting-track [81] and cascade [82] events, respectively.

in Fig. 6. Looking at the contours that indicate the 1σ and 2σ confidence regions of the fit for NGC 1068, it becomes apparent that it is challenging to find values of $P_{\text{CR}}/P_{\text{th}}$ and η for which the resulting diffuse flux spectrum agrees with the IceCube ESTES flux while the single-source neutrino spectrum for NGC 1068 still provides a good fit to the public IceCube data. This is also illustrated in Fig. 10, which shows a selection of diffuse neutrino spectra for parameter values within the 1σ and 2σ confidence regions of the fit result for NGC 1068 marked in Fig. 9. Here, the upper limit at ~ 2 TeV lies just at the edge of the 2σ band.

As shown in the bottom panel in Fig. 9, the constraints imposed by the diffuse flux spectrum derived from cascade events are weaker, mainly because its sensitive energy range starts only at ~ 7 TeV. Our model spectrum for the best-fit parameters obtained for NGC 1068 exceeds the observed flux by only 1.4σ . In addition, there is a range of parameter values for which our model predictions for both the single-source spectrum of NGC 1068 and the diffuse flux are in agreement with the observations. This can also be seen in Fig. 10, where the segmented power-law flux derived from cascade events lies within or above the 1σ band in all energy bins.

Here, we compare our results only to the IceCube measurements using starting-track and cascade events and do not perform a similar comparison to the one using through-going muon tracks [83]. The sensitive energy range of this analysis only starts at 15 TeV, but our diffuse neutrino spectra peak at energies $\lesssim 10$ TeV. Therefore, the diffuse flux spectrum derived from through-going muon tracks is not suited to constrain our model. At energies

below ~ 100 TeV, the result of the combined fit using cascades and muon tracks presented in Ref. [84] is very similar to the spectrum obtained in the cascades-only analysis, since both analyses use the same underlying cascade event sample and the cascade events drive the combined fit at low energies. Thus, comparing our diffuse flux predictions to the combined fit spectrum would yield a result almost identical to that of the comparison to the spectrum derived from cascade events only.

VI. DISCUSSION

Our results show that the neutrino observations of NGC 1068 are well described by our neutrino emission model, which is based on the disk-corona model presented in Ref. [16]. However, extrapolating our model to the entire population of Seyfert galaxies, assuming that all sources share the same parameters as NGC 1068, yields a diffuse neutrino flux that exceeds current upper limits at TeV energies. This suggests that only a small fraction of Seyfert galaxies, at most $\sim 25\%$,² can be as efficient neutrino emitters as NGC 1068.

A. Coronal properties of NGC 1068

As shown in Sec. III A, the best-fit model parameters obtained for NGC 1068 lie at the upper edge of the physically allowed parameter space. For the normalization of the neutrino spectrum, our fit yields the maximum allowed value for the pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$, indicating that the corona of NGC 1068 must host an extremely efficient hadronic particle accelerator for our model to explain the observed neutrino emission (see also discussion in Sec. II B). This is in agreement with the results of previous studies, which compare the predictions of different neutrino emission models to the IceCube power-law spectrum of NGC 1068 and find similarly high pressure ratios [16,17,20,29,65]. Abbasi *et al.* [4] fit the model for NGC 1068 and also other nearby Seyfert galaxies presented in Ref. [20] to the proprietary IceCube data consisting of 11 years of through-going muon tracks from the northern sky. Fitting only the normalization and keeping the overall shape of the neutrino spectrum fixed, they also find a pressure ratio of $P_{\text{CR}}/P_{\text{th}} \simeq 0.5$ for NGC 1068.

The size of the corona is strongly constrained by the Fermi-LAT observations of NGC 1068 [70]. As shown in Sec. III B, in order for our neutrino emission model for NGC 1068 not to exceed the upper limit on the γ -ray flux at ~ 30 MeV, we need a coronal radius of $R_c \lesssim 5R_S$. This is consistent with the findings of Das *et al.* [29], who obtain $R_c \lesssim (3-15)R_S$ depending on the magnetic field strength inside the corona. In general, such small and compact

coronae are consistent with the results of various observational methods used to estimate the size of AGN coronae [73]. Thus, not only neutrino data, but also γ -ray observations in the MeV range are essential for a better understanding of AGN coronae as possible neutrino production sites.

B. Description of the number density of x-ray bright AGNs

As described in Sec. IV, the predictions of the U14 XLF and the B15 XLF for the constant-slope prior agree with the BASS catalog within one order of magnitude (see also Table III). Consequently, replacing the simulated bright sources with sources from the BASS catalog has a negligible effect on the resulting diffuse neutrino flux. This is not the case for the B15 XLF for the constant-value prior, which predicts more than 2 orders of magnitude more sources with $F_X \geq 10^{-11}$ erg cm $^{-2}$ s $^{-1}$ than are present in the BASS catalog. In particular, the simulated population contains 46 sources that are as bright or brighter than NGC 1068, some even with intrinsic x-ray fluxes of $F_X \geq 10^{-8}$ erg cm $^{-2}$ s $^{-1}$. Assuming that these sources are described by similar model parameters as NGC 1068, at least some of them should have already been discovered by IceCube. Since this is not the case, we conclude that, of the two XLFs derived by Buchner *et al.*, the one for the constant-slope prior provides a more realistic description of the number density of AGNs, especially at low redshifts and for high x-ray luminosities. This is consistent with the findings of Refs. [76] and [79].

C. Implications of our diffuse flux predictions

Our results for all three XLFs indicate that Seyfert galaxies can account for the diffuse astrophysical neutrino flux at energies below ~ 10 TeV without exceeding their energy budget. To explain the observed diffuse neutrino flux at higher energies, at least one additional population of neutrino sources, such as starburst galaxies [85,86] or AGN outflows [87,88], is needed. A study of different populations of candidate neutrino sources presented in Ref. [89] comes to the conclusion that low-luminosity AGNs, misaligned radio-loud AGNs, and galaxy clusters could also contribute significantly to the diffuse neutrino flux. The possible contribution from blazars, on the other hand, is limited to $\lesssim 10\%$ at 100 TeV [90,91]. For further details on the contribution of blazars to the astrophysical neutrino flux, see [92] and references therein.

Evaluating how well our model predictions agree with the diffuse flux measured by IceCube, we obtain detailed constraints on a neutrino emission model for Seyfert galaxies based on both observations of NGC 1068 as a nearby point source candidate and the cumulative neutrino emission of the entire source population. Our results show that if all sources were described by the same model parameters as NGC 1068, their neutrino emission would

²Under the assumption that the remaining sources do not emit any neutrinos, up to $\sim 25\%$ of all sources could have the same model parameters as NGC 1068 without the resulting diffuse neutrino flux exceeding the IceCube observations by more than 1σ .




FIG. 11. Best-fit all-flavor neutrino spectra of NGC 1068 (left panel), NGC 4151 (middle panel), and CGCG 420-015 (right panel) obtained using skyLLH. The shaded bands correspond to the 1σ and 2σ uncertainty regions. For CGCG 420-015, the 2σ band extends down to a neutrino flux of zero, reflecting that the neutrino excess associated with this source is not statistically significant and a scenario with no associated signal events ($n_s = 0$) is consistent with the IceCube data within 2σ . The corresponding best-fit values of the model parameters are listed in Table IV. For each source, the gray line indicates the maximum allowed neutrino emission for which the corresponding diffuse flux would still be compatible with the IceCube observations within 1σ .

significantly exceed current diffuse flux measurements at TeV energies, in particular, the upper limit at ~ 2 TeV from the IceCube ESTES analysis [81].

For the population of Seyfert galaxies to account for the bulk of the observed diffuse flux below ~ 100 TeV without violating this upper limit, the majority of sources would have to have a much lower pressure ratio of $P_{\text{CR}}/P_{\text{th}} \sim 0.01$ and a very high level of turbulence with $\eta \sim 1$. However, as indicated by the red solid and dashed contours in Fig. 9, the corresponding single-source neutrino spectrum for NGC 1068 is not compatible with the IceCube data. It is generally not possible to find values of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η for which our model provides a good fit to the public IceCube data for NGC 1068 without significantly overproducing the observed diffuse neutrino flux. Thus, we conclude that NGC 1068 is an extraordinarily powerful Seyfert galaxy and that most Seyfert galaxies in the Universe must be significantly less efficient neutrino emitters.

A similar conclusion can be drawn for two other nearby Seyfert galaxies, NGC 4151 and CGCG 420-015, for which the IceCube Collaboration has also reported evidence for neutrino emission [4,5]. In addition, neutrino and γ -ray emission from NGC 4151 was studied by [93].

Following the procedure described in Sec. III A for NGC 1068, we fit our neutrino emission model for these two sources to the public IceCube data. The best-fit neutrino spectra of NGC 4151 and CGCG 420-015 shown in Fig. 11 peak at slightly higher energies than that of NGC 1068 but have a lower overall normalization. The obtained number of signal events and the corresponding best-fit parameter values are listed in Table IV. For both sources, our fit yields the maximum allowed value of the pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$ but smaller values of η than for NGC 1068, namely, $\eta = 4$ for NGC 4151 and $\eta = 1$ for CGCG 420-015.

However, as shown in Fig. 9, the resulting diffuse neutrino flux for these parameter values would dramatically exceed the IceCube observations. This is further illustrated by the fact that the best-fit neutrino spectra for both NGC 4151 and CGCG 420-015 lie well above the gray lines shown in Fig. 11, which indicate the maximum neutrino emission compatible with the observed diffuse flux. Values of the model parameters $P_{\text{CR}}/P_{\text{th}}$ and η for which the single-source neutrino spectra lie fully below the gray line result in diffuse neutrino spectra that exceed the IceCube measurements by less than 1σ . So, if we assumed that all sources were described by the same model parameters as

TABLE IV. Distances and intrinsic x-ray luminosities as well as best-fit model parameters for the nearby Seyfert galaxies NGC 1068, NGC 4151, and CGCG 420-015. The total CR proton luminosity needed to achieve the given pressure ratio is denoted by L_p .

	$d_L [\text{Mpc}]$	$L_X [\text{erg/s}]$	$\hat{n}_s$	$\hat{\eta}$	$P_{\text{CR}}/P_{\text{th}}$	$L_p [\text{erg/s}]$	L_p/L_{Edd}
NGC 1068	11.14	4.2×10^{43}	46.4	56	0.5	6.0×10^{43}	0.009
NGC 4151	15.8 ^a	$2.5 \times 10^{42\text{b}}$	23.5	4	0.5	1.7×10^{43}	0.020
CGCG 420-015	128.8 ^c	$1.0 \times 10^{44\text{c}}$	8.6	1	0.5	4.7×10^{44}	0.038

^aFrom [94].

^bCalculated from the intrinsic 2–10 keV x-ray flux reported in [11].

^cFrom [11].

NGC 4151 or CGCG 420-015, the tension between our diffuse flux predictions and the IceCube observations would become even stronger. Padovani *et al.* [31] come to a similar conclusion based on the fact that NGC 1068 has a much lower neutrino-to-x-ray ratio than both NGC 4151 and CGCG 420-015.

Thus, those Seyfert galaxies that start to emerge as neutrino point sources in the IceCube data appear to be above-average neutrino emitters, and it is very unlikely that all sources in the population share the same neutrino emission properties. To obtain a more realistic estimate of the contribution of Seyfert galaxies to the diffuse neutrino flux, it is necessary to assume a distribution of $P_{\text{CR}}/P_{\text{th}}$ and η among the population. For Seyfert galaxies to be able to account for the observed diffuse flux up to ~ 100 TeV without exceeding the upper limit at ~ 2 TeV, a distribution for which the majority of sources have a low pressure ratio ($P_{\text{CR}}/P_{\text{th}} \lesssim 0.05$) and a high level of turbulence ($\eta \lesssim 10$) is needed. The resulting diffuse neutrino spectra for two distributions that fulfill this requirement, as well as for a uniform distribution of $P_{\text{CR}}/P_{\text{th}}$ and η , are shown in Appendix C. Nevertheless, the exact form of such a distribution remains unknown.

In a stacking analysis, the IceCube Collaboration has derived an upper limit on the diffuse neutrino emission from Seyfert galaxies, which provides slightly stronger constraints than the total observed diffuse flux at energies higher than ~ 20 TeV [5]. However, this limit was calculated under the assumption that the neutrino spectra of all sources follow a single power law with a fixed spectral index and that the neutrino luminosity of a source is proportional to its x-ray luminosity. Therefore, a direct comparison with predictions of other models is possible to only a limited extent.

Murase *et al.* [16], whose description of AGN coronae our neutrino emission model is based on, have also studied the contribution of Seyfert galaxies to the diffuse neutrino flux. Their model neutrino spectrum for NGC 1068 has a lower normalization and a higher cutoff energy than ours, reflecting their focus on explaining the high-energy end of the observed neutrino emission. Consequently, their diffuse neutrino spectrum also has a lower normalization and peaks at higher energies than our result for the best-fit parameters for NGC 1068, shown in Fig. 6. This allows them to reproduce the observed diffuse flux in the ~ 7 – 70 TeV range while remaining consistent with observations of the diffuse MeV γ -ray background. Another difference lies in the treatment of the nonthermal proton spectrum. We approximate it as a power law with an exponential cutoff (see Sec. II B), whereas Murase *et al.* solve the Fokker-Planck equation. For sources with low L_X and high η , for which proton acceleration is limited by diffusive escape, this yields cutoff energies that are ~ 10 – 30 times higher than those that we obtain by equating the acceleration time and the total loss time [93]. Consequently, this leads to a

slightly different mapping between η and the cutoff energy of the proton and neutrino spectra in our framework. Moreover, for sources with high L_X and low η , for which acceleration is stopped by $p\gamma$ and Bethe-Heitler interactions, Murase *et al.* obtain very peaked spectra with a pileup just below the cutoff. Thus, in our model, these sources require somewhat higher CR pressures to achieve the same peak neutrino flux.

Padovani *et al.* [31] estimated the diffuse flux contribution from nonjetted AGNs based on the model neutrino spectrum for NGC 1068 presented in Ref. [32], simply assuming that the neutrino flux of a source scales linearly with its intrinsic x-ray flux. In contrast, under the assumption that the CR-to-thermal pressure ratio is constant among all sources, we find that the neutrino luminosity increases slower than linearly with the x-ray luminosity [see Eq. (20)]. With this assumption, we also ensure that all sources have a pressure ratio lower than $P_{\text{CR}}/P_{\text{th}} = 0.5$, above which the AGN corona becomes unstable (see Sec. II B).

Using a neutrino emission model based on nonresonant stochastic acceleration in strong turbulence with a much higher level of magnetization than in our model, Fiorillo *et al.* [34] can explain the observed diffuse neutrino flux in the ~ 10 – 100 TeV energy range. In addition, they find a scaling of L_ν with L_X that is similar to ours. However, contrary to our work, they do not quantify the consistency of their model with the IceCube observations of NGC 1068. Our approach is also different from that by Ambrosone [33]. Fitting a simple analytical leaky-box model to the neutrino luminosities of four nearby Seyfert galaxies, they obtain a diffuse neutrino flux that exceeds the IceCube observations below ~ 30 TeV by almost one order of magnitude.

VII. CONCLUSIONS

We have investigated the potential contribution of Seyfert galaxies to the astrophysical diffuse neutrino flux, assuming that, inside AGN coronae, protons undergo stochastic acceleration to energies of a few hundred TeV and subsequently interact with the surrounding medium producing both neutrinos and γ rays. Fitting our neutrino emission model to the public IceCube data, we find that explaining the observed neutrino flux from the nearby Seyfert galaxy NGC 1068 requires this source to be an exceptionally efficient hadronic accelerator characterized by the maximum allowed value of the CR-to-thermal pressure ratio of $P_{\text{CR}}/P_{\text{th}} = 0.5$ and an inverse turbulence strength of $\eta = 56$. In addition, we were able to constrain the size of the corona to less than five Schwarzschild radii by comparing the predicted cascaded γ -ray emission to the latest Fermi-LAT observations of NGC 1068.

To estimate the total diffuse neutrino flux from Seyfert galaxies, we have extrapolated our model to the full source population based on the x-ray luminosity function of AGNs. Comparing our results to observations, we have placed the

first detailed constraints on a physically motivated neutrino emission model for Seyfert galaxies informed by neutrino and γ -ray data of NGC 1068 as a nearby benchmark source as well as measurements of the diffuse neutrino flux, accurately incorporating uncertainties.

Our results indicate that Seyfert galaxies can account for a sizable fraction of the observed diffuse neutrino flux below ~ 10 TeV. However, assuming that all sources share the same model parameters as NGC 1068 leads to an overproduction of the observed diffuse flux by 3.8σ . Explaining the neutrino emission of NGC 4151 and CGCG 420-015, two other nearby Seyfert galaxies that are beginning to emerge as neutrino point sources, requires even more extreme parameter values for which the tension between our model predictions and the observed diffuse flux would become even stronger. This suggests that those Seyfert galaxies for which IceCube starts to see evidence for neutrino emission are not representative of the broader population but rather outliers with unusually efficient neutrino production. Under the assumption of the neutrino emission model presented in Sec. II, we can rule out scenarios in which the majority of sources have a high CR pressure ($P_{\text{CR}}/P_{\text{th}} \gtrsim 0.2$) in combination with a high level of turbulence ($\eta \lesssim 40$).

Neutrino observations of other nearby Seyfert galaxies, including those in the southern sky, will be particularly important to gain a better understanding of the uniqueness of NGC 1068 in the context of the entire population. Next-generation neutrino telescopes such as IceCube-Gen2 and KM3NeT will play a key role in this. IceCube-Gen2 is a planned expansion of the instrumented volume of IceCube from 1 to 8 km³ [95]. This will increase the effective area for through-going muon tracks by a factor of 5 [96] and improve the angular resolution, since longer muon tracks can be contained within the detector volume. KM3NeT is a cubic kilometer neutrino telescope currently under construction in the Mediterranean [97]. Using seawater as a detection medium, it will have a better angular resolution than IceCube due to the longer scattering length of Cherenkov photons. In addition, KM3NeT will be more sensitive to neutrino point sources in the southern sky due to its location in the Northern Hemisphere. This is particularly interesting, since several Seyfert galaxies located in the southern sky are intrinsically very bright in x-rays. Assuming that the bulk of the observed x-ray flux is produced in the AGN corona, the most promising targets for KM3NeT include the Circinus galaxy, ESO 138-G001, NGC 7582, and Centaurus A [11,20]. Baikal-GVD, which is currently under construction in Lake Baikal [98], and planned neutrino telescopes such as P-ONE in the Northeast Pacific Ocean [99], TRIDENT in the South China Sea [100], and HUNT [101] will further enhance the chances of identifying neutrino point sources in both the northern and the southern sky.

In addition, γ -ray observations at MeV energies provide a means to test the neutrino emission model underlying our

calculations. If Seyfert galaxies contribute to the diffuse neutrino flux at TeV energies, they may also account for a significant fraction of the diffuse MeV γ -ray background [16,18]. Moreover, observations of nearby individual sources in the MeV band can be used to further constrain possible neutrino production sites. Future MeV γ -ray telescopes such as newASTROGAM [102], AMEGO-X [103], and COSI [104] will, therefore, be crucial for improving our understanding of Seyfert galaxies as hidden neutrino sources.

ACKNOWLEDGMENTS

The authors thank Martina Karl and Chiara Bellenghi for their support with using SkyLLH to fit our model to the public IceCube data. We also thank Xavier Rodrigues for insight into incorporating the diffusive escape mechanism (see Sec. II C) into the AM³ framework. Furthermore, we are grateful to Martin Lemoine for valuable discussions on particle acceleration in magnetized turbulence.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: PROTON TIMESCALES

As explained in Sec. II C, nonthermal protons inside AGN coronae can lose energy via pp interactions, $p\gamma$ interactions, Bethe-Heitler pair production, and synchrotron radiation. In addition, they can escape from the corona via infall onto the SMBH and diffusion.

Figure 12 shows the proton interaction and escape timescales as well as the acceleration timescale for different x-ray luminosities, an inverse turbulence strength of $\eta = 56$, which corresponds to the best-fit value found for NGC 1068 in Sec. III, and a coronal radius of $r_c = 5$. Sources with higher x-ray luminosities have a larger corona and, thus, a weaker magnetic field and lower densities of target protons and photons (see Table I). This leads to overall larger interaction and escape timescales and a slower acceleration process. As mentioned in Sec. II C, the cutoff energy of the proton spectrum is calculated as the energy at which the acceleration timescale is equal to the total loss timescale. This corresponds to the energy at which the green dash-dotted line and the thin black line in Fig. 12 intersect. For intermediate values of the inverse turbulence strength, like $\eta = 56$, the cutoff energy derived in this way does not exhibit a strong dependence on L_X . For sources with $L_X \gtrsim 10^{44}$ erg/s, particle acceleration is mainly stopped by $p\gamma$ and Bethe-Heitler interactions, while for lower luminosities, escape via diffusion becomes equally important.




FIG. 12. Proton timescales for different x-ray luminosities of $L_X = 10^{42}$ erg/s (left panel), 10^{44} erg/s (middle panel), and 10^{46} erg/s (right panel). The inverse turbulence strength is set to $\eta = 56$ and the dimensionless coronal radius to $r_c = 5$. The thin black line corresponds to the total loss timescale defined in Eq. (19).

Because of the relatively small coronal radius of $r_c = 5$, the infall timescale is slightly shorter than the pp interaction timescale at all energies and for all luminosities. Nevertheless, both pp and $p\gamma$ interactions are important for the production of neutrinos, with $p\gamma$ interactions becoming more relevant for higher L_X . Energy losses via synchrotron radiation are negligible for all considered x-ray luminosities.

APPENDIX B: SENSITIVE ENERGY RANGE OF THE MODEL FIT FOR NGC 1068

In order to determine the sensitive energy range of the fit of our neutrino emission model for NGC 1068 to the public IceCube data, we estimate the central 68% interval of neutrino energies that contribute the most to the excess of events over the background [67–69].

To this end, we select all detected neutrino events within a spatial box of size $\pm 10^\circ$ around the position of NGC 1068 and use SkyLLH [63,64] to determine the contribution of each of these events to the test statistic (TS). The latter is defined as

$$\begin{aligned} \text{TS} &\equiv -2 \log \frac{\mathcal{L}(n_s = 0)}{\mathcal{L}(\hat{n}_s, \hat{\eta})} \\ &= 2 \sum_{i=1}^N \log \left[\frac{\hat{n}_s}{N} \left(\frac{S(\hat{E}_i, \hat{\theta}_i | \hat{\eta})}{B(\hat{E}_i, \hat{\theta}_i)} - 1 \right) + 1 \right], \end{aligned} \quad (\text{B1})$$

where $\mathcal{L}$ is the likelihood given by Eq. (21). To map the reconstructed muon energy $\hat{E}$ to the true neutrino energy E_ν , we simulate neutrino events from a point source at the position of NGC 1068 and create a sample of 1.2×10^8 events with a flat distribution in $\log E_\nu$ between 100 GeV and 1 EeV. This ensures sufficient statistics across the entire energy range. For each detected event, we select all simulated events whose reconstructed energies and angular uncertainties differ by less than 10% from those of the detected event and construct the corresponding distribution in true neutrino energy [68]. Then, we compute the weighted sum of these distributions over all detected events, using their individual contributions to the test

statistic as weights, and determine the central 68% energy range from the resulting distribution. For our best-fit model neutrino spectrum of NGC 1068, this yields a sensitive energy range of 350 GeV–190 TeV.

In addition, we use SkyLLH to perform a fit of a simple power-law spectrum for NGC 1068 to the public IceCube data. In agreement with the results presented in [64], we obtain $\hat{n}_s = 56.7$ signal events and a spectral index of $\hat{\gamma} = 3.2$, with the sensitive energy range of the analysis extending from 290 GeV to 110 TeV. The resulting best-fit power-law spectrum is shown in Fig. 13 together with our best-fit model neutrino spectrum for NGC 1068. Overall, the results of the two analyses are compatible with each other. At energies $\gtrsim 500$ GeV, the spectra agree within their 2σ uncertainties and at energies $\gtrsim 1$ TeV within their 1σ




FIG. 13. Best-fit power-law neutrino spectrum of NGC 1068 (blue) together with our best-fit model neutrino spectrum (red). The shaded bands correspond to the 1σ and 2σ uncertainty regions, and their extent along the x axis indicates the central 68% energy range of the respective analysis. The 2σ uncertainty band of the power-law spectrum was calculated with the spectral index restricted to $1.5 \leq \gamma \leq 5.0$. The black solid line shows the best-fit power-law spectrum for NGC 1068 reported in [3], with the dashed lines indicating its 95% confidence region. All fluxes are shown as all-flavor neutrino fluxes.

uncertainties. However, the accuracy of both analyses is limited by the coarse binning of the publicly available instrument response function (IRF) in true neutrino energy and declination [105]. This is the reason why the result of the power-law fit performed with *skyLLH* differs from that reported in [106], although the underlying event sample is the same. The limited energy resolution of the IRF also explains why we obtain relatively large sensitive energy ranges for both analyses.

APPENDIX C: DIFFUSE NEUTRINO FLUX FOR DIFFERENT DISTRIBUTIONS OF $P_{\text{CR}}/P_{\text{th}}$ AND η

In the calculation of the diffuse neutrino flux in Sec. V, we assume that all sources share the same values of the pressure ratio and η . However, under this assumption, the neutrino emission from NGC 1068 cannot be explained without the diffuse neutrino flux produced by the entire population of Seyfert galaxies exceeding the observations by IceCube. Therefore, in a more realistic scenario, both parameters will likely follow some distribution among the population.

Figure 14 shows the diffuse neutrino flux predicted by our model for three possible distributions of $P_{\text{CR}}/P_{\text{th}}$ and η : a uniform distribution (left panel), a log-uniform distribution (middle panel), and a half-normal distribution (right panel). The parameters used for the different distributions are given in Table V. In general, changing from a fixed value of η across the population to a distribution affects both the cutoff energy and the normalization of the diffuse neutrino flux (see blue dashed lines). The higher the abundance of sources with small values of η , the higher the cutoff energy and the normalization of the resulting

spectrum. This is due to the fact that the neutrino spectra of sources with a higher level of turbulence not only extend to higher energies, but also have a higher overall normalization than those of sources with a less turbulent magnetic field (see Fig. 1). In contrast, assuming a distribution of the pressure ratio instead of a fixed value for all sources in the population affects only the normalization, while the overall shape of the spectrum stays the same (see blue dotted lines).

The diffuse flux spectrum obtained for a uniform distribution of both $P_{\text{CR}}/P_{\text{th}}$ and η has a slightly lower normalization and a higher cutoff energy compared to that for fixed parameter values. However, it still exceeds the IceCube observations at energies $\lesssim 10$ TeV, suggesting that this distribution is unlikely to represent the actual underlying one. This is different for the other two distributions, both of which peak at low values of $P_{\text{CR}}/P_{\text{th}}$ and η and lead to diffuse neutrino spectra that are broadly consistent with the IceCube flux measurements up to ~ 100 TeV. Here, we have chosen the minimum value of $P_{\text{CR}}/P_{\text{th}}$ for the log-uniform distribution as well as the widths of the half-normal distributions of $P_{\text{CR}}/P_{\text{th}}$ and η such that the resulting contribution of the Seyfert population to the observed diffuse flux is maximized without exceeding the upper limit at ~ 2 TeV.

Overall, a distribution for which most sources have a low pressure ratio ($P_{\text{CR}}/P_{\text{th}} \lesssim 0.05$) and a high level of turbulence ($\eta \lesssim 10$) is necessary for the population of Seyfert galaxies to account for the observed diffuse neutrino flux up to energies of ~ 100 TeV. In such a scenario, NGC 1068 with $P_{\text{CR}}/P_{\text{th}} = 0.5$ and $\eta = 56$ would stand out as an exceptional source, regardless of the exact form of the underlying distribution.




FIG. 14. Diffuse all-flavor neutrino flux for a uniform (left panel), a log-uniform (middle panel), and a half-normal distribution (right panel) of $P_{\text{CR}}/P_{\text{th}}$ and η among the population, calculated for the U14 XLF. Red lines show the flux for the case in which both parameters follow the given distribution. Dashed (dotted) blue lines correspond to the case where only η ($P_{\text{CR}}/P_{\text{th}}$) varies, while $P_{\text{CR}}/P_{\text{th}} = 0.5$ ($\eta = 56$) is fixed for all sources. For comparison, the flux obtained when fixing both parameters to the best-fit values for NGC 1068 is shown in gray. The black, dark gray, and light gray data points represent the diffuse neutrino flux measured by IceCube using starting tracks [81], cascade events [82], and northern sky tracks [83], respectively. The parameters of the different distributions are given in Table V.

TABLE V. Parameters of the different distributions of $P_{\text{CR}}/P_{\text{th}}$ and η underlying the diffuse flux spectra shown in Fig. 14. Listed are the lower and upper bounds of the considered parameter ranges as well as the mean and the standard deviation for the normal distribution.

Distribution	$(P_{\text{CR}}/P_{\text{th}})_{\text{min}}$	$(P_{\text{CR}}/P_{\text{th}})_{\text{max}}$	η_{min}	η_{max}	$\mu_{P_{\text{CR}}/P_{\text{th}}}$	$\sigma_{P_{\text{CR}}/P_{\text{th}}}$	μ_{η}	σ_{η}
Uniform	0.0	0.5	1	100	...	...	...	...
Log-uniform	10^{-7}	0.5	1	100	...	...	...	...
Half-normal	0.0	0.5	1	100	0.0	0.05	1	10

- [1] M. G. Aartsen *et al.*, *Phys. Rev. Lett.* **111**, 021103 (2013).
[2] M. G. Aartsen *et al.*, *Science* **342**, 1242856 (2013).
[3] R. Abbasi *et al.*, *Science* **378**, 538 (2022).
[4] R. Abbasi *et al.*, *Astrophys. J.* **988**, 141 (2025).
[5] R. Abbasi *et al.*, *Astrophys. J.* **981**, 131 (2025).
[6] A. Neronov, D. Savchenko, and D. V. Semikoz, *Phys. Rev. Lett.* **132**, 101002 (2024).
[7] G. Sommani, A. Franckowiak, M. Lincetto, and R.-J. Dettmar, *Astrophys. J.* **981**, 103 (2025).
[8] P. Padovani *et al.*, *Nat. Astron.* **8**, 1077 (2024).
[9] N. A. Tikhonov and O. A. Galazutdinova, *Astrophys. Bull.* **76**, 255 (2021).
[10] A. Marinucci *et al.*, *Mon. Not. R. Astron. Soc.* **456**, L94 (2016).
[11] C. Ricci *et al.*, *Astrophys. J. Suppl. Ser.* **233**, 17 (2017).
[12] F. Aharonian *et al.*, *Astron. Astrophys.* **441**, 465 (2005).
[13] V. A. Acciari *et al.*, *Astrophys. J.* **883**, 135 (2019).
[14] A. Albert *et al.*, *Astrophys. J.* **907**, 67 (2021).
[15] K. Murase, D. Guetta, and M. Ahlers, *Phys. Rev. Lett.* **116**, 071101 (2016).
[16] K. Murase, S. S. Kimura, and P. Mészáros, *Phys. Rev. Lett.* **125**, 011101 (2020).
[17] B. Eichmann, F. Oikonomou, S. Salvatore, R.-J. Dettmar, and J. B. Tjus, *Astrophys. J.* **939**, 43 (2022).
[18] Y. Inoue, D. Khangulyan, S. Inoue, and A. Doi, *Astrophys. J.* **880**, 40 (2019).
[19] Y. Inoue, D. Khangulyan, and A. Doi, *Astrophys. J. Lett.* **891**, L33 (2020).
[20] A. Kheirandish, K. Murase, and S. S. Kimura, *Astrophys. J.* **922**, 45 (2021).
[21] D. F. G. Fiorillo, M. Petropoulou, L. Comisso, E. Peretti, and L. Sironi, *Astrophys. J. Lett.* **961**, L14 (2024).
[22] D. Karavola, M. Petropoulou, D. F. G. Fiorillo, L. Comisso, and L. Sironi, *J. Cosmol. Astropart. Phys.* **04** (2025) 075.
[23] A. Lamastra, F. Fiore, D. Guetta, L. A. Antonelli, S. Colafrancesco, N. Menci, S. Puccetti, A. Stamerra, and L. Zappacosta, *Astron. Astrophys.* **596**, A68 (2016).
[24] E. Peretti, A. Lamastra, F. G. Saturni, M. Ahlers, P. Blasi, G. Morlino, and P. Cristofari, *Mon. Not. R. Astron. Soc.* **526**, 181 (2023).
[25] S. Inoue, M. Cerruti, K. Murase, and R.-Y. Liu, *arXiv:2207.02097*.
[26] K. Fang, E. Lopez Rodriguez, F. Halzen, and J. S. Gallagher, *Astrophys. J.* **956**, 8 (2023).
[27] H. Tagawa, S. S. Kimura, and Z. Haiman, *Astrophys. J.* **955**, 23 (2023).
[28] D. Hooper and K. Plant, *Phys. Rev. Lett.* **131**, 231001 (2023).
[29] A. Das, B. T. Zhang, and K. Murase, *Astrophys. J.* **972**, 44 (2024).
[30] K. Yasuda, N. Sakai, Y. Inoue, and A. Kusenko, *Phys. Rev. Lett.* **134**, 151005 (2025).
[31] P. Padovani, R. Gilli, E. Resconi, C. Bellenghi, and F. Henningsen, *Astron. Astrophys.* **684**, L21 (2024).
[32] K. Murase, *Astrophys. J. Lett.* **941**, L17 (2022).
[33] A. Ambrosone, *J. Cosmol. Astropart. Phys.* **09** (2024) 075.
[34] D. F. G. Fiorillo, L. Comisso, E. Peretti, M. Petropoulou, and L. Sironi, *Astrophys. J.* **989**, 215 (2025).
[35] F. Haardt and L. Maraschi, *Astrophys. J. Lett.* **380**, L51 (1991).
[36] K. A. Miller and J. M. Stone, *Astrophys. J.* **534**, 398 (2000).
[37] A. Merloni and A. C. Fabian, *Mon. Not. R. Astron. Soc.* **321**, 549 (2001).
[38] B. F. Liu, S. Mineshige, and K. Shibata, *Astrophys. J. Lett.* **572**, L173 (2002).
[39] Y. Io and T. K. Suzuki, *Astrophys. J.* **780**, 46 (2014).
[40] Y.-F. Jiang, J. M. Stone, and S. W. Davis, *Astrophys. J.* **784**, 169 (2014).
[41] Y.-F. Jiang, O. Blaes, J. M. Stone, and S. W. Davis, *Astrophys. J.* **885**, 144 (2019).
[42] A. A. Zdziarski, W. N. Johnson, and P. Magdziarz, *Mon. Not. R. Astron. Soc.* **283**, 193 (1996).
[43] F. Haardt, L. Maraschi, and G. Ghisellini, *Astrophys. J.* **476**, 620 (1997).
[44] M. Lemoine and M. A. Malkov, *Mon. Not. R. Astron. Soc.* **499**, 4972 (2020).
[45] M. Lemoine, K. Murase, and F. Rieger, *Phys. Rev. D* **109**, 063006 (2024).
[46] M. Lemoine and F. Rieger, *Astron. Astrophys.* **697**, A124 (2025).
[47] M. Klinger, A. Rudolph, X. Rodrigues, C. Yuan, G. Fichtel de Clairfontaine, A. Fedynitch, W. Winter, M. Pohl, and S. Gao, *Astrophys. J. Suppl. Ser.* **275**, 4 (2024).
[48] J. A. Mayers *et al.*, *arXiv:1803.06891*.
[49] B. Trakhtenbrot *et al.*, *Mon. Not. R. Astron. Soc.* **470**, 800 (2017).
[50] C. Ricci *et al.*, *Mon. Not. R. Astron. Soc.* **480**, 1819 (2018).

- [51] P. F. Hopkins, G. T. Richards, and L. Hernquist, *Astrophys. J.* **654**, 731 (2007).
- [52] E. Kara, J. A. García, A. Lohfink, A. C. Fabian, C. S. Reynolds, F. Tombesi, and D. R. Wilkins, *Mon. Not. R. Astron. Soc.* **468**, 3489 (2017).
- [53] N. Kamraj *et al.*, *Astrophys. J.* **927**, 42 (2022).
- [54] Y. Inoue, D. Khangulyan, and A. Doi, *Galaxies* **9**, 36 (2021).
- [55] Y. Inoue, S. Takasao, and D. Khangulyan, *Publ. Astron. Soc. Jpn.* **76**, 996 (2024).
- [56] S. S. Kimura, K. Murase, and P. Mészáros, *Phys. Rev. D* **100**, 083014 (2019).
- [57] P. A. Becker, T. Le, and C. D. Dermer, *Astrophys. J.* **647**, 539 (2006).
- [58] Ł. Stawarz and V. Petrosian, *Astrophys. J.* **681**, 1725 (2008).
- [59] G. Ghisellini, *Radiative Processes in High Energy Astrophysics* (Springer International Publishing, New York, 2013).
- [60] G. Sigl, *Astroparticle Physics: Theory and Phenomenology* (Atlantis Press, Paris, 2017), Vol. 1.
- [61] IceCube Collaboration, All-sky point-source IceCube data: Years 2008–2018, [10.7910/DVN/VKL316](https://arxiv.org/abs/10.7910/DVN/VKL316) (2021).
- [62] J. Braun, J. Dumm, F. De Palma, C. Finley, A. Karle, and T. Montaruli, *Astropart. Phys.* **29**, 299 (2008).
- [63] M. Wolf, *Proc. Sci., ICRC2019* (**2019**) 1035 [[arXiv:1908.05181](https://arxiv.org/abs/1908.05181)].
- [64] C. Bellenghi *et al.*, *Proc. Sci., ICRC2023* (**2023**) 1061 [[arXiv:2308.12733](https://arxiv.org/abs/2308.12733)].
- [65] C. Blanco, D. Hooper, T. Linden, and E. Pinetti, *Phys. Rev. D* **112**, 123016 (2025).
- [66] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [67] S. Schoenen, Discovery and characterization of a diffuse astrophysical muon neutrino flux with the IceCube neutrino observatory, Ph.D. thesis, RWTH Aachen University, 2017.
- [68] C. Bellenghi, The emergence of a new sky: First associations of IceCube high-energy neutrinos with Active Galactic Nuclei, Ph.D. thesis, Technical University of Munich, 2024.
- [69] T. Kontrimas, Methods for the search of high-energy neutrino sources: From IceCube data to source modeling, Ph.D. thesis, Technical University of Munich, 2025.
- [70] M. Ajello, K. Murase, and A. McDaniel, *Astrophys. J. Lett.* **954**, L49 (2023).
- [71] F. Capel, J. Kuhlmann, C. Haack, M. Ha Minh, H. Niederhausen, and L. Schumacher, *Astrophys. J.* **976**, 127 (2024).
- [72] S. Abdollahi *et al.*, *Astrophys. J. Suppl. Ser.* **247**, 33 (2020).
- [73] S. Laha, C. Ricci, J. C. Mather, E. Behar, L. Gallo, F. Marin, R. Mbarek, and A. Hankla, *Front. Astron. Space Sci.* **11**, 1530392 (2025).
- [74] S. Salvatore, B. Eichmann, X. Rodrigues, R. J. Dettmar, and J. Becker Tjus, *Astron. Astrophys.* **687**, A139 (2024).
- [75] Y. Ueda, M. Akiyama, G. Hasinger, T. Miyaji, and M. G. Watson, *Astrophys. J.* **786**, 104 (2014).
- [76] J. Buchner *et al.*, *Astrophys. J.* **802**, 89 (2015).
- [77] A. Akylas, A. Georgakakis, I. Georgantopoulos, M. Brightman, and K. Nandra, *Astron. Astrophys.* **546**, A98 (2012).
- [78] J. Burgess and F. Capel, *J. Open Source Software* **6**, 3257 (2021).
- [79] T. T. Ananna, E. Treister, C. M. Urry, C. Ricci, A. Kirkpatrick, S. LaMassa, J. Buchner, F. Civano, M. Tremmel, and S. Marchesi, *Astrophys. J.* **871**, 240 (2019).
- [80] M. Adachi and M. Kasai, *Prog. Theor. Phys.* **127**, 145 (2012).
- [81] R. Abbasi *et al.*, *Phys. Rev. D* **110**, 022001 (2024).
- [82] M. G. Aartsen *et al.*, *Phys. Rev. Lett.* **125**, 121104 (2020).
- [83] R. Abbasi *et al.*, *Astrophys. J.* **928**, 50 (2022).
- [84] R. Naab *et al.*, *Proc. Sci., ICRC2023* (**2023**) 1064 [[arXiv:2308.00191](https://arxiv.org/abs/2308.00191)].
- [85] E. Peretti, P. Blasi, F. Aharonian, G. Morlino, and P. Cristofari, *Mon. Not. R. Astron. Soc.* **493**, 5880 (2020).
- [86] A. Ambrosone, M. Chianese, D. F. G. Fiorillo, A. Marinelli, G. Miele, and O. Pisanti, *Mon. Not. R. Astron. Soc.* **503**, 4032 (2021).
- [87] D. Ehlert, F. Oikonomou, and E. Peretti, *Mon. Not. R. Astron. Soc.* **539**, 2435 (2025).
- [88] R. Abbasi *et al.*, *Phys. Rev. D* **106**, 022005 (2022).
- [89] K. M. Groth and M. Ahlers, *Phys. Rev. D* **111**, 103052 (2025).
- [90] M. G. Aartsen *et al.*, *Astrophys. J.* **835**, 45 (2017).
- [91] M. Huber, *Proc. Sci., ICRC2019* (**2019**) 916 [[arXiv:1908.08458](https://arxiv.org/abs/1908.08458)].
- [92] P. Giommi and P. Padovani, *Universe* **7**, 492 (2021).
- [93] K. Murase, C. M. Karwin, S. S. Kimura, M. Ajello, and S. Buson, *Astrophys. J. Lett.* **961**, L34 (2024).
- [94] W. Yuan *et al.*, *Astrophys. J.* **902**, 26 (2020).
- [95] M. G. Aartsen *et al.*, *J. Phys. G* **48**, 060501 (2021).
- [96] B. A. Clark, *J. Instrum.* **16**, C10007 (2021).
- [97] S. Adrián-Martínez *et al.*, *J. Phys. G* **43**, 084001 (2016).
- [98] V. M. Aynutdinov *et al.*, *Proc. Sci., ICRC2023* (**2023**) 976 [[arXiv:2309.16310](https://arxiv.org/abs/2309.16310)].
- [99] M. Agostini *et al.*, *Nat. Astron.* **4**, 913 (2020).
- [100] Z. P. Ye *et al.*, *Nat. Astron.* **7**, 1497 (2023).
- [101] T.-Q. Huang, Z. Cao, M. Chen, J. Liu, Z. Wang, X. You, and Y. Qi, *Proc. Sci., ICRC2023* (**2023**) 1080.
- [102] D. Berge, M. N. Mazziotta, M. Tavani, V. Tatischeff, and U. Oberlack, *Proc. Sci., ICRC2025* (**2025**) 572 [[arXiv:2507.08133](https://arxiv.org/abs/2507.08133)].
- [103] R. Caputo *et al.*, *J. Astron. Telesc. Instrum. Syst.* **8**, 044003 (2022).
- [104] J. Tomsick *et al.*, *Proc. Sci., ICRC2023* (**2023**) 745 [[arXiv:2308.12362](https://arxiv.org/abs/2308.12362)].
- [105] R. Abbasi *et al.*, [arXiv:2101.09836](https://arxiv.org/abs/2101.09836).
- [106] M. G. Aartsen *et al.*, *Phys. Rev. Lett.* **124**, 051103 (2020).