



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/165064/>

Version: Submitted Version

Article:

Vahsen, S.E., O'Hare, C.A.J., Lynch, W.A. et al. (Submitted: 2020) CYGNUS : Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos. arXiv. (Submitted)

© 2020 The Author(s). For reuse permissions, please contact the Author(s).

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

CYGNUS: Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos

S. E. Vahsen,¹ C. A. J. O’Hare,² W. A. Lynch,³ N. J. C. Spooner,³ E. Baracchini,^{4,5,6} P. Barbeau,⁷ J. B. R. Battat,⁸ B. Crow,¹ C. Deaconu,⁹ C. Eldridge,³ A. C. Ezeribe,³ M. Ghrear,¹ D. Loomba,¹⁰ K. J. Mack,¹¹ K. Miuchi,¹² F. M. Mouton,³ N. S. Phan,¹³ K. Scholberg,⁷ and T. N. Thorpe^{1,6}

¹*Department of Physics and Astronomy, University of Hawaii, Honolulu, Hawaii 96822, USA*

²*The University of Sydney, School of Physics, NSW 2006, Australia*

³*Department of Physics and Astronomy, University of Sheffield, S3 7RH, Sheffield, United Kingdom*

⁴*Istituto Nazionale di Fisica Nucleare, Laboratori Nazionali di Frascati, I-00040, Italy*

⁵*Istituto Nazionale di Fisica Nucleare, Sezione di Roma, I-00185, Italy*

⁶*Department of Astroparticle Physics, Gran Sasso Science Institute, L’Aquila, I-67100, Italy*

⁷*Department of Physics, Duke University, Durham, NC 27708 USA*

⁸*Department of Physics, Wellesley College, Wellesley, Massachusetts 02481, USA*

⁹*Department of Physics, Enrico Fermi Inst., Kavli Inst. for Cosmological Physics, Univ. of Chicago, Chicago, IL 60637, USA*

¹⁰*Department of Physics and Astronomy, University of New Mexico, NM 87131, USA*

¹¹*Department of Physics, North Carolina State University, Raleigh, NC 27695, USA*

¹²*Department of Physics, Kobe University, Rokkodaicho, Nada-ku, Hyogo 657-8501, Japan*

¹³*Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM 87545, USA*

(Dated: August 31, 2020)

Now that conventional weakly interacting massive particle (WIMP) dark matter searches are approaching the neutrino floor, there has been a resurgence of interest in detectors with sensitivity to nuclear recoil directions. A large-scale directional detector is attractive in that it would have sensitivity below the neutrino floor, be capable of unambiguously establishing the galactic origin of a purported dark matter signal, and could serve a dual purpose as a neutrino observatory. We present the first detailed analysis of a 1000 m³-scale detector capable of measuring a directional nuclear recoil signal at low energies. We propose a modular and multi-site observatory consisting of time projection chambers (TPCs) filled with helium and SF₆ at atmospheric pressure. By comparing several available readout technologies, we identify high-resolution strip readout TPCs as the optimal tradeoff between performance and cost. We estimate that suitable angular resolution and head-tail recognition is achievable down to helium recoil energies of ~ 6 keV_r. Depending on the readout technology, an average of only 4-5 detected 100-GeV c^{-2} WIMP-fluorine recoils above 50 keV_r are sufficient to rule out an isotropic recoil distribution at 90% CL. An average of 10-20 helium recoils above 6 keV_r or only 3-4 helium recoils above 20 keV_r would suffice to distinguish a 10 GeV c^{-2} WIMP signal from the solar neutrino background. High-resolution TPC charge readout also enables powerful electron background rejection capabilities well below 10 keV. We detail background and site requirements at the 1000 m³-scale, and identify materials that require improved radiopurity. The final experiment, which we name CYGNUS-1000, will be able to observe ~ 10 –40 neutrinos from the Sun, depending on the final energy threshold. With the same exposure, the sensitivity to spin independent cross sections will extend into presently unexplored sub-10 GeV c^{-2} parameter space. For spin dependent interactions, already a 10 m³-scale experiment could compete with upcoming generation-two detectors, but CYGNUS-1000 would improve upon this considerably. Larger volumes would bring sensitivity to neutrinos from an even wider range of sources, including galactic supernovae, nuclear reactors, and geological processes.

I. INTRODUCTION

A wide range of astrophysical observations across galactic and cosmological scales indicate that dark matter (DM) dominates the mass budget of the Universe. While the gravitational evidence for DM is now overwhelmingly strong (for a review, see Ref. [1]), its particle identity remains unknown. A definitive detection of DM is expected to be a gateway to physics beyond the Standard Model. Being of such fundamental importance, the quest to uncover the nature of DM remains one of the most important experimental challenges in contemporary physics [2].

Experimental efforts in direct DM detection have

largely focused on the possibility that it consists of weakly interacting massive particles (WIMPs) whose scattering cross sections with ordinary matter are small but nonzero, detectable in low-background nuclear recoil experiments. After years of incremental improvements, large swaths of WIMP parameter space are now ruled out, including those in which detections were previously claimed [3–12]. As detectors become sensitive to lower masses and weaker cross sections, the previously negligible neutrino background will become important, potentially swamping any WIMP signal. Models with weak signals that are hidden beneath the neutrino background are said to live below the “neutrino floor”, a theoretical boundary beyond which conventional DM detectors can-

not efficiently probe.

DM detectors with *directional sensitivity*—the ability to reconstruct the directions of WIMP-induced nuclear recoil events—have the potential to address two of the main challenges to current detection techniques: (1) the difficulty of positively identifying detected interactions as being caused by the DM making up the galactic halo, and (2) the potential for the saturation of the signal by the neutrino background. Directionality addresses both challenges by identifying the direction from which interacting particles originate, thus distinguishing between events coming from the Milky Way DM halo and those originating elsewhere. While this is a powerful background discrimination technique in principle, in practice it remains challenging to build a directional detector that is large enough to compete with non-directional experiments.

This article lays out the science case and goals for a large directional nuclear recoil observatory, and details the challenges that need to be faced in order to achieve those goals with existing technology. The motivation for this work draws on results summarized in three extensive reviews: the first summarized the status of existing directional detection projects [13], the second detailed the full discovery potential for directional detectors [14]; and the third presented a broad summary of available direction-sensitive readout technologies [15]. Here, we combine all present knowledge in the field and develop a single strategy for a directional detector named CYGNUS. This work bridges the gap between, on the one hand, theoretical literature on directional detection which often involves many idealized assumptions about detector capabilities; and on the other, purely experimental studies which are relatively far removed from a potential future utility in a competitive DM search.

The primary goal of CYGNUS will be to discover DM even if its cross section lies beyond the neutrino floor. WIMP sensitivity at this level also brings a secondary motivation: the study of the neutrino background itself. There is substantial physics motivation for the detection of natural or human-made neutrinos, via either coherent neutrino-nucleus scattering ($\text{CE}\nu\text{NS}$) or neutrino-electron scattering, both of which will be measurable in CYGNUS. In particular, directional measurements of these processes would be a highly-novel additional signal for many terrestrial and astrophysical sources of neutrino. CYGNUS, therefore, retains a strong science case, even in the absence of a positive DM signal.

Observing the neutrino background and WIMP cross sections lying below the neutrino floor requires experiments with very large total masses. Hence the final goal for CYGNUS will be a network of modular and multi-site gas experiments each individually based on the strategies set out here. Organizing the experiment in a modular and multi-site manner alleviates some of the issues regarding the very large volumes needed for a gas target to reach the required mass. A schematic for a style of modularity is shown in Fig. 1, which shows how several smaller “back-to-back” gas time projection chambers (TPCs) can

be placed within the same background shielding. Our study focuses on the design of a single module that can be scaled up in this fashion. In Fig. 2 we show the locations of some of the candidate sites under consideration. A globally distributed experiment is also advantageous as it provides an additional method to deal with seasonal or location-dependent backgrounds as well as help to overcome possible site-restrictions on the overall size of a detector.

The organization of this article is as follows: in Sec. II we describe the advantages of directionality for the robust discovery and characterization of a DM signal, and to contribute to studies of solar, supernova, atmospheric, and geological neutrinos. Section III reviews the existing and proposed detector technologies that may enable directional detection, including the gas TPC, which forms the basis of CYGNUS. In Sec. IV, we describe how different TPC gases and readout technologies compare in their detection capabilities while keeping in mind practical considerations for construction. Section V discusses the feasibility of achieving zero-background in large-scale TPCs, while Sec. VI describes the requirements for detector location and installation. Finally in Sec. VII, we conclude the technological considerations described in the above sections and outline a specific direction-sensitive detector that is both feasible and cost-effective.

II. SCIENCE CASE FOR A LARGE NUCLEAR RECOIL OBSERVATORY

Past and ongoing directional detection experiments such as DRIFT [16], MIMAC [17], D^3 [18], DMTPC [19] and NEWAGE [20] have demonstrated the low-pressure gas TPC concept, and have made many impressive technological advances over the years. However, current stringent WIMP limits from non-directional experiments dictate a significant scaling-up of these technologies, which will pose new experimental challenges. This section provides the motivation for facing those challenges by laying out the strong science case for a direction-sensitive “recoil observatory.”

A. Motivation

The most popular technique to directly detect WIMP DM is to search for nuclear recoils with energies $\mathcal{O}(1 - 100)$ keV [21, 22]. Experimental strategies using a wide range of targets have been developed, exploiting charge, light, or heat signals (for a review see *e.g.* Ref. [21]). After decades of progress, the tightest limits now challenge favored models from supersymmetry [23]. Several experiments have in the past reported detections that are consistent with a DM interpretation, but were in tension with limits from other experiments. The only outstanding WIMP detection claim remaining is from the DAMA/LIBRA collaboration who report a




FIG. 1. Schematic of a modular $N \times 10 \text{ m}^3$ directional detector made up of N back-to-back TPC modules. Each module would have two readout planes and a central cathode to ensure a short maximum drift distance of 50 cm.

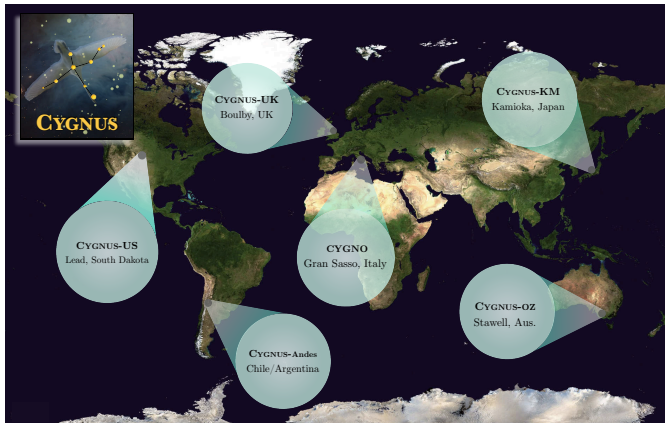


FIG. 2. Proposed sites for a network of CYGNUS detectors. Site considerations for a large-scale TPC are discussed in Sec. VI.

9.3σ annual modulation of the event rate in their NaI(Tl) crystal scintillators [24, 25]. While the annual modulation they observe is ostensibly DM-like, such an interpretation is without support from any other experiment. The annual modulation of galactic DM was often touted as a clear signature of DM, however; the DAMA/LIBRA result has demonstrated this approach can be entirely frustrated by systematics. So persistent is this disagreement that there are now several experiments around the world designed to test the DAMA result in the most direct way possible: by replicating the experiment (DM-Ice [26], KIMS [27], SABRE [28], COSINE-100 [12, 29],

ANAIS [30, 31]) and COSINUS [32, 33]. Of the currently released results from ANAIS and COSINE-100, both are still consistent with a modulation amplitude of zero. However they will need several more years of exposure to begin to definitively rule out the presence of a modulation at the level of DAMA.

Though experience has uncovered the limitations of the annual modulation signature – the distribution of recoil directions promises more robust and convincing evidence for DM. The directional DM signature arises from the motion of the Earth (and solar system) through the galactic halo. The flux of DM particles incident on Earth is strongly anisotropic. This so-called DM “wind” peaks close to the star Deneb in the constellation Cygnus [34]. For DM-induced recoil directions, this leads to two prominent phenomena. Firstly, in galactic coordinates the angular distribution of recoils should be a dipole pointing back towards this well-understood direction. Secondly, in lab-centric coordinates, the anisotropy will cycle across the sky due to the Earth’s rotation, with a period of one *sidereal* day. By definition, these effects are observable in *direction-sensitive* experiments, but not otherwise. The measurement of recoil directions is therefore critical to establishing unequivocally that the source of some signal excess in events is due to the same particle that makes up the DM in the Milky Way. Such a detector would also have a strongly-enhanced ability to remove backgrounds. In an idealized situation, directional capability would make a WIMP search maximally reliable and robust by eliminating all possible backgrounds [35–37] and demonstrating the galactic origin of the detected

particle [38].

As anticipated long before the advent of ton-scale detectors, natural sources of neutrinos are the ultimate irreducible background to WIMP searches. Indeed, the upcoming generation of direct detection experiments are anticipated to be so large and so sensitive that the coherent scattering between neutrinos and the target will already be one of the most important backgrounds. Neutrinos are unshieldable so without a strategy to discriminate against them, the discovery of some well-motivated low-mass WIMP models [39] would require prohibitively large exposures. However, even in the eventuality that no WIMP signal is detected, the neutrino background itself constitutes an interesting signal. A measurement of the angular distribution of neutrinos is of great interest in the context of supernovae [40], geoneutrinos [41], and atmospheric neutrinos [42]. Finally, we note that in the event of a positive WIMP detection, a directional experiment would be able to probe some more exotic dark particle physics as well as to map the currently unknown three-dimensional velocity structure of the DM distribution around the solar system [14, 43–47]. This latter achievement would greatly improve our understanding of the cosmological accretion and merger history of our galaxy, making such an experiment an instrument of astrophysics as well as particle physics.

B. Dark matter

1. WIMP scattering review

The event rate for WIMP-induced nuclear recoils is derived by integrating the product of the incoming flux of DM and the cross section σ for the relevant WIMP-nucleus interaction. This is usually written in terms of the differential event rate R per unit detector mass, as a function of recoil energy E_r and time t ,

$$\frac{dR}{dE_r}(E_r, t) = \frac{\rho_0}{m_\chi m_A} \int_{v > v_{\min}} v f(\mathbf{v}, t) \frac{d\sigma}{dE_r}(E_r, v) d^3v, \quad (1)$$

where ρ_0 is the local DM mass density, m_χ is the WIMP mass, m_A is the nucleus mass and \mathbf{v} is the DM velocity in the detector rest frame. The integral is performed for WIMP speeds v larger than $v_{\min}(E_r) = \sqrt{m_A E_r / 2\mu_{\chi A}^2}$ which is the minimum speed kinematically permitted to produce a recoil with energy E_r . The factor $\mu_{\chi A}$ is the WIMP-nucleus reduced mass. The integral over velocities is weighted by the flux of DM particles $v f(\mathbf{v}, t)$ where f is the galactic distribution of DM velocities. For DM in the form of WIMPs, over the timescales of observation, the velocity distribution is constant in time. In this formula however, it picks up a time dependence after a boost into the laboratory rest frame by \mathbf{v}_{lab} , the velocity with which we are moving through the galactic halo. The differential cross section is proportional to the squared matrix element for a particular WIMP-nucleus

interaction, so is therefore specifically model-dependent. However one can work with general formulae. The most common approach is to divide the interaction into two possible channels, both of which may contribute to the total rate,

$$\frac{d\sigma}{dE_r} = \frac{m_A}{2\mu_{\chi A}^2 v^2} (\sigma_0^{\text{SI}} F_{\text{SI}}^2(E_r) + \sigma_0^{\text{SD}} F_{\text{SD}}^2(E_r)). \quad (2)$$

Here, the first term describes spin-independent (SI) interactions such as those arising from scalar or vector WIMP-quark couplings, whereas the second includes the spin-dependent (SD) contributions from, for example, axial-vector couplings. The cross sections $\sigma_0^{\text{SI,SD}}$ are defined at zero momentum transfer, and nuclear form factors $F_{\text{SI,SD}}^2$ (where $F(0) \equiv 1$) are invoked to describe the nuclear structure and its response to WIMP scattering events at different energies: in particular the suppression of the rate due to the loss in coherence over the nucleus at high momentum transfers. Note that the form factors depend entirely on nuclear physics, and all WIMP model dependence is contained in $\sigma_0^{\text{SI,SD}}$. Keep in mind that this is only the baseline set of WIMP interactions which are most commonly used to compare experimental results. The list of all the physically-allowed non-relativistic WIMP-nucleus interaction operators is much longer than the two mentioned so far, and accounting for them reveals a much richer array of possible signals (as will be discussed in Sec. II B 4). But for simply setting and comparing null results, the convention has been to focus on the SI and SD cross sections. Simplifying things even further, it is also commonplace to write all nuclear cross sections as proportional to those on the proton $\sigma_p^{\text{SI,SD}}$,

$$\sigma_0^{\text{SI}} = |Z + (f_n/f_p)(A - Z)|^2 \sigma_p^{\text{SI}}, \quad (3)$$

$$\sigma_0^{\text{SD}} = \frac{4}{3} \frac{J + 1}{J} |\langle S_p \rangle + (a_n/a_p)\langle S_n \rangle|^2 \sigma_p^{\text{SD}}, \quad (4)$$

where A is the nucleus mass number, Z is the nucleus atomic number, $\langle S_{p,n} \rangle$ are the average of the proton and neutron spins in the nucleus, and J is the total nuclear spin. The ratios of the couplings to the proton and neutron are written as f_p/f_n and a_p/a_n for SI and SD interactions respectively. In the SI case, the total nuclear cross section is enhanced by the number of nucleons squared (assuming, as we do here, isospin conserving DM where $f_n/f_p = 1$), meaning that large target nuclei tend to have the potential to set the most stringent limits. In the spin-dependent case, the interaction probability is not amplified by mass number, but depends instead on the spin content of the target nuclei, hence; constraints tend to be weaker.

In the case of directional detection we want to know how the differential event rate depends on a recoil's direction as well as its energy. We can derive this by considering the kinematic relationship between the incoming

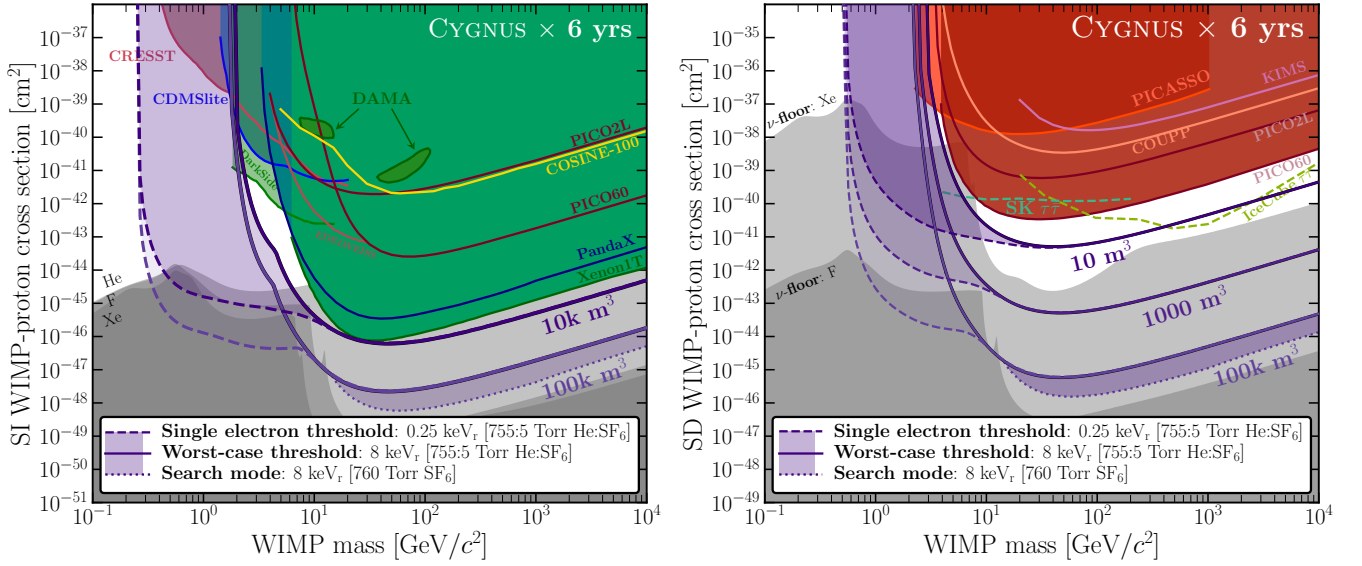


FIG. 3. Constraints on the spin-independent WIMP-nucleon (left) and spin-dependent WIMP-proton (right) cross sections. We show the existing constraints and detections from various experiments as labeled (see text for the associated references). In purple solid and dashed lines we show our projected 90% CL exclusion limits for the CYGNUS experiment operating for six years with 10 m^3 up to $100,000 \text{ m}^3$ of He:SF₆ gas at 755:5 Torr (where 6 years \times 1000 m^3 corresponds to a ~ 1 ton-year exposure). For each volume we show limits for two possible thresholds, ranging from the worst-case electron discrimination threshold of 8 keV_r to a very best-case minimum threshold corresponding to a single electron, 0.25 keV_r . We emphasize however that we anticipate electron discrimination well below 8 keV_r . For the 100k m^3 limits we add a third dotted line which corresponds to a mode with purely SF₆ gas at 760 Torr. This experiment would have a significantly higher total mass but would come at the cost of any directional sensitivity. This ‘search mode’ could be used to extend the high mass sensitivity to just within reach of the neutrino floor. For the SI panel, we shade in gray the neutrino floor for helium, fluorine, and xenon targets (top to bottom), and for SD we show only fluorine and xenon. We define the neutrino floor as the cross section limit at which the rate of improvement with increasing exposure is the slowest in standard direct detection—the effect that CYGNUS aims to circumvent. This definition corresponds to $\mathcal{O}(100)$ neutrino events.

WIMP velocity, \mathbf{v} , with the outgoing recoil direction $\hat{\mathbf{r}}$,

$$\mathbf{v} \cdot \hat{\mathbf{r}} = \sqrt{\frac{m_A E_r}{2\mu_{\chi A}^2}} \equiv v_{\min}, \quad (5)$$

which can be enforced in the differential cross section with a delta function,

$$\frac{d^2\sigma}{dE_r d\Omega} = \frac{d\sigma}{dE_r} \frac{1}{2\pi} v \delta(\mathbf{v} \cdot \hat{\mathbf{r}} - v_{\min}), \quad (6)$$

where $d\Omega$ is the solid angle element around $\hat{\mathbf{r}}$. The event rate then has the structure,

$$\frac{d^2 R}{dE_r d\Omega}(E_r, \hat{\mathbf{r}}, t) = \frac{\rho_0}{4\pi\mu_{\chi p}^2 m_\chi} (\sigma_0^{\text{SI}} F_{\text{SI}}^2(E) + \sigma_0^{\text{SD}} F_{\text{SD}}^2(E)) \times \hat{f}(v_{\min}, \hat{\mathbf{r}}, t), \quad (7)$$

where the velocity distribution now enters in the form of its Radon transform [48, 49],

$$\hat{f}(v_{\min}, \hat{\mathbf{r}}, t) = \int \delta(\mathbf{v} \cdot \hat{\mathbf{r}} - v_{\min}) f(\mathbf{v}, t) d^3\mathbf{v}. \quad (8)$$

The characteristic angular structure of the DM flux on Earth is the reason why directional detection could

be such a powerful means to discover DM. The primary signal is a dipole anisotropy towards the direction $\hat{\mathbf{r}} = -\hat{\mathbf{v}}_{\text{lab}}$, leading to an $\mathcal{O}(10)$ forward-backward asymmetry in the number of recoil events. The strength of this dipole means that in ideal circumstances (*i.e.* perfect recoil direction reconstruction) an isotropic null hypothesis for the recoil direction distribution can be rejected at 90% confidence with only $\mathcal{O}(10)$ events [43, 50]. With $\mathcal{O}(30)$ recoil directions it becomes possible to point back towards Cygnus and confirm the galactic origin of the signal [44]. Secondary signals such as a ring feature at low energies [51], and the aberration of recoil directions over time [52], may also aid in the confirmation of a DM signal, as well as for characterizing astrophysical properties of the DM halo.

Existing limits and projections for CYGNUS

Figure 3 shows a selection of constraints from direct detection experiments along with our headline result: the WIMP reach of CYGNUS. Constraints exist for WIMPs with masses larger than $\sim 1 \text{ GeV } c^{-2}$ and SI cross sections larger than $\sim 10^{-46} \text{ cm}^2$. Underneath these limits lies the neutrino floor, below which

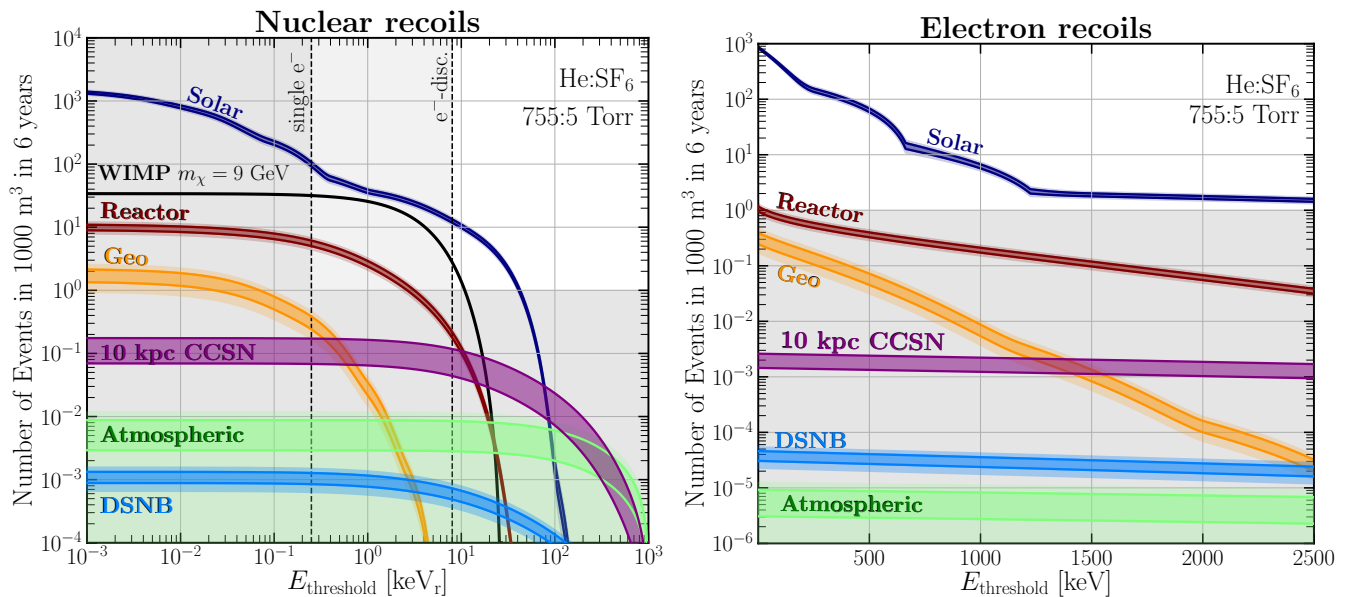


FIG. 4. Number of neutrino-nucleus (left) and neutrino-electron (right) recoil events observed in a CYGNUS-1000 m³ detector filled with atmospheric pressure He:SF₆ at a 755:5 Torr ratio (the event rates are summed over each target nuclei). We calculate the expected number of observed events by integrating the event rate for each background component above a lower energy threshold $E_{\text{threshold}}$. The background components are shown as darker and lighter shaded regions indicating the 1 and 2σ uncertainties from the predicted flux. For comparison we also show the nuclear recoil event rate expected from a $m_\chi = 9 \text{ GeV } c^{-2}$ WIMP with a SI WIMP-proton cross section of $\sigma_p^{\text{SI}} = 5 \times 10^{-45} \text{ cm}^2$ as a black line. For the reactor and geoneutrinos, we assume the entire 1000 m³ is located at Boulby, UK. The purple region indicates the range of expected numbers of events from the neutrino bursts from 11–27 M_\odot core-collapse supernovae located 8 kpc away from Earth. To add further clarity we shade in gray parts of the plot which give fewer than one event in this exposure. In the left panel we also show as dashed lines the 0.25 and 8 keV_r best-case and worst-case thresholds respectively.

WIMP models are rendered practically unidentifiable due to the saturation of their signal by the background from CE ν NS (to be discussed in Sec. II B 2). The limits we display are from the following experiments: CRESST-II [3], CDMSLite [4], COSINE-100 [29], EDELWEISS-III [5], PICO-2L [6], PICO-60 [7], DarkSide-50 [53], PandaX [9], XENON1T [11], PICASSO [54], KIMS [55] and COUPP [56]. For comparison, in the SD case we also display two competitive (but model dependent) limits from neutrino telescopes SK [57] and IceCube [58]. These limits are based on searches for the annihilation of WIMPs captured by the Sun. In the SI case the closed detection regions correspond to the disputed DAMA/LIBRA [24] annual modulation signal.

Also in Fig. 3 we display the final result of this paper: the potential sensitivities of CYGNUS. The gas mixture is assumed to be He:SF₆ at 755 and 5 Torr respectively. This means that the relevant target nuclei are ¹⁹F and ⁴He for the SI search, and just ¹⁹F for the SD search¹. We have assumed a detector with angular resolution, head/tail recognition efficiency, energy resolution,

and charge detection efficiency derived for a strip-based readout (see Sec. IV). The shaded regions indicate limits when varying the hard lower cut imposed on recoil energy. These range from 8 keV_r (above which the electron background can be definitively rejected by factors greater than 10^3 – 10^4), down to 0.25 keV_r (the smallest physically possible threshold as it corresponds to the detection of a single electron). We emphasize that 8 keV is a worst-case scenario. As we will show, we expect electron rejection to be possible down to much smaller recoil energies. We show several possible volumes, including the benchmark 1000 m³ as used in our background study (see Sec. V) and a further-future 100k m³. The former would begin to break into the neutrino floor at low masses whereas the latter would allow us to further characterize the neutrino background and study in greater detail a possible WIMP signal if detected. Even an earlier-stage 10 m³-scale experiment would already be able to set the most sensitive limits on SD-proton WIMP-nucleus interaction, thanks to the large number of fluorine nuclei. Beneath the largest 100k m³ cases we also show a hypothetical ‘search mode’ limit. This corresponds to a mode in which the TPC contains entirely SF₆ at atmospheric pressure. This mode would have no directional sensitivity due to the high gas density, however the potential exclusion limits at high masses would greatly improve. The derivation

¹ Fluorine is an attractive target for SD-proton searches since its most common natural isotope possesses a relatively high value of $\langle S_p \rangle$.

of all the limits displayed here is the subject of the remaining sections of this paper.

2. WIMP detection below the neutrino floor

It was anticipated in early work on direct DM detection that large detectors would eventually become sensitive to CE ν NS [59]. For the keV nuclear recoil energies observed in direct detection experiments solar, diffuse supernovae and atmospheric neutrinos all constitute a significant background for detector exposures beyond the ton-year scale [60–62]. Because neutrinos are impossible to shield against, they represent the ultimate background for the direct detection of WIMPs. Moreover, because the nuclear recoil energy spectra induced by CE ν NS mimic the spectra for WIMPs of certain masses, the discovery of these characteristic masses is limited due to the sizeable systematic uncertainty on the expected neutrino flux. We give more details on CE ν NS and the various neutrino backgrounds in Sec. II C.

The limiting WIMP-nucleon cross section below which experimental sensitivities are impacted by the neutrino background is known as the “neutrino floor” [63]. The shape of the neutrino floor is dependent on the flux of each neutrino background component as well as, importantly, the uncertainty on this flux. The relevant neutrino backgrounds for WIMP searches using nuclear recoils are summarised in the left-hand panel of Fig. 4. The most notable and threatening feature of the neutrino floor is the shoulder below $\sim 10 \text{ GeV } c^{-2}$ arising from the large flux and low energies of solar neutrinos² (see Sec. II C 2). The most important of these are the neutrinos originating from ^8B decay. In a fluorine experiment the nuclear recoil signal due to a $9 \text{ GeV } c^{-2}$ WIMP with an SI cross section around $5 \times 10^{-45} \text{ cm}^2$ is well matched by recoils from ^8B neutrinos. Towards slightly larger masses ($10\text{--}30 \text{ GeV } c^{-2}$) the neutrino floor is set by the diffuse supernova neutrino background (DSNB): the cumulative emission of neutrinos from a cosmological history of supernovae. The expected flux of the DSNB is extremely low ($\sim 80 \text{ cm}^2 \text{ s}^{-1}$ [65]) so the neutrino floor at these intermediate masses falls by several orders of magnitude in cross section. Towards masses beyond $100 \text{ GeV } c^{-2}$ the neutrino floor is induced by the low-energy tail of atmospheric neutrinos from cosmic ray interactions in the upper atmosphere. Atmospheric neutrinos (see Sec. II C 4) are the only significant background contributing neutrino energies above 100 MeV. The low-energy tail of atmospheric neutrinos is difficult to both measure and the-

oretically predict [66] so currently has uncertainties of around 20% [67].

The shape of the neutrino floor is different depending on the target nucleus, because the CE ν NS energy spectrum also depends upon m_A . In this article, unless otherwise specified, we use the definition of the neutrino floor discussed in, for example, Refs. [42, 68]. As opposed to a limit which corresponds to a fixed expected number of neutrino events, we instead show a slightly lower floor which corresponds to the cross sections at which the rate of change in a discovery limit scales the slowest with increasing exposure. This allows us to account for the effect of the neutrino flux uncertainty and properly display which WIMP masses are most severely impacted by each background. Approximately though, the floor can be interpreted as a limit below which an experiment would observe more than ~ 100 neutrino events.

Given that the next generation of ton-scale experiments are expected to become sensitive to CE ν NS, it is pertinent to search for alternative and more powerful methods of subtracting the background. The most basic approach to alleviate the background is to exploit the complementarity between target nuclei of differing masses and nuclear content. For the SI neutrino floor it has been shown that this approach only leads to a marginal improvement in discovery limits, but in the case of SD interactions the differences in nuclear spin contents make complementarity a more viable strategy [69]. It was also shown that the use of event time information also allows the low-mass neutrino floor to be overcome with slightly lower statistics [68, 70]. This exploits both the annual modulation of the DM signal due to the relative motion of the Earth and the Sun, as well as the annual modulation in the solar neutrino flux due to the eccentricity of the Earth’s orbit.

Directionality presents the most attractive prospect for circumventing the neutrino floor because the unique angular signatures of both DM and solar neutrinos allows optimum discrimination between signal and background [35, 36]. This has also been explored in non-gas TPC directional experiments, such as with nuclear emulsions [71], spin-polarized helium-3 [72] and electron-hole pair excitations in semiconductors [73]. The consensus is that: *in a directional experiment there should be effectively no neutrino floor, provided that directionality is well-measured.*

We display the angular distribution of the DM and neutrino-nucleus recoils in Fig. 5. The crucial factor that enables their discrimination is that over the course of the year the Sun does not pass through the constellation of Cygnus. The angular distance between Cygnus and the Sun undergoes a sinusoidal modulation which peaks in September at around 120° and is a minimum during March at around 60° . Because solar neutrino recoils can only point with angles less than 90° from the solar direction, this implies that over long periods during a year there are large patches of sky where the event rate of solar neutrinos is zero (before accounting for the finite

² This shoulder extends down to very low WIMP masses, $< 1 \text{ GeV } c^{-2}$, where the lowest energy solar neutrino sources from pp and CNO reactions become important. These neutrinos are important for DM searches which rely on electronic recoils [64], but are probably out of reach for nuclear recoil analyses for the foreseeable future.

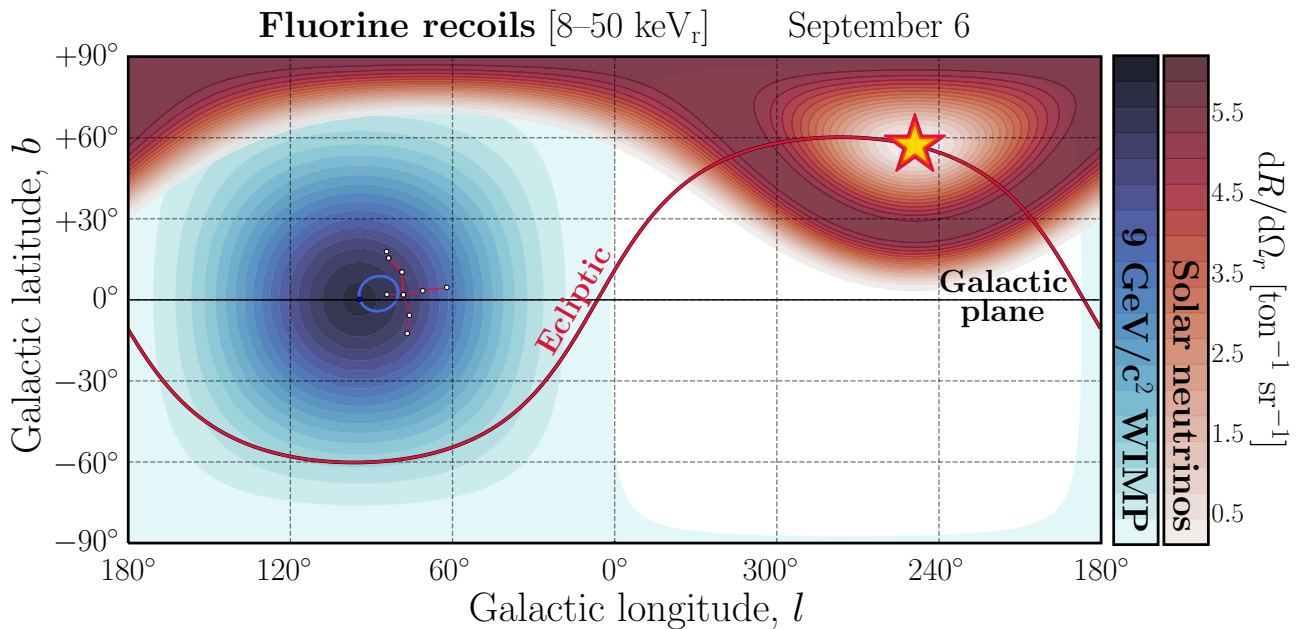


FIG. 5. Total angular distribution across the sky of WIMP-induced (blue contours) and neutrino-induced (red contours) fluorine recoils on September 6. We show the distributions in galactic coordinates (l, b) where the line for $b = 0$ corresponds to the galactic plane. Both distributions have been integrated over recoil energy between 8 and 50 keV_r. We choose a WIMP mass of $9 \text{ GeV } c^{-2}$ and a SI cross section of $5 \times 10^{-45} \text{ cm}^2$ so that its signal is of a similar size to that of the ^8B neutrinos. For reference we also show the ecliptic in red: the path along which the Sun, shown by a star, moves over the year. Towards the center of the WIMP recoil distribution we also show the stars of the Cygnus constellation. The blue line which encircles the star Deneb shows the variation in the peak direction of the DM wind over the year.

angular resolution of the detector). The separation is large enough that it is possible to discriminate the two signals even if the recoil vectors cannot be oriented in galactic coordinates as shown here. This would be the case for experiments in which recoil vectors were only measured after being projected on to a plane [36], or if timing information was unavailable [37].

3. WIMP astrophysics

Predicting WIMP signals requires astrophysical input in the form of the DM velocity distribution, $f(\mathbf{v})$, as well as the local density of DM ρ_0 . While the former is not known concretely, the latter can be determined with astronomical data. Decades of attempts to constrain this value have begun to settle towards $\rho \approx 0.5 \text{ GeV cm}^{-3}$ (see, *e.g.* Refs. [74–78] and Ref. [79] for a review on methods). However the standard value used by experimental collaborations is 0.3 GeV cm^{-3} . Upcoming analyses with the extremely high and precise statistics of the astrometric survey *Gaia* [80] are expected to lead to even more tightly constrained values in the next few years. Since the local dark matter density only amounts to a multiplicative factor which can be absorbed into the (also

unknown) DM cross section³, its precise value is of low importance before a detection is made.

On the other hand the great deal of uncertainty surrounding the velocity distribution of the DM is generally more important to understand. The benchmark model assumed since the very first direct detection experiments were conducted is the Standard Halo Model (SHM), in which the Milky Way’s DM forms an isothermal sphere. The motivation is that it is the simplest model that gives rise to flat rotation curves. The SHM has a Gaussian distribution for $f(\mathbf{v})$ (or Maxwellian distribution for $f(v)$),

$$f(\mathbf{v}) = \frac{1}{(2\pi\sigma_v^2)^{3/2} N_{\text{esc}}} \exp\left(-\frac{|\mathbf{v}|^2}{2\sigma_v^2}\right) \Theta(v_{\text{esc}} - |\mathbf{v}|). \quad (9)$$

Under the isothermal SHM, the dispersion is related to the local rotation speed of circular orbits: $\sigma_v = v_0/\sqrt{2}$. The benchmark used for this speed is the now similarly out-of-date value of $v_0 \sim 220 \text{ km s}^{-1}$. A more recent analysis indicates a value of $v_0 = 235 \text{ km s}^{-1}$ [82]. As is convention, the velocity distribution has been truncated at the escape speed v_{esc} , with the constant N_{esc} used to renormalize the distribution after this truncation. Experimental analyses typically have assumed $v_{\text{esc}} = 544$ or 533 km s^{-1} , the latter being the more recent RAVE

³ Although, see Ref. [81] for a specific case where this is not true.

value [83]. Again with *Gaia* this is undergoing revision, however newly found high speed substructure is introducing additional complexity to its determination [84]. The escape speed in effect controls the highest energy WIMP observable, but due to the exponential suppression this has a marginal impact on the event rate.

Astrophysical uncertainties in the form of $f(\mathbf{v})$ impact the reliability of signal modeling and hence feed into the measurements of DM particle properties [85, 86] and the calculation of exclusion limits [68, 87]. Indeed much effort has been spent in devising methods to integrate out this uncertainty so that accurate limits can be made. See for example the extensive literature on “halo-independent” methods for the calculation of exclusion limits in Refs. [88–94] and many others. Much of this interest has been driven by previous claims that the anomalous DAMA/LIBRA result could be explained by particular choices for the speed distribution.

Directional experiments, however, will offer a novel resolution to the uncertainty on $f(\mathbf{v})$ that is not present in conventional approaches. In addition to having a unique smoking gun signal that is difficult, or even impossible, to mimic with backgrounds, it has been shown that the prospects for directional detectors to *measure* the DM velocity structure greatly exceed that of an equal-standing *non*-directional detector [45, 47, 95]. This is primarily because recoil energy information alone cannot be used to access the full three-dimensional form of $f(\mathbf{v})$ and is instead only sensitive to the one-dimensional speed distribution, $f(v)$. Since our study is largely comparative in nature we adopt the SHM and the out of date astrophysical parameter values for consistency with past studies and limits. However in the future as more ton-scale detectors begin taking data, it has been advised that the SHM be updated to include recent refinements. A summary of these recommendations can be found in Ref. [82].

It should be emphasized additionally that the structure of the local halo of the Milky Way is itself also of great interest. The measurement of anisotropies in the velocity distribution may provide insights into the archaeology of the Milky Way’s formation, as well as the fundamental properties of DM. In particular, directional detectors are well suited to detect kinematically localized substructures such as DM streams [96]. Less prominent velocity substructures such as debris flow [97], asymmetry in the velocity ellipsoid of the Milky Way [82], and the possible influence of the Large Magellanic Cloud [98], will all require many more events to detect, hence very large detectors will be essential. This is also true for very low-mass streams from smaller DM substructures or any other low-level structures that are likely to have undetectable levels of influence on luminous matter; an opportunity to better understand the level of phase space clumpiness and non-Gaussianity of the Milky Way DM halo is therefore a key potential advantage of directional detection.

While it is useful to remain speculative about the kinds of structures one could hope to see in a future exper-

iment, in the advent of *Gaia*, our concrete knowledge about our local distribution is improving. For instance the expanding catalog of nearby populations of clumps and streams as well as the overall kinematic structure of the stellar halo is providing us greater insight into the formation of the Milky Way halo and the construction of its gravitational potential. In fact we already know that at a large chunk of the local DM halo is likely to be part of a highly radially anisotropic feature, variously called the Gaia Sausage or Gaia-Enceladus [82, 99–106]. Some of the most interesting prospects for directional detectors are the tidal streams intersecting our galactic position [107, 108]. These have been discovered recently thanks to *Gaia*, some prominent examples are the local streams S1 [109] and S2 [110, 111], see also the examples studied in Ref. [112]. Streams will be subdominant contributions to the local density ρ_0 , hence will be secondary signals. However their unique characteristics relative to the bulk of the halo DM are almost entirely washed out in their eventual nuclear recoil signals—unless one has directional information [82, 109, 110]. The ability to resolve these substructures is another major advantage of such a detector [96].

4. Particle models and directionality

As well as simply detecting DM, we also require that a future large-scale experiment be able to uncover properties of the particle itself. This is particularly challenging as there are many competing models that may be degenerate with respect to the signals they produce in the usual direct detection schemes. Various classes of particle models give rise to unique directional signals that would go undetected in a conventional experiment. We outline a few of these here.

Inelastic DM (IDM) models are those in which DM can have a lower or higher energy excited state to which it can down- or up-scatter via a nuclear recoil event. IDM models were proposed to reconcile the DAMA annual modulation signal [113]. The availability of excited states and the suppression of elastic scattering means that heavier nuclei are favored, the low energy recoil spectrum is modified, and the annual modulation is enhanced. It has also been shown that IDM models can give rise to enhanced signal discrimination power in directional detectors [114]. This is because the recoils are more focused in the forward direction, since slower WIMPs cannot scatter with enough energy to induce an excited state.

Directional detectors can also disentangle elastic and inelastic scattering events in DM models that allow for both [115], and if the detector is equipped with any photodetector technology then the luminous decay of the excited state may also be observed [116]. In this latter example, the sensitivity to this process is facilitated by the large volumes of gas-based detectors, as opposed to only their directional sensitivity. This is also the case for proposals for super-heavy DM particles with masses

extending up to the Planck scale and beyond [117].

Another possibility is if DM exists in the form of “darkonium” bound states composed of two or more particles (as is predicted in some configurations of asymmetric DM models). It has been shown that there may be angular signatures observable in directional experiments that can constrain the properties of these bound states [118].

Several years ago a novel scheme to generalize the calculation of signals in direct detection experiments was developed. This uses a non-relativistic effective field theory description of the DM-nucleus interaction to posit a set of operators that describe general processes beyond the simple spin-independent and spin-dependent. The basic operators include all Hermitian, Galilean, and rotation-invariant interactions constructed out of the low energy degrees of freedom involved in a WIMP-nucleus interaction [119]. Certain examples, those which are dependent on the DM particle’s transverse velocity $\mathbf{v}^\perp = \mathbf{v} + \mathbf{q}/2\mu_{\chi A}$, give rise to unique ring-like angular signatures [120, 121]. This means that directional detectors would be more powerful than conventional experiments in distinguishing between these particular operators. An important consequence of these features is that it enables an experiment to distinguish spin-0 DM particles from spin-1/2 or 1 [122]. The information which permits this turns out to be found in the angular dependence of certain effective field theory operators. However, many non-standard operators tend to be suppressed, meaning event rates are inherently low. Large scale detectors will therefore be required for studying the complete phenomenology of DM.

5. Axions

Direct detection experiments searching for WIMPs are not limited to a single class of DM candidate. Another very popular class of candidate is the axion and its generalization, the axion-like particle (ALP). The motivation for axions originates in the dynamical solution of Peccei and Quinn [123] (PQ) to the strong-CP problem of quantum chromodynamics (QCD), see *e.g.* Ref. [124] for a recent review. The ALP is inspired phenomenologically by the QCD axion but could come from a variety of theoretical origins when spontaneously broken PQ-like symmetries are embedded in higher energy theories. This is the case most notably in string theory where both the QCD axion and a slew of phenomenologically similar particles are predicted [125]. The masses of axions and ALPs can span many orders of magnitude but their extremely weak couplings to the Standard Model make them attractive DM candidates. Axions produced non-thermally in the early Universe via vacuum misalignment [126–128], decaying topological defects [129, 130] or in the form of axion stars [131] or miniclusters [132–134] have been shown to be able to match the required properties and cosmological abundance of cold DM (see *e.g.* Ref. [135]). Axions and ALPs are by construction coupled to the Stan-

dard Model through quark loops which gives rise to a number of potentially observable interactions: for example, axion-photon conversion inside magnetic fields, absorption by atomic electrons (the axioelectric effect) and spin-precession of nuclei. In the case of the QCD axion the strength of the coupling and its linear relationship to the axion mass is prescribed by theory, but for the generalized ALP, the coupling may take any value. See Ref. [136] for a recent review of experimental searches for axions.

Existing WIMP direct detection experiments such as LUX [137], XENON [138–140], EDELWEISS [141], CDMS [142] and PandaX [143] have already constrained axions and ALPs in the search for their interactions with electrons via the axioelectric effect [144–146]. The coupling most accessible to a WIMP experiment is the axion-electron coupling g_{ae} , as opposed to the photon coupling measured in the most mature axion experiments using resonant cavities such as ADMX [147] and CAST [148] (although in QCD axion models these couplings are related). Note that most of these analyses can also be recast as constraints on vector bosonic dark matter particles, the dark photon being the most notable example [149].

ALPs may be searched for as both a DM candidate as well as solely a modification to the Standard Model. If ALPs with masses in the range $m_a \sim 1 - 40$ keV comprise a significant fraction of the local DM density, then they should stream into a detector and induce electron emission from target atoms with a sharp spectrum located at m_a . Since the spectral width of this signal would be well below the energy resolution of any detector, astrophysical uncertainties might only be resolved with the angular spectrum. CYGNUS may also be particularly advantageous in this regard because of the excellent electron/nuclear recoil discrimination that could be achievable at low energies. It should be mentioned that the keV-mass QCD axion couplings are already ruled out, so low-background recoil experiments would only be able to detect DM in the form of an ALP and not an axion.

On the other hand, these experiments may in the future be able to see the QCD axion, since it is expected that they should be produced in the Sun with \sim keV energies, and their precise incoming flux and spectrum is well understood [150].⁴ As is the case with solar axion telescopes such as CAST, if the ALP mass is much less than a keV c^{-2} the axions are produced relativistically and the signal is dominated by the energy of the solar emission. This means an experiment is consistently sensitive over a large range of arbitrarily small masses. Even in the event of the detection of a DM or solar ALP, a large directionally-sensitive experiment could be novel

⁴ Recently an excess of electronic recoil events below 7 keV was reported by XENON1T [140] with a similar spectrum to the solar axion flux. However the size of the excess would require couplings in violation of astrophysical bounds [151–153]

because, as with WIMPs, there may be unique angular signatures. While directional sensitivity is desirable with regards to dedicated axion experiments [154, 155], in the case of the axioelectric effect this is yet to be studied in detail.

C. Neutrinos

1. Coherent neutrino-nucleus scattering

CE ν NS was predicted over 40 years ago with the realization of the neutral weak currents [156]. This Standard-Model process went unobserved for many years due to daunting detection requirements: \sim keV nuclear recoil thresholds, kilogram to ton-scale target masses, and low backgrounds. Recently COHERENT has made the first measurements of the CE ν NS cross section in agreement with the Standard Model [157–159]. Due to the small weak charge of the proton, the coherence results in an enhanced neutrino-nucleon cross section that is approximately proportional to the square of the number of neutrons in the nucleus. A few years after the CE ν NS prediction, and ironically before the conception of the first DM direct detection experiments, the possibility of using this enhanced process to develop a “neutrino observatory” was put forward [160]. A cornucopia of physics searches was envisioned using neutrinos from stopped-pion beams, reactor neutrinos, supernovae, solar neutrinos, and even neutrinos of a geological origin. See *e.g.* Ref. [161] for a summary of natural sources of neutrino.

Shortly thereafter, the first generation of DM experiments began to search for the scattering of WIMPs, where the signature was a low-energy nuclear recoil. Today the irony lies in the fact that the unshieldable recoils that result from CE ν NS will soon be a source of background for the next generation of DM direct detection experiments. But an experiment that can successfully separate and identify these neutrino events can not only proceed past the neutrino floor, but can also realize the long-awaited vision of a “neutrino observatory”. A detector with directional sensitivity has the potential to do just that.

In CE ν NS, coherence is only satisfied when the initial and final states of the nucleus are identical, limiting this enhancement to neutral current scattering. The coherence condition, where the neutrino scatters off all nucleons in a nucleus in phase, is also only maintained when the wavelength of the momentum transfer is larger than that size of the target nucleus. A high level of coherence across all recoil energies is only guaranteed for low energy neutrinos: less than tens of MeV, depending on the mass of the target nucleus. The Standard Model total cross section for the process can be approximated (neglecting subdominant axial-vector terms that arise from unpaired nucleons) as

$$\sigma = \frac{G_F^2}{4\pi} [A + Z(4\sin^2\theta_W - 2)]^2 E_\nu^2 F^2(E_r), \quad (10)$$

where G_F is the Fermi constant, θ_W is the Weinberg angle and E_ν is the energy of the incoming neutrino [156]. It is evident that the cross section increases with the square of the energy of the neutrinos; however, while the form-factor condition, which comes in as $|F(E_r)|^2$, is easily satisfied for solar neutrinos, the total cross section begins to suffer from decoherence for supernova and atmospheric neutrinos. As can be seen in Fig. 4, a detector with an energy threshold of zero can expect to see several hundred to a few thousand recoils from solar neutrinos per ton-year of exposure, depending on the mass of the target nucleus [160].

The differential cross section with recoil energy can be approximated as [156]:

$$\frac{d\sigma}{dE_r} = \frac{G_F^2}{8\pi} [Z(4\sin^2\theta_W - 1) + N]^2 m_A \left(2 - \frac{E_r m_A}{E_\nu^2}\right). \quad (11)$$

Assuming a ^{19}F target, for example, and a 5 (10) keV threshold for observing nuclear recoils. This results in an expectation of \sim 90 (15) background recoils per ton-year, from solar neutrinos alone [162].

2. Solar neutrinos

On Earth, the most prominent source of neutrinos is our Sun with a total flux at Earth of $6.5 \times 10^{11} \text{ cm}^{-2} \text{ s}^{-1}$ [163]. Due to the eccentricity of the Earth’s orbit, the Earth-Sun distance has an annual variation leading to a modulation in the solar neutrino flux Φ ,

$$\frac{d^2\Phi(t)}{dE_\nu d\Omega_\nu} = \frac{d\Phi}{dE_\nu} \left[1 + 2e \cos\left(\frac{2\pi(t - t_\nu)}{T_\nu}\right)\right] \times \delta(\hat{\mathbf{r}}_\nu + \hat{\mathbf{r}}_\odot(t)), \quad (12)$$

where t is the time from January 1st, $e = 0.016722$ is the eccentricity of the Earth’s orbit, $t_\nu = 3$ days is the time at which the Earth-Sun distance is minimum, $T_\nu = 1$ year, $\hat{\mathbf{r}}_\nu$ is a unit vector in $d\Omega_\nu$, and $\hat{\mathbf{r}}_\odot(t)$ is a unit vector pointing towards the Sun. The directional event rate is found by convolving this directional flux, with the directional cross section for CE ν NS. The cross section with respect to the lab-frame recoil angle θ can be written as [156, 160]

$$\frac{d\sigma}{d(\cos\theta)} = \frac{G_F^2}{8\pi} [Z(4\sin^2\theta_W - 1) + N]^2 E_\nu^2 (1 + \cos\theta). \quad (13)$$

The resulting recoils are thus biased to the forward direction, away from the location of the Sun.

The spectra of solar neutrinos $d\Phi/dE_\nu$ come in various distinct forms depending on the nuclear fusion reaction involved in their production. Neutrinos from the initial proton-proton fusion reaction, pp , make up 86% of the solar emission [164]. Despite the huge flux of pp neutrinos, they yield nuclear recoils well below the threshold of any direct detection experiment; however they would be

the dominant source of electron recoils. Secondary fusion of $p + e^- + p$ and ${}^3\text{He} + p$ produce neutrinos, labeled pep and hep , extend to energies beyond pp neutrinos but with lower flux. There are also two monoenergetic lines associated with ${}^7\text{Be}$ electron capture with $E_\nu = 384.3\text{ keV}$ and 861.3 keV . The latter of these is principally responsible for limiting WIMP discovery for $m_\chi < 1\text{ GeV } c^{-2}$ [165]. At higher energies there are neutrinos due to the decay of ${}^8\text{B}$ which extend up to $E_\nu \sim 10\text{ MeV}$ and within the reach of nuclear recoil WIMP searches, as already discussed. Finally, extending to relatively high energies but with much weaker fluxes, are neutrinos arising from the carbon-nitrogen-oxygen (CNO) cycle labeled by the decay from which they originate: ${}^{13}\text{N}$, ${}^{15}\text{O}$ and ${}^{17}\text{F}$.

The theoretical uncertainties on the solar neutrino fluxes range from 1% (pp flux) to 14% (${}^8\text{B}$ flux). For all except ${}^8\text{B}$, the theoretical uncertainty is smaller than the measurement uncertainty. The theoretical uncertainty originates largely from the uncertainty in the solar metallicity, and in order to establish a self-consistent set of solar neutrino fluxes one must assume a metallicity model. The Standard Solar models (SSMs) of Grevesse & Sauval [166] are generally split into two categories: “high- Z ” and “low- Z ,” based on the assumed solar metallicity. Both models have historically disagreed with some set of observables such as neutrino data, helioseismology, or surface helium abundance [167]. More recent generations of SSMs [168] have a mild preference towards a high- Z configuration, though neither are free from some level of disagreement with the various solar observables. DM detection experiments may shed further light on the solar metallicity issue (see *e.g.* Refs. [165, 169, 170]). The measurement of CNO neutrinos will be essential for this, and may be possible in future DM experiments [170, 171]. Recently CNO neutrinos were measured for the very first time by Borexino [172], though currently with insufficient statistics to resolve the solar metallicity problem. The advantage of directional detection in performing these science goals is the vastly improved background rejection capabilities and the ability to reconstruct neutrino energies on an event-by-event basis.

In Fig. 4 we show the expected neutrino background for all components mentioned in this section, for both nuclear and electron recoils (although the latter do not enter into the main discussion of the paper). As in later examples we assume a 1000 m^3 gas TPC located at Boulby Underground Laboratory⁵ with a $\text{He}:\text{SF}_6$ gas mixture held at atmospheric pressure with a ratio of 755:5 Torr. Again, the justification for this choice of target, volume and site are the subject of the remaining sections of the paper. We show the expected number of events as a function of the lower limit of integration in recoil energy, *i.e.* a hard

threshold but with detection efficiency and resolution ignored.

As we have discussed, solar neutrinos are expected—and indeed desired—to be the dominant background for CYGNUS. In the case of nuclear recoils, the majority of the event rate above reasonable $\mathcal{O}(1\text{--}10\text{ keV}_r)$ thresholds originates from ${}^8\text{B}$ component of the flux. One could expect up to 100 events in six years if the threshold can be lowered to 0.25 keV_r . Electron recoils are sensitive to much lower energy neutrinos, so all components of the solar flux contribute to the background leading to a very large event rate that persists even at high energies. If non-neutrino electron recoil backgrounds are able to be suppressed to low levels (see Sec. V) then one would expect pp and ${}^7\text{Be}$ neutrinos to contribute significantly to the electron recoil signal in CYGNUS. In fact it is possible that these fluxes could be measured to even lower energies than in Borexino. We have shown the spectrum of neutrino-electron recoils for reference and future interest.

The directionality of neutrino-electron recoils is not studied here, but will be crucial to understand if they are to be extracted from the large background of electron recoils from other sources. Nevertheless, the electron recoils at the energies shown in the right-hand panel of Fig. 4 will have very long tracks in a gas target, so their directions should be easily measurable [173, 174]. A full investigation of neutrino-electron recoil signals in CYGNUS will be the subject of follow-up work.

3. Supernovae

Another potential signal of interest in directional DM detectors is the enormous burst of neutrinos from a core-collapse supernova (CCSN), which sheds the binding energy of the resulting compact remnant almost entirely in the form of neutrinos over a timescale of a few tens of seconds. Such a collapse-induced burst is expected a few times per century in the Milky Way. The neutrinos in a supernova burst will have a few to a few tens of MeV of energy, and will include all flavors of neutrinos and antineutrinos with roughly equal luminosity [175, 176].

Dark matter detectors with low recoil energy thresholds are sensitive to supernova neutrino bursts via $\text{CE}\nu\text{NS}$ [177]. The order of magnitude is a handful of events per ton of detector material for a supernova at $\sim 10\text{ kpc}$ (just beyond the center of the galaxy, and close to the most likely distance to the supernova [178]). Observed numbers of events scale as the inverse square of the distance to the supernova. In Fig. 4 we showed the expected number of neutrino events in CYGNUS for such a typical CCSN. We source total neutrino luminosities and moments of the energy spectra from the 1d simulations of Ref. [179] and compute the neutrino spectra using the fitting formulae from Ref. [180]. We show the range of expected event numbers for the explosions of stars with masses between 11 and $27 M_\odot$ which covers most of the range of possible supernova-progenitor

⁵ The Boulby mine is the site of the DRIFT experiment, and is one of several candidate labs for future CYGNUS detectors, see Fig. 2.

masses. For CYGNUS-1000 m³ the typical galactic supernovae at 10 kpc would be just out of reach, however > 1 nuclear or electron recoil events could be observed for supernovae closer than 3 or 0.5 kpc respectively. All of the neutrino burst events would be concentrated around a $\mathcal{O}(10\text{s})$ period. For very close supernovae, at sub 100 pc distances, there may even be detectable pre-supernova neutrino events generated ~ 10 hours prior to the explosion during the dying star’s silicon burning phase [181].

A detection of a supernova explosion via CE ν NS would be valuable due to its sensitivity to the total, all-flavor burst flux. Most other detectors that are currently online are primarily sensitive to either the $\bar{\nu}_e$ (in water and scintillator detectors) or ν_e (in argon and lead detectors) components of the flux [176, 182]. Furthermore, some neutrino spectral information and therefore properties of the supernova could be reconstructed from the measured recoil energy spectrum [183].

The advantages of directionality for the detection of supernova burst neutrinos via CE ν NS are several: first, and most clearly, directional information about the source will be of value to observers in electromagnetic wavelengths and multimessenger channels who want to make prompt observations of the supernova event in real-time. Currently, only detectors able to make directional measurements of elastic scattering on electrons have good pointing ability [184, 185] (and Super-Kamiokande is the only current instance) [186]. Even if there is no obviously bright supernova event (as may be the case for a failed supernova), directional information will be helpful to narrow down the possible progenitors. Finally, if the direction to the supernova is known via astronomical measurements, then the neutrino direction information can be used to reconstruct neutrino energies on an event-by-event basis.

The diffuse supernova neutrino background (DSNB) is also of interest as an astrophysical target [65, 187–189]. A measurement of the DSNB flux would be an independent probe of cosmology, as well as representing an average of supernova bursts. Unfortunately, both the nuclear and electron recoil rates are expected to be immeasurably small in CYGNUS. This can be seen in light-blue curves in Fig. 4; clearly an increase in detector size of several orders of magnitude from 1000 m³ would be needed to realistically observe the signal. Even then, since the flux is expected to be very close to isotropic, isolating the DSNB would be difficult, and is therefore not a primary goal for CYGNUS.

4. Atmospheric neutrinos

Cosmic ray interactions in the upper atmosphere have long been a reliable flux of 10 MeV–PeV neutrinos. Atmospheric neutrinos principally originate from the decays of π and K mesons in hadronic showers created by collisions between high-energy cosmic rays and air molecules. The emission of atmospheric electron and muon neutrinos

and antineutrinos has been well-studied for energies above 100 MeV, and was important historically in the discovery of neutrino oscillations. The spectrum in the 10 GeV – TeV region is roughly a power law with $E_\nu^{-2.7}$. However at low energies, especially $\lesssim 100$ MeV, the geomagnetic field causes suppression in the flux [66, 67]. Current studies of atmospheric neutrinos are based on Monte Carlo simulations such as HKKM [67], Bartol [190] and FLUKA [191], which are informed by cosmic ray data and atmospheric nuclear interaction models. Recent observations by Super-Kamiokande [192] and IceCube [193] measure the spectrum between 100 MeV–10 TeV with uncertainties now comparable to the models, enabling tight constraints on sterile neutrinos [194, 195] and non-standard interactions [196–198]. The limited availability of secondary muon data at low energies make the flux uncertainty for the 10 MeV–1 GeV neutrino tail much larger [62], but a nuclear recoil observatory would be most sensitive to atmospheric neutrinos in this very region. The nuclear recoil rate above these energies is very small due to the falling flux and nuclear form factor suppression. A future measurement of the low-energy flux by a nuclear recoil experiment would be beneficial for the improvement of atmospheric neutrino flux models. This is important as atmospheric neutrinos at these energies will be a background for future studies of the less well-understood diffuse supernova background. As with solar neutrinos, a key advantage of observation via nuclear recoils is the sensitivity to the three-flavor flux, improving the measurement of the total normalization.

In addition to spectral information, the angular distribution of atmospheric neutrinos is also of great interest. The flux is known to peak towards the horizon over the full energy range due to the longer flight path for primary cosmic rays through the atmosphere. The directionality of neutrinos is also controlled by the cutoff in the geomagnetic rigidity (pc/Ze for a particle with momentum p and atomic number Z), which sets the minimum energy required for a cosmic ray to have been deflected to a given direction. For the higher energies when cosmic rays are generally more energetic than the rigidity cutoff, the flux is symmetric in zenith angle about the horizon; however for low energies below the cutoff, the flux is expected to become enhanced in upward-going directions [67]. The spatial dependence of the rigidity cutoff also induces an east-west dipole asymmetry in both ν_e and ν_μ directions as well as a strong dependence in the flux normalization on latitude. The subsequent nuclear recoil distributions therefore also weakly depend on angle: a calculation relevant for directional detectors can be found in [42]. A measurement of these angular features at new locations would serve as an additional test for flux simulations as well as again supplying an enhanced background rejection capability for a directional experiment. The science case for a nuclear recoil observatory and atmospheric neutrinos is therefore well motivated. However the relevant flux at low energies is still extremely small, $\sim 10 \text{ cm}^{-2} \text{ s}^{-1}$ [67], with roughly 10-20% uncer-

tainties (see the green curves in Fig. 4). This would yield a nuclear recoil rate in a 5 keV_r threshold ¹⁹F target experiment around of 0.1 ton⁻¹ year⁻¹. Due to the high energies of the atmospheric neutrino flux, the electron recoil rate on the other hand is negligible. Atmospheric neutrinos may be measurable via nuclear recoils if further refinements to these calculations put the flux in the upper end of the current uncertainty, but this would still be a late-stage goal of a CYGNUS-like experiment, requiring DUNE-sized [199] ~100,000 m³ scale volumes.

5. Geoneutrinos

The radioactive decay of uranium, thorium and potassium in the Earth’s crust and mantle is believed to power a large fraction of the internal heat flow of the Earth [200, 201]. Such decays will also be a significant source of antineutrinos for energies below 4.5 MeV. The impact of these geoneutrinos on DM experiments has been discussed in the past [60], but as a background they are subdominant when compared with the large flux of solar neutrinos at similar energies. Recently however, Ref. [41] explored the science potential of *directional* geoneutrino measurements in particular. Naturally, the flux of neutrinos originating from the Earth will be strongly anisotropic in an upward going direction, approximately azimuthally symmetric and constant in time. As with a DM search, the distinct angular signature and the lack of modulation for geoneutrinos (with respect to solar neutrinos for instance) would give a directional experiment strong background rejection capabilities in this context.

As shown in Fig. 4 (left) the CEνNS recoil rate from geoneutrinos in CYGNUS (assuming a location at Boulby) is expected to subdominant, and difficult to observe without lowering the threshold significantly. However there may be scope for using electron recoils as can be seen in Fig. 4 (right). The event rate for electron recoils from geoneutrinos is around 1–10 per 10,000 m³ in six years. Boulby in fact has a relatively low flux of geoneutrinos compared with alternative sites [202]. The main issue in observing this flux would be the fact that the electron recoil rate is completely swamped by solar neutrinos. This is precisely why directional experiments would be desirable, since the geoneutrino and solar neutrino fluxes are almost completely separable in zenith angle during the daytime, assuming head/tail recognition.

The physics motivation for studying geoneutrinos is substantial. Past measurements of ²³⁸U and ²³²Th geoneutrinos by KamLAND [203] and Borexino [204] have relied on antineutrino capture by protons. However the 1.8 MeV threshold for free proton inverse beta decay has rendered neutrinos from the decays of ⁴⁰K and ²³⁵U undetectable in this manner. A measurement of the threshold-free neutrino-electron scattering from these lower energy sources could tighten the constraint on the radioactive contribution to the Earth’s surface heat flow,

which currently stands at $38 \pm 23\%$ [205]. A better understanding of the heat flow from an improved knowledge of the abundance and spatial distribution of radioactive sources would be invaluable for tracing the thermal history of the planet. Since inverse beta decay measurements have a very weak angular dependence, a directional geoneutrino search would again be particularly advantageous in this regard. In Ref. [41] it was found that a 10 ton-scale detector operating for 10 years would be capable of a 95% CL measurement of the ⁴⁰K flux. With even larger detectors the angular recoil spectrum of geoneutrinos would provide insight into the source of the Earth’s magnetic field. Crucially this requires knowledge of the composition and distribution of radioactive elements in the core. Distinguishing neutrinos from the core and the mantle is not possible without directional information. Detecting geoneutrinos will therefore be an important secondary goal of a CYGNUS-100,000 m³ scale experiment, particularly if this total volume is distributed over multiple sites.

6. Science with source and detector

Artificial sources of neutrinos can be used in conjunction with a CEνNS-sensitive detector for multiple physics purposes. Potential artificial sources include— in approximately increasing order of neutrino energy— radioactive sources (typically ~MeV or less), nuclear reactors [206] (several MeV), isotope decay at rest [207] ($\lesssim 15$ MeV), stopped-pion neutrinos [208] (up to 52 MeV) and low-energy beta beams [209] (tunable, up to tens of MeV or more). Many of these were first proposed in Ref. [160] in 1983. For neutrino-nucleus interactions to be dominated by CEνNS, the momentum transfer Q must be less than the inverse nuclear size, which is largely the case for medium-size nuclei for neutrino energies up to around 50 MeV.

Physics possibilities with an artificial-neutrino-source CEνNS experiment are extensive [208]. They include Standard Model tests [210–214], neutrino electromagnetic properties [211, 215, 216], sterile neutrino oscillation searches [217, 218], as well as nuclear form factor [219, 220] and neutron radius measurements [221]. For some experimental setups, light DM produced at the source can be probed [222–225].

In Fig. 4 we showed the expected number of nuclear and electron recoil events in CYGNUS over six years from nearby nuclear reactors. The rates shown here are relatively large since the Boulby site benefits from its close proximity to the Hartlepool nuclear power station. As with geoneutrinos, the major advantage of a directionally sensitive experiment is in its ability to reject the otherwise irreducible background from other sources of neutrinos, namely the large rate of solar neutrinos at similar energies. As with natural sources, the measured recoil directions of artificial neutrinos, in combination with the known source location would enable event-by-event

reconstruction of neutrino energies and improve background rejection. The improvements brought by directionality would in general enhance sensitivity to any of the array of physics measurements accessible to artificial-source CE ν NS experiments [226].

Of the various artificial neutrino source possibilities, the currently-available ones are reactors, which offer huge fluxes of $\bar{\nu}_e$, and stopped-pion sources. The latter has yielded the first measurements of CE ν NS [157–159] by the COHERENT experiment. Stopped-pion sources produce neutrinos from the weak decays of charged pions at rest. They emit neutrinos with a well-understood spectrum with a maximum energy of half the mass of the muon, 52.8 MeV, overlapping well with typical supernova neutrino energies. These neutrinos are produced copiously at accelerators when \sim GeV-scale protons collide with matter, producing pions; when the predominant π^+ decay after stopping in the matter, they yield monochromatic 30 MeV ν_μ and muons; these at-rest muons decay on 2.2 μ s timescales to create $\bar{\nu}_\mu$ and ν_e with a few tens of MeV. Stopped-pion sources have been used in experiments in the past [227, 228]. Currently the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory, [211, 229] is in use for CE ν NS measurements [208], and other potential sources might be available in the future [220]. The SNS produces about 5×10^{14} neutrinos per flavor per second with few-hundred-ns pulse width at 60 Hz, which results in a factor of $10^3 - 10^4$ background rejection. While no current detector there has directional capability, this source would be suitable for such experiments. This setup could serve as a detector technology test, given that the expected recoil distributions are well known.

7. Exotic models

DM experiments will also be able to explore novel neutrino sector physics. The recently measured CE ν NS cross sections [157–159] appears to agree with the Standard Model prediction currently, but there may be additional non-standard interactions that would affect the recoil energy spectra observable in future DM experiments. For example Ref. [170] explored the prospects for ton-scale experiments to perform novel solar neutrino measurements, such as measuring the pp or ${}^8\text{B}$ flux via the neutral current, as well as constraining the running of the electroweak mixing angle and the possible existence of additional mediators from some light dark sector. By modifying the recoil energy spectra, additional exotic interactions involved with both DM and neutrinos will affect the shape of the neutrino floor [230, 231]. Certain mediators may also increase the number of expected events from the neutrino background by several orders of magnitude [232]. This further emphasizes the need for directional experiments with low energy sensitivity. If the CE ν NS cross section is enhanced in this way by non-standard interactions then potentially a much greater

range of WIMP cross sections at low masses are saturated by the neutrino background. It was also shown in Ref. [169] that direct detection experiments would be able to make complementary constraints on sterile neutrinos if both coherent nuclear and electron scattering of solar neutrinos is measurable. Again, the fact that the initial directions of solar neutrinos are known means that a directionally sensitive experiment can reconstruct the neutrino energy spectrum on an event-by-event basis.

D. Summary of science case

A large-scale directional nuclear recoil detector like CYGNUS will offer many opportunities in a great number of research avenues, both in particle physics and astrophysics. We have described the major goals for such an experiment functioning as an observatory for DM as well as neutrinos. As a DM detector, a recoil observatory offers one of the most sensitive tests of the putative WIMP in the sense that it can attain the most powerful background rejection capabilities and confirm the galactic origin on the detected signal. Additionally, we have discussed here that such an experiment will also be capable of searching for other candidate particles for DM, probe non-standard and exotic particle interactions and map the local DM phase space structure of the Milky Way galaxy around us. Furthermore, while we would desirably have the experiment detect the DM, we emphasize that the presence of the unavoidable neutrino background in nuclear recoil experiments brings with it novel opportunities for discovery in its own right.

III. EXISTING DIRECTIONAL DETECTION TECHNOLOGIES

The directions of nuclear recoils can be inferred in a detector by direct or indirect means. A direct reconstruction of the recoil track, such as in a tracking detector, can be achieved if the nuclear recoil geometry is measurable. An indirect measure on the other hand may be possible if there exists some proxy for the recoil direction, such as a detector response that depends upon some angle between the recoil and a detector axis. In this section we summarize the various available readout methods based on these two broad ideas. In doing so we outline and provide references to the experimental data that supports the simulation parameters used in the following Sec. IV. For a more detailed and critical assessment of these technologies, including those not under consideration for CYGNUS, see Ref. [15].

A. Detectors with recoil track reconstruction

All currently active directional experiments aim to reconstruct the geometry of recoil tracks. Most of these

make use of a low-pressure gas time projection chamber (TPC), in which the mm-scale track geometry is measured in 1, 2, or 3 dimensions. In addition to gas-based TPCs, track reconstruction at the 100 nm scale has been demonstrated in solid-state nuclear emulsions [233]. More exotic, and at this point unvalidated, technologies such as nitrogen-vacancy centers in diamond [234], DNA strands [235], and columnar recombination [236] have also been proposed.

1. Gaseous TPCs

The low-pressure gas TPC is the most mature directional detection technology. In this scheme, the WIMP-induced recoil generates a track of ionization in the gas volume, and an electric field transports the resulting charge to an amplification and readout plane. The full three dimensions of a recoil track can be reconstructed by combining the 2d measurement of the ionization charge distribution on the readout plane, with the third dimension inferred by sampling the transported signal as a function of time. The projection of the track along this third dimension, parallel to the drift field, is found by multiplying the duration of the signal with the known drift velocity of the charge in the gas. In some designs, the ionization electrons can be transported directly. In others, a target gas with high electron affinity is used where the ionization electrons rapidly combine with surrounding gas molecules near the interaction site to form negative ions which instead drift to the readout plane. This latter method, so-called ‘negative ion drift’ (NID), can help suppress the diffusion of the ionization track to the thermal limit and thus preserve more of the track geometry prior to readout [237]. Recent work has explored NID operation at atmospheric pressure as well [238].

Given the $\mathcal{O}(\text{keV})$ energies expected for WIMP-induced nuclear recoils, and the typical energy required to create an electron-ion pair in gas targets, $W \approx 20 \text{ eV}$, a recoil will produce $\mathcal{O}(10^2\text{--}10^3)$ primary ionization electrons. To enhance this signal, a gas amplification device is used, consisting of a carefully designed region of high electric field where avalanche multiplication occurs. In some cases, this amplification device is integral to the readout—Multi-Wire Proportional Chambers (MWPCs) and Micromegas for instance—while in others, the gas amplification and readout are distinct, like with silicon pixel chips or an optical readout. Indeed, multiple gas amplification techniques, and readouts, may be used together in the same detector. The recoil measurement may combine charge readout with other detection channels. For example, if the gas target scintillates during amplification, then optical readout can be used with, or in place of, charge readout.

a. Multi-Wire Proportional Chambers Early work in directional detection used MWPCs, which provide both gas amplification and spatial information. The

DRIFT collaboration pioneered this method [239, 240], and currently holds the leading DM limit set by a directional detector [241, 242]. DRIFT-IIId achieved zero-background operation over more than 100 live-days using a low-pressure gas mixture (30:10:1 Torr of $\text{CS}_2:\text{CF}_4:\text{O}_2$) in a back-to-back TPC configuration, each with drift length of 50 cm. An energy resolution of 40%, a threshold for head/tail recognition for fluorine recoils (in a statistical sense, not event-by-event) of 40 keV_r , and a fluorine recoil detection energy threshold of 20 keV_r , were all achieved experimentally. The energy threshold was ultimately limited by the achievable gas amplification of $\sim 10^3$ [243, 244]. The spatial resolution in the readout plane is given by the amplification wire spacing, or pitch, of 2 mm. Mechanical instabilities arising from inter-wire electrostatic interactions prevent smaller wire spacings. The spatial granularity along the drift direction, on the other hand, is substantially finer: $60 \mu\text{m}$, for a 1 MHz sampling rate and negative ion drift speed of $60 \mu\text{m}/\mu\text{s}$. In an MWPC, the capacitance per wire is very low ($\ll 1 \text{ pF}$ for a wire length of 1 m), and the noise will depend on the readout electronics. For example, using front-end ASICs developed for LAr TPCs [245] coupled to a custom digitizer, a noise of 250 e^- with a $10 \mu\text{s}$ averaging time has been demonstrated [246]. Based on previous experience, a $2 \times 2 \text{ m}^2$ MWPC would cost approximately US\$15k.

b. Micro-Patterned Gas Detectors with strip readout The development of micro-patterned gas detectors (MPGDs) [247, 248] has enabled improved spatial resolution relative to MWPCs, which enhances the track reconstruction capability of gaseous TPCs. MPGDs are defined as high spatial granularity gaseous devices with sub-mm gaps between their anode and cathode electrodes, fabricated using microelectronics technology. Examples include the $\mu\text{-PIC}$ [249], and the Micromegas [250].

Large-area MPGDs have become commodity items. For example, a $2 \times 1 \text{ m}^2$ resistive-strip micromegas with 2d strip readout ($200 \mu\text{m}$ pitch) is available for purchase from CERN for US\$30k [251]. The measured strip capacitance is 500 pF per meter [251]. A vigorous R&D effort is underway to produce low-background MPGDs for rare-event searches (see *e.g.* Ref. [252]).

MPGDs with segmented readout are in use for directional dark matter detection already. For example, the MIMAC collaboration [17, 253] uses a low-pressure electron-drift gas mixture of CF_4 , CHF_3 , and C_4H_{10} (70%/28%/2% at 50 mbar) in a back-to-back TPC configuration with 25-cm drift length. MIMAC demonstrated an event detection threshold of 1 keV_{ee} , and an energy resolution of 10% at 5.9 keV_{ee} for a gas amplification of 20,000 [254]. In addition, the NEWAGE collaboration [255], uses a $\mu\text{-PIC}$ with $400 \mu\text{m}$ pitch and a readout sampling time of 10 ns. To enhance the gas amplification, a Gas Electron Multiplier (GEM) [256] is used, for a typical total gas gain of 6000. NEWAGE has run a TPC with 40 cm drift length and CF_4 gas

at 100 Torr. They have achieved an energy resolution of 16% (σ/μ), and an angular resolution of 40° at their recoil energy threshold of 50 keV_{ee} with an underground detector [255]. NEWAGE has also demonstrated head/tail sensitivity down to 75 keV_{ee} with a test chamber [257].

c. MPGD with pad readout Driven by the needs of large-scale LAr neutrino detectors, the community has seen new developments in large-area MPGDs. For example, a segmented readout plane with integrated readout electronics has been developed [258]. The pad pitch of 3 mm is large compared to the needs for directional DM detection, but the pad geometry could be optimized in the future. The model considered here has a pitch of 3 mm, a pad capacitance of 0.25 pF, and an rms readout noise of $375 e^-$ in a $1 \mu\text{s}$ sampling window, with an estimated cost of US\$5k per square meter readout plane [259].

d. Pixel chips Even finer spatial resolution than MPGDs can be obtained with pixelated silicon chips. Examples of pixel-based readouts include the FE-I4b ATLAS chips [260], QPIX [261] and Medipix [262] pixel chip families, which provide spatial granularity as fine as $50 \mu\text{m}$. Rapid digitization of the readout plane provides spatial information along the drift direction. The fine spatial resolution and low per-pixel capacitance (0.01–0.2 pF) are clear strengths of this technology, though a drawback is the small size of each chip (the current-generation ATLAS pixel chip is $2.0 \times 1.68 \text{ cm}$). Building a square-meter area readout plane would be highly costly and complex, though certainly possible. For example, the ATLAS Phase II pixel detector will employ enough chips to cover more than 10 square meters of readout area [263]. R&D work has demonstrated the capability of pixel chips for directional DM detection. For example, in a TPC filled with one atmosphere of He:CO₂ (70:30) — an electron-drift gas — and using double GEMs for gas amplification, an energy resolution of 20% and single electron threshold (for primary ionization) was demonstrated [264–267]. Presently, the cost of the ATLAS pixel chips (model RD53B) is US\$90k/m², not inclusive of the cost of a gas amplification device or readout system.

e. Optical imaging The ionization from charge amplification in the gas may be accompanied by the emission of photons, which can be imaged optically (*e.g.* with a CCD or CMOS camera). This was the first approach used in 1994 for directional dark matter detection with TPCs [268]. The technique was then revived in the 2000s [269], and was employed by the DMTPC collaboration [270–272] who used four CCD cameras (Apogee Alta U6, 1024×1024 pixels, $24 \times 24 \mu\text{m}^2$ each) to image a TPC with 27 cm drift length filled with CF₄ at 60 Torr (typical). The proven energy resolution was 35% at 80 keV_r , with a recoil energy threshold of 20 keV_{ee} . The angu-

lar resolution was 15° at 20 keV_{ee} , with head/tail sensitivity above 40 keV_{ee} [271]. The DMTPC cameras had $10 e^-$ readout noise. More recent high-end commercially available cameras (*e.g.* Hamamatsu ORCA-Flash4.0 V3 CMOS, costing US\$9k) could offer substantial performance improvements.

Exposure and readout times of cameras are long compared to the signal generation in the detector, so a disadvantage of optical readouts is that only a 2d projection of a recoil track in the readout plane is measurable. One would need to combine this slow imaging technique with a faster readout (*e.g.* photomultipliers) to recover the third dimension of the track [273–275]. Another drawback is that the ratio of photon production to charge production can be low, *e.g.* one photon per three electrons in CF₄ [276]. More importantly, the geometric acceptance for photons is very small because very few photons produced in the amplification region make it onto the image sensor. Orders of magnitude more gas amplification are therefore required to compensate, relative to charge readout. However, cameras do have advantages over charge-based readouts, in particular a lower burden on radiopurity since the cameras are located outside of the active volume of the detector. Additionally, high 2d spatial resolution is possible, and data acquisition is trivial, *e.g.* a USB cable to a PC.

Optical readouts are not a major focus of this paper. Instead, the CYGNO project [277, 278] is working with and alongside CYGNUS to investigate separately the prospects of optically based readouts and electron drift. CYGNO employs a CMOS camera coupled to a TPC with triple thin GEMs for gas amplification in a 60:40 He:CF₄ mixture at 1 bar [275]. This configuration provides the necessary high gas gain of $\mathcal{O}(10^6)$, with about one photon produced for every ten electrons [279]. An energy threshold of 2 keV_{ee} and 20% energy resolution at 5.9 keV_{ee} has been demonstrated [280]. In LEMOn, the largest CYGNO prototype, the drift length is 20 cm, achievable thanks to the low electron diffusion in He:CF₄. Fiducialization may also be possible because the high spatial granularity can allow the diffusion of ionization cloud to be measured, which in turn is dependent on the absolute track position along the drift direction. Preliminary measurements with LEMOn have demonstrated directional and head/tail sensitivity down to about 20 keV_{ee} [281].

2. Nuclear Emulsions

The nuclear emulsion most well developed for low energy applications consists of silver halide (AgBr) crystals with a 2.7 eV semiconducting band dispersed in a polymer. The crystal grains work as sensors of charged particles by producing several-nanometer diameter silver clusters in response to a nuclear recoil track. The emulsions need to be developed successively with a catalytic process, so that a two-dimensional projection of

the track trajectory onto the surface can be eventually reconstructed with an optical or x-ray microscope. Recently, Nagoya University managed to produce a super-high-resolution nuclear emulsion called a Nano-Imaging Tracker (NIT), with a mean crystal size of 20-40 nm [282]. Currently no head-tail sensitivity has yet been demonstrated in nuclear emulsions.

The Nuclear Emulsions for WIMP Search project (NEWS) is a directional DM experiment with NIT and a fully automated optical scanning system. It demonstrated absolute tracking efficiencies for carbon recoils at 60, 80, and 100 keV of 30%, 61% and 73%, respectively, by implanting collimated monoenergetic ions (*e.g.* C, O, Kr, F, B etc.). With a directional analysis based on the use of plasmon resonance [283] and a prototype microscope, NEWS managed to establish a position accuracy of 10 nm for a single grain [233] for 100 keV C ions. Background discrimination methods are under investigation to suppress the high expected ^{14}C contamination present in the gelatin producing β decays: the current rejection power is 10^{-6} , but the required is 10^{-8} . An R&D experiment of 10 grams of NIT is currently being prepared for installation at LNGS.

Since the measurement of the nuclear emulsion must be performed after the exposure time of the experiment, any directional signal will be washed out by many rotations of the Earth during that time. NEWS have proposed mounting the experiment on an equatorial telescope but at great financial cost. It was shown in Ref. [37], however, that there is still a directional signal present after time-integration, but the strength of the anisotropy is reduced. A factor of 2–3 more events would be required to distinguish the WIMP signal from an isotropic background if the mounting strategy were not implemented.

3. Crystal defect spectroscopy

Nitrogen vacancy (NV) centers in diamond are quantum systems that are highly sensitive to nearby magnetic fields as well as to localized crystal strain. A recent proposal suggests that this emerging technology could be used for directional DM detection [234]. WIMP-induced nuclear recoils in the diamond target would create a damage trail in the crystal. The trail alters the strain pattern which could be measured using spectroscopic interrogation of nearby NV centers. This technology has the benefit of being solid-state, so a large potential target mass. With ultra-fine nanometer-scale spatial resolution, preliminary calculations suggest a sensitivity to a head/tail signature as well, though no experimental work has been reported yet.

4. Graphene targets

Nuclear recoil directions in 3d (bulk) targets often get scrambled through multiple interactions with the sur-

rounding medium. The recoil direction may be more directly measured in 2d targets. Such targets could be fabricated from semiconductor materials in which the excitation energy is on the order of ~ 1 eV, allowing even MeV-scale WIMPs to initiate electronic excitations. A recent proposal suggests that 2d graphene could serve as a directional detector of sub-GeV WIMPs [284]. This is a particularly interesting idea, especially given that no other directional technology can probe this WIMP mass scale. There has not yet been an experimental demonstration of this technology, although it may be possible within the PTOLEMY relic neutrino search [285].

B. Detectors that indirectly determine the recoil direction

1. Anisotropic scintillators

Solid scintillators (*e.g.* NaI and CsI) are commonly used in particle detection, and specifically in DM detection. Because of their large target mass and high nucleon content, they are ideal for SI WIMP searches. Some scintillators, such as ZnWO_4 and stilbene, have been shown to exhibit a response that depends on the recoil ion direction relative to the crystal axes. In principle, this scintillation anisotropy can be used to infer the nuclear recoil track direction without direct reconstruction of the track geometry. Several groups have explored the possibility of using anisotropic scintillators for a directional dark matter search [286–291], though the magnitude of the anisotropy is too small (less than 10%) for anything particularly sensitive. Moreover, none of these studies have demonstrated anisotropic scintillation for nuclear recoils at low enough energies. Therefore, all quoted energy resolutions, thresholds, and general performances are for general detection of alpha, beta, and gamma radiation and not necessarily relevant for discussion in the context of DM.

2. Columnar recombination

When a recoil ionizes the detector medium, a track of electrons and ions is created. If no electric field is present, these ionization products recombine, producing scintillation light and suppressing the charge signal. With an external electric field, the amount of light produced will depend on the relative orientation of the recoil track and the field: a large angle results in a small recombination fraction, which enhances the charge signal relative to the light signal. A precise measurement of the charge-to-light-ratio could be a proxy for recoil track direction [236, 292].

Evidence for columnar recombination with high-energy alpha tracks was observed long ago in dense xenon gas [293]. Recent simulations [294] suggest that it could be present for tracks as short as $2\ \mu\text{m}$ (corresponding to a

threshold of 30 keV_r at a pressure of 10 bar), if the ionized electrons are properly thermalized. This is not possible however in pure Xe due to the lack of inelastic scattering below 7 eV. Early experimental efforts to cool ionization electrons with trimethylamine strongly suppressed the primary scintillation light, so that a columnar recombination measurement was not possible [295]. In liquid argon, however, there has been a marginal detection of the signature of columnar recombination for nuclear recoils with energies above 57 keV [296]. If the effect was shown to be strong enough in liquid noble experiments then it may marginally help in extending the discovery below the neutrino floor at WIMP masses around $100 \text{ GeV } c^{-2}$ and above [42]. R&D work on columnar recombination is ongoing (see *e.g.* Ref. [292]), but given the low technological readiness at this time, we do not evaluate it further here.

3. Carbon nanotubes

Single-wall aligned carbon nanotubes (CNTs) have been recently proposed as a DM target due to their expected anisotropic response to neutral particles [297]. Under the right conditions, when a C ion scatters off of the CNT walls it sees the tube as empty and can travel with nearly no loss of energy: an effect known as channeling. Different orientations of the CNT axis with respect to the DM wind would give different channeling probabilities and therefore produce significantly different C ion currents at the end of the nanotube. The proposed detector concept in Ref. [297] is a brush of CNTs array closed at one end and open at the other, inserted in a (low-pressure) TPC to detect the outgoing C ions down to $\sim 10 \text{ keV}$. An R&D effort is currently underway in Italy to test the channeling hypothesis for neutral particle scattering and the TPC detector approach.

IV. DETERMINING THE DIRECTIONAL SENSITIVITY OF CYGNUS

To determine and optimize the directional sensitivity of a large scale experiment like CYGNUS, we will focus on gas TPCs, which constitute the most mature directional detection technology. In the previous section, we have seen that a large number of TPC readout technologies have already been successfully demonstrated on smaller scales. Our primary goal here therefore, is to determine the most appropriate TPC readout technology to realize the science goals set out in Sec. II.

Previous theoretical work already compared the ability of directional detectors to discover galactic dark matter, and to set limits beyond the neutrino floor [14, 36]. Here we go one step further, by simulating TPCs with specific readout technologies from the ground up, accounting for inevitable physics and detector effects such as nuclear straggling, diffusion of drift charge, readout quantization,

and readout noise. This allows us to obtain an energy-dependent description of the detector performance including angular resolution, head/tail efficiency, energy resolution, and detection efficiency. The gradual variation of the detector performance with recoil energy is ultimately what will determine how well CYGNUS can utilize recoil directions to distinguish between dark matter and neutrinos, or between dark matter and an unexpected background of nuclear recoils. There is also the question of electron background rejection, which is also strongly energy-dependent, and which will determine the recoil energy threshold above which CYGNUS can be expected to remain background free. We will use the TPC simulations to estimate electron rejection factors, which will inform the later discussion (see Sec. V) of backgrounds for a 1000 m^3 scale experiment. To further inform our readout choices for CYGNUS, we will also evaluate the cost of each simulated technology. Ultimately the ideal detector is the one that maximizes science sensitivity per unit cost. This will be the manner in which we arrive at a final readout decision.

A. Choice of gas

The parameter space for gas TPCs is large, as a number of readout technologies, gas mixtures, and operational parameters are possible. Furthermore, for each readout technology the spatial segmentation can be varied, for each gas mixture the pressure can be varied, and there are many operational parameters such as drift field strength, avalanche device voltage, and gas temperature.

Our goal here is to identify the optimal choice of readout, with the goal of uniting investigators in the field to pursue a single, optimal, detector design. This is complicated by the fact that the optimal choice of readout depends somewhat on gas mixture and operating parameters. For example, the gas mixture and pressure affect the maximum avalanche gain and recoil length, which affect the charge sensitivity and segmentation required of the readout, respectively.

To make progress, we partition the problem, and here compare multiple existing charge readout technologies with typical performance parameters while holding the gas mixture and detector avalanche gain constant.

As mentioned in Sec. III A 1, NID gases are well known to improve the performance of dark matter TPCs by limiting diffusion of the drift charge. Moreover, it has also been discovered that minority carriers in NID gases can enable fiducialization in the drift direction. Full 3d fiducialization in NID was first demonstrated with CS_2 gas [298], and more recently with SF_6 gas [299]. This ability to locate events within the detector volume is crucial for background rejection. SF_6 gas also has a number of additional properties that when combined make it a particularly suitable TPC fill gas for a DM search: is non-toxic, non-corrosive, and contains ^{19}F —one of the most powerful targets to set limits on the SD WIMP-

proton cross section. For these reasons, the NID gas SF₆ has received substantial recent interest from the directional dark matter detection community [238, 300–302].

One potential drawback of SF₆, is that it tends to result in approximately two orders of magnitude lower gas avalanche gains than is typical for commonly used electron drift gases. As a result, fixing SF₆ as a design choice biases our comparison of TPC readouts in favor of more sensitive charge readout technologies. This is important to mention because alternative design strategies based on electron-drift gases may also be competitive. In detectors using electron drift, the substantially higher avalanche gain will generally lead to greatly improved charge sensitivity and performance of low mass WIMP and neutrino experiments. With a high-resolution charge readout, 3d fiducialization is also possible in electron drift detectors — albeit with different techniques [264]. However charge tends to diffuse more in electron drift detectors which reduces directionality at the $\mathcal{O}(1\text{--}10\text{ keV})$ energies needed to detect light WIMPs and solar neutrinos. This aspect therefore favors NID gases.

A new approach for this paper — one which is already under experimental investigation by the authors [280, 303] — is the addition of helium to the fill gas. We expect this to greatly improve WIMP sensitivity and directionality at the lowest recoil energies. When used as a scattering target, helium has greater energy transfer than fluorine for particle masses near the proton mass. This means our DM mass sensitivity could be extended down to very light and unconstrained $\sim 1\text{ GeV } c^{-2}$ WIMPs, while increasing the solar neutrino event rate. Helium gas has a much lower mass density than fluorine at the same pressure, so its inclusion does not greatly affect the track lengths of fluorine recoils. Finally, helium ions have lower specific ionization than heavier nuclei and therefore longer tracks, further improving directionality at low recoil energies. For these reasons a He:SF₆ or He:CF₄ mixture may be optimal as it could both improve the directional sensitivity in general as well as extend the dark matter search into unexplored parameter space.

In this work, we have focused on He:SF₆ mixtures, and the final gas mixture was identified iteratively. Three different mixtures were simulated, using the gas parameters in Table I, and the performance was evaluated. The angular resolution for nuclear recoils depends strongly on these gas parameters (see Ref. [267], Eq. 5): the angular resolution is proportional to the readout point resolution and inversely proportional to the track length. The point resolution in turn grows with diffusion, and the track length is inversely proportional to the gas density. Hence angular resolution improves with smaller diffusion and with lower mass-density. In addition, recoil length at fixed energy also increases as the nuclear mass is lowered. Hence angular resolution can also improve when low-mass target nuclei are introduced, as long as other gas parameters do not degrade.

The first gas simulated, pure SF₆, has been used extensively experimentally, but we found here that the simu-

lated fluorine recoils do not have satisfactory angular resolution at low recoil energies relevant to WIMP sensitivity near the solar neutrino floor. We then added helium to arrive at a 740:20 Torr He:SF₆ mixture. This leads to two improvements: The helium recoils have substantially better angular resolution than the fluorine recoils, and an atmospheric pressure gas mixture would negate the need for low-pressure vessels, reducing cost. We found, however, that the addition of helium increased the mass density to a point where angular resolution for low energy recoils was still worse than desirable. For that reason, we performed a third set of simulations with a 755:5 Torr He:SF₆ gas mixture. This mixture has the combined benefits of low mass density, high helium content, and atmospheric pressure, enabling good directional performance at low recoil energies while keeping cost low.

The final readout comparison and sensitivity analysis here is therefore performed with the 755:5 He:SF₆ gas mixture. Due to substantial computing time required to perform detector simulations, not all the other studies below have been performed for all gas mixtures. Results with pure SF₆ and 740:20 He:SF₆ should be considered lower limits on performance.

Improved performance and sensitivity per unit cost is likely to result from further optimization of the gas mixture. The final choice of gas mixture is also constrained by operational stability, which cannot be reliably simulated. An experimental investigation of gases suitable for directional recoil detection, for electron drift gases at atmospheric and sub-atmospheric pressures, can be found in Ref. [266].

B. Simulation

Our simulation of a directional detector consists of the following stages: (1) nuclear and electron recoil momentum vectors are generated from the expected distribution of WIMP and neutrino event rates; (2) primary ionization distributions are generated for each generated recoil event; (3) the primary ionization is propagated through the detector; and finally (4) the simulation of the gain stage and charge readout. Each stage is described in more detail below. Once we have obtained a simulation of the signal measured by each readout technology under consideration, a fit to obtain a recoil direction for each event is performed, and the result is put through an analysis framework to derive physics results. These latter steps will be detailed in the subsequent sections.

1. Momentum vector generation

WIMP nuclear recoil vectors are generated assuming a Gaussian velocity distribution Eq. (9), with circular velocity $v_0 = 220\text{ km s}^{-1}$, escape speed $v_{\text{esc}} = 533\text{ km s}^{-1}$ and DM density $\rho_0 = 0.3\text{ GeV cm}^{-3}$ at the solar system. For SI recoils, we assume the standard Helm form fac-

Nuclear recoils (20 keV_r)

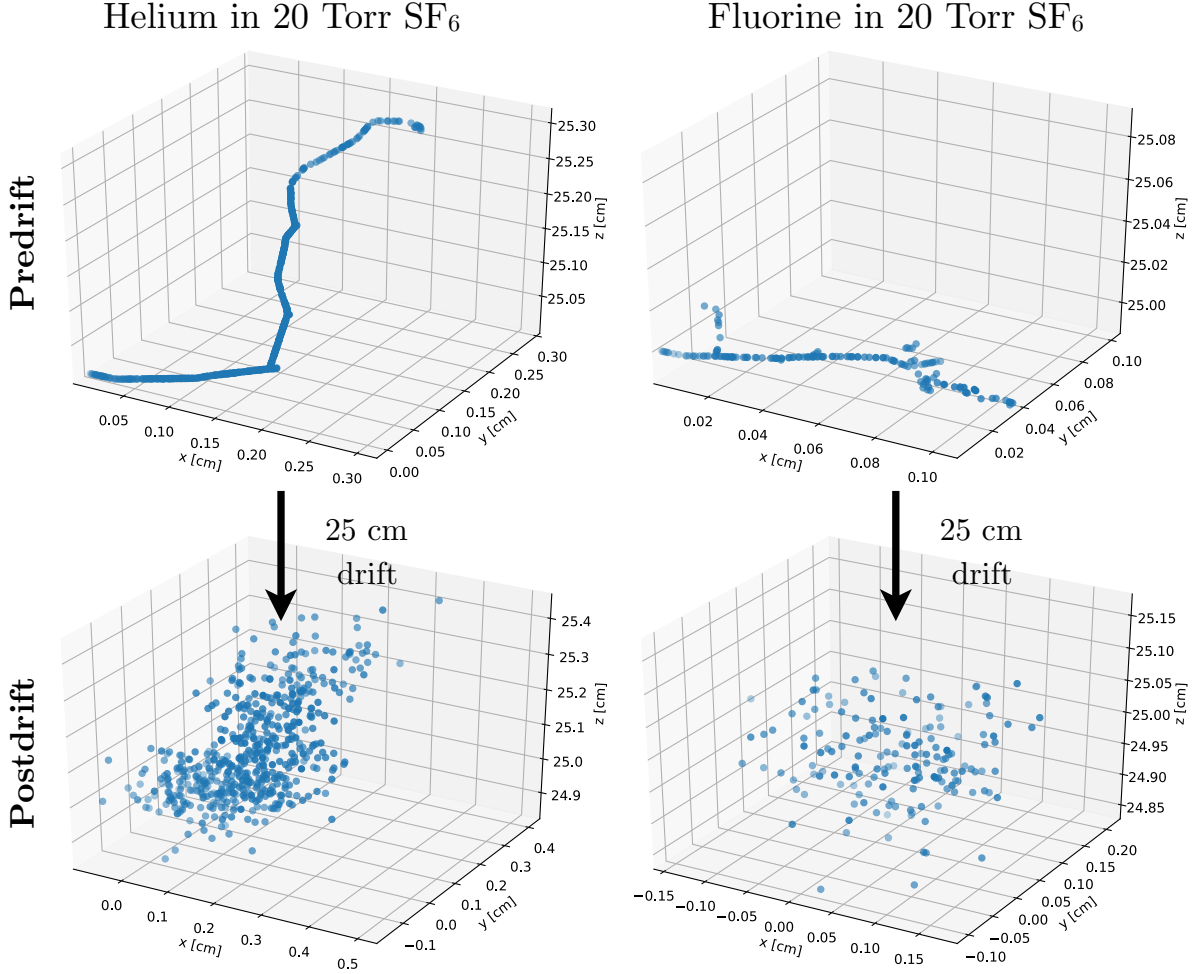


FIG. 6. Distribution of recoil ionization for 20 keV_r helium (left column) and fluorine (right column) nuclear recoil events in 20 Torr of SF₆, generated with TRIM. Both recoils start out at $(x, y, z) = (0, 0, 25 \text{ cm})$ with initial velocity vectors in the direction of the positive x axis. The top row displays the primary ionization prior to drift, and the bottom row displays the negative ion distribution after 25 cm of diffusion. The recoil direction of the helium nucleus is better preserved after diffusion of drift charge than the fluorine recoil. This is mostly because helium recoils are longer than fluorine recoils of the same energy. In addition, helium nuclei suffer smaller energy losses to nuclear collisions, which results both in more visible ionization and straighter recoils.

tor (although in the case of helium this has a negligible impact on the distribution). For fluorine SD recoils, we take the two-body corrected values for $\langle S_p \rangle = 0.421$ and $\langle S_n \rangle = 0.045$ from Ref. [304] combined with a shell model form factor calculation [305]. We assume equal SI and SD couplings to protons and neutrons. Fluorine and helium recoils from coherent elastic neutrino-nucleus scattering, on the other hand, are generated using the analytic expression for the angular differential event rate calculated in Ref. [36]. We again use the Helm form factor and assume a ${}^8\text{B}$ flux normalization in accordance with a global analysis of solar and terrestrial neutrino data [306].

WIMP and neutrino nuclear recoil vectors are generated in the lab (North-West-Zenith) coordinate system at the latitude and longitude of Boulby, UK (54.5534° N , 0.8245° W). We use the full computation of $\mathbf{v}_{\text{lab}}(t)$ and $\hat{\mathbf{r}}_{\odot}$ in this system (see Ref. [14]) which ensures the vector generation automatically includes annual and daily modulation signals for both WIMPs and neutrinos. The events are generated assuming uniform up-time throughout the calendar year. The simulated position distribution of recoils in the detector is uniform. Electron momentum vectors finally, are generated uniformly in the detector and isotropically with respect to

Electron recoils (20 keV)

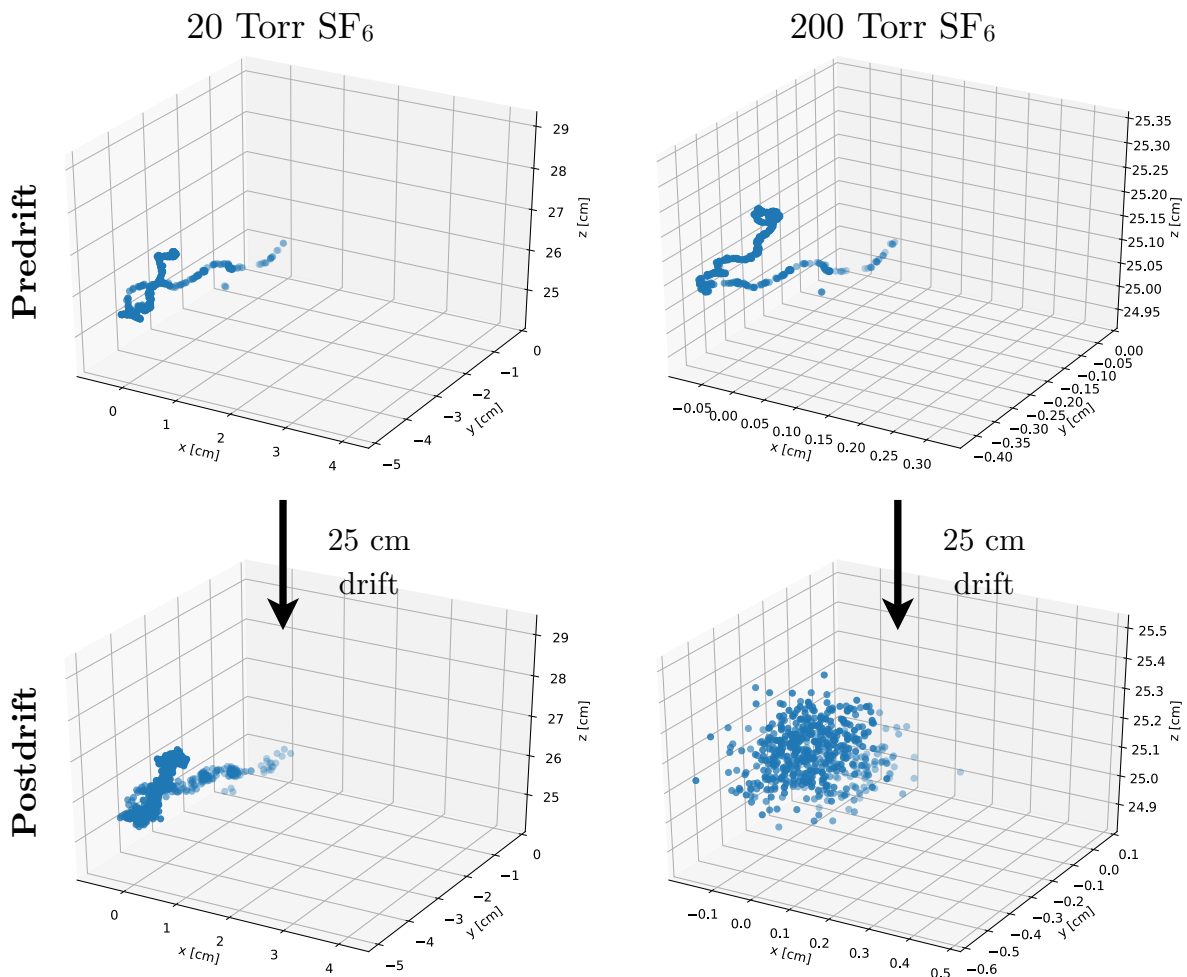


FIG. 7. As in Fig. 6 but now for 20 keV electron recoil events in SF₆, generated with DEGRAD. The left- and right-hand columns compare two different pressures: 20 Torr and 200 Torr respectively (note the different spatial scales between the two columns). Then, as in the previous figure, the top row displays the primary ionization prior to drift, and the bottom row displays the negative ion distribution after 25 cm of diffusion. In 20 Torr, the electron recoil is essentially preserved after the diffusion of drift charge, but when the pressure is increased to 200 Torr the topology is washed out. This illustrates that a gas with low mass density improves not only directionality but also particle identification.

this coordinate system.

2. Generation of primary ionization distributions

To generate the primary ionization distribution of nuclear recoils and electron recoils, we utilize the event generators TRIM [307] and DEGRAD [308], respectively. Both generators take as input the momentum vector of the particle to be simulated, and configuration files that specify the gas mixture. DEGRAD directly outputs a 3d distribution of ionized electrons, while the TRIM output needs to be post-processed. At low energies, ionization

from secondaries cannot be neglected. To estimate the full ionization distribution for individual recoils, we generate TRIM recoils in the most detailed mode and configure TRIM to output all primary and secondary collisions. From the detailed collision output, we reconstruct the trajectories and momenta of all particles in the recoil cascade, from which the 3d ionization is estimated. Further details on this procedure are available in Ref. [309]. The top rows of Figs. 6 and 7 show examples of generated events as they appear at this stage. Figure 8 then shows distributions of resulting event-level quantities that can be computed at this stage: the quenching factor for nuclear recoils (left), and the fitted length of

the charge cloud versus the deposited ionization energy (right). The track lengths versus energy show clear separation between helium, fluorine and electron recoils. This will be important later in this Section when we discuss electron discrimination.

3. Simulation of diffusion and gain

Gas mixture	SF ₆	He:SF ₆	He:SF ₆
Pressure [Torr]	20	740:20	755:5
Density [kg/m ³]	0.16	0.32	0.20
W [eV/ion pair]	35.5	38.0	40.0
Trans. diffusion [$\mu\text{m}/\sqrt{\text{cm}}$]	116.2	78.6	78.6
Long. diffusion [$\mu\text{m}/\sqrt{\text{cm}}$]	116.2	78.6	78.6
Drift velocity [mm/ μs]	0.140	0.140	0.140
Mean avalanche gain	9×10^3	9×10^3	9×10^3

TABLE I. Various gas-dependent parameters assumed in the TPC detector simulation. The values are sourced as follows: the W factor for pure SF₆ is from a measurement with alpha particles [310], while the W factors for the He:SF₆ and He:CF₄ mixtures are calculated using Eq.(1) of Ref. [266]. The diffusion values and drift velocity in 20 Torr of pure SF₆ were measured in Ref. [299]. For the He:SF₆ mixtures, no measurements or reliable simulations exist, so we use the 40 Torr pure SF₆ diffusion from Ref. [299] and then assume the electric field can be adjusted to keep the drift velocity constant. The avalanche gain assumed for pure SF₆ has been achieved with THGEMs in Ref. [311] and triple thin GEMs in Ref. [312], and is also used for He:SF₆ mixtures.

After generating primary ionization distributions, we assume that all electrons attach to fluorine atoms, converting into negative ions, and then drift towards the readout plane. We simulate the combined effect of this attachment, the drift of individual ions, and the subsequent gas avalanche multiplication in the TPC using the diffusion and gain parameters listed in Table I. For pure SF₆, these parameters have already been experimentally established. For the He:SF₆ mixture, although the exact gas parameters have not been measured, SRIM predictions for recoil properties are similar to what is obtained by simulating helium and fluorine recoils in pure SF₆ of the same total density. In the diffusion and gas gain simulation, we assume that by adjusting the electric field strength and the gain stage, the same diffusion, drift velocity, and gas gains as measured in pure SF₆ can be achieved. Experimental work has shown the feasibility of obtaining gain in He:SF₆ and He:CF₄:SF₆ mixtures, at both low (~ 100 Torr [313]) and high (~ 600 Torr [238]) pressure.

The simulated drift length of each electron (negative ion in the case of NID) differs. The distance of the WIMP interaction vertex position to the readout plane is pulled from a uniform distribution ranging from 0 to 50 cm. Each electron (ion for NID) in the generated recoil is then placed at the appropriate position in the detector and smeared with a Gaussian diffusion dependent on its

distance to the readout plane. Edge effects in the drift direction are included by deleting charges that start outside the fiducial volume. Amplification is also simulated at the single electron (or ion) level, using an exponential distribution with mean equal to the avalanche gain. This produces an energy-dependent fractional gain resolution at the event level, which in turn results in an energy-dependent energy resolution, with $\sigma_E/E \sim \mathcal{O}(10\%)$ at 5.9 keV_{ee}, typical for gas detectors. Readout noise broadens this resolution further, and is discussed later. For each gas, the diffusion and gain parameters from Table I are fixed for all readouts to ensure a fair comparison. This means that we are essentially comparing detectors with identical gain stages, but different charge readout technologies. While many combinations of gain stages and readouts are possible, we keep the gain stage fixed so as to focus on how the readout affects the final directional performance.

Figure 6 shows 20 keV_r helium and fluorine recoils in 20 Torr of SF₆ gas. After 25 cm of drift the fluorine recoil direction is rather washed out, but it can still be approximately determined by eye in the case of the helium recoil. Figure 7 shows a 20 keV_{ee} electron event before and after the initial ionization distribution has drifted 25 cm in the detector. For a gas pressure of 20 Torr, the event topology is still preserved after diffusion of drift charge, and can be reconstructed with a detector capable of high-resolution charge readout. For a gas pressure of 200 Torr, such an electron event would have a reduced range, and the topology is therefore mostly washed out by diffusion after 25 cm of drift.

Comparing the bottom rows of Figs. 6 and 7 it can be seen that nuclear and electron recoils with similar amounts of ionization (accounting for the quenching factor) differ both in the length and topology of the charge cloud. In particular the charge cloud from the electron recoil is longer and less uniform. Together these results show that in 20 Torr SF₆, the ionization distribution can be used to determine the recoil direction and discriminate between nuclear and electron recoils—even after charge diffusion from realistic drift lengths. Helium is a preferable target for extending our low WIMP mass reach. When adding helium, we must keep the overall mass density low, to preserve recoil lengths and hence directionality and particle identification capabilities. Here we accomplish this by maintaining atmospheric pressure and reducing the partial pressure of SF₆. An alternative strategy is to operate at sub-atmospheric pressure.

4. Simulation of charge readout

A realistic detector requires a charge readout that is capable of preserving at least part of the topology of the charge cloud. We compare the performance of six charge readout technologies in doing this. We discuss the advantages and disadvantages of each technology below, in order of increasing cost, complexity and performance (the




FIG. 8. Two different features of the primary ionization at generator level. **Left:** quenching factors for TRIM fluorine (blue) and helium (red) recoils in 755:5 He:SF₆ gas, versus recoil energy. Error bars indicate the standard deviation of the quenching factor, resulting from fluctuations of the primary ionization. **Right:** track length versus ionization energy for DEGRAD electron recoils (black) and TRIM fluorine (blue) and helium recoils (red). Track length is defined as the projected length of the primary ionization distribution along a 3d track fit. Neither diffusion nor detector effects are included at this stage.

Readout type	Dimensionality	Segmentation ($x \times y$)	Capacitance [pF]	σ_{noise} in 1 μs	Threshold/ σ_{noise}
planar	1d (z)	10 cm \times 10 cm	3000	18000 e^-	3.09
wire	2d (yz)	1 m wires, 2 mm pitch	0.25	800 e^-	4.11
pad	3d (xyz)	3 mm \times 3 mm	0.25	375 e^-	4.77
optical	2d (xyz)	200 μm \times 200 μm	n/a	2 photons	5.77
strip	3d (xyz)	1 m strips, 200 μm pitch	500	2800 e^-	4.61
pixel	3d (xyz)	200 μm \times 200 μm	0.012 - 0.200	42 e^-	5.77

TABLE II. List of readout-specific parameters that are used in the simulation of each technology we consider here. The capacitance, which determines the noise level, is listed as that for a single detector element. For the optical readout, a yield of 7.2×10^{-6} photons per avalanche electron is used to account for the combined effects of photon yield, geometric optical acceptance, optical transparency, and quantum efficiency.

parameters used to simulate each one are also summarized in Table II).

Planar denotes a simple but cost-effective readout where only the time-dependence of the avalanche charge is recorded, with 10×10 cm segmentation. This results in a signal that is effectively a 1d projection of the charge cloud. One implementation of such a readout would be to directly read the signal from the gain stage, for instance using a GEM with a digitizer. This approach has the advantage of having very low cost per unit area but also has several drawbacks resulting from the low segmentation transverse to the drift axis. This leads to reduced directionality, reduced electron rejection, and high capacitance, which in turn means high readout noise and high energy thresholds.

Wire readouts are traditional MWPCs (see Sec. III). They are of interest because they can be made highly

radiopure, and have very small capacitance and hence low readout noise. One drawback however is that constraints on wire spacing limit the minimum segmentation of the readout plane.

Pad readouts with mm-scale feature size are under consideration for liquid-argon neutrino detectors [258]. The smaller segmentation leads to much lower noise, but also much higher channel counts and higher instrumentation cost than for planar readout. The transverse feature size is still larger than ideal for directional detection of WIMP and neutrino nuclear recoils.

Optical readouts (CCDs or CMOS) can be used to image scintillation light from the amplification region. This approach may lead to lower backgrounds because the cameras can be placed outside optical ports in the gas vessel. Optical sensors provide exceptionally high segmentation and low noise but have comparatively poor

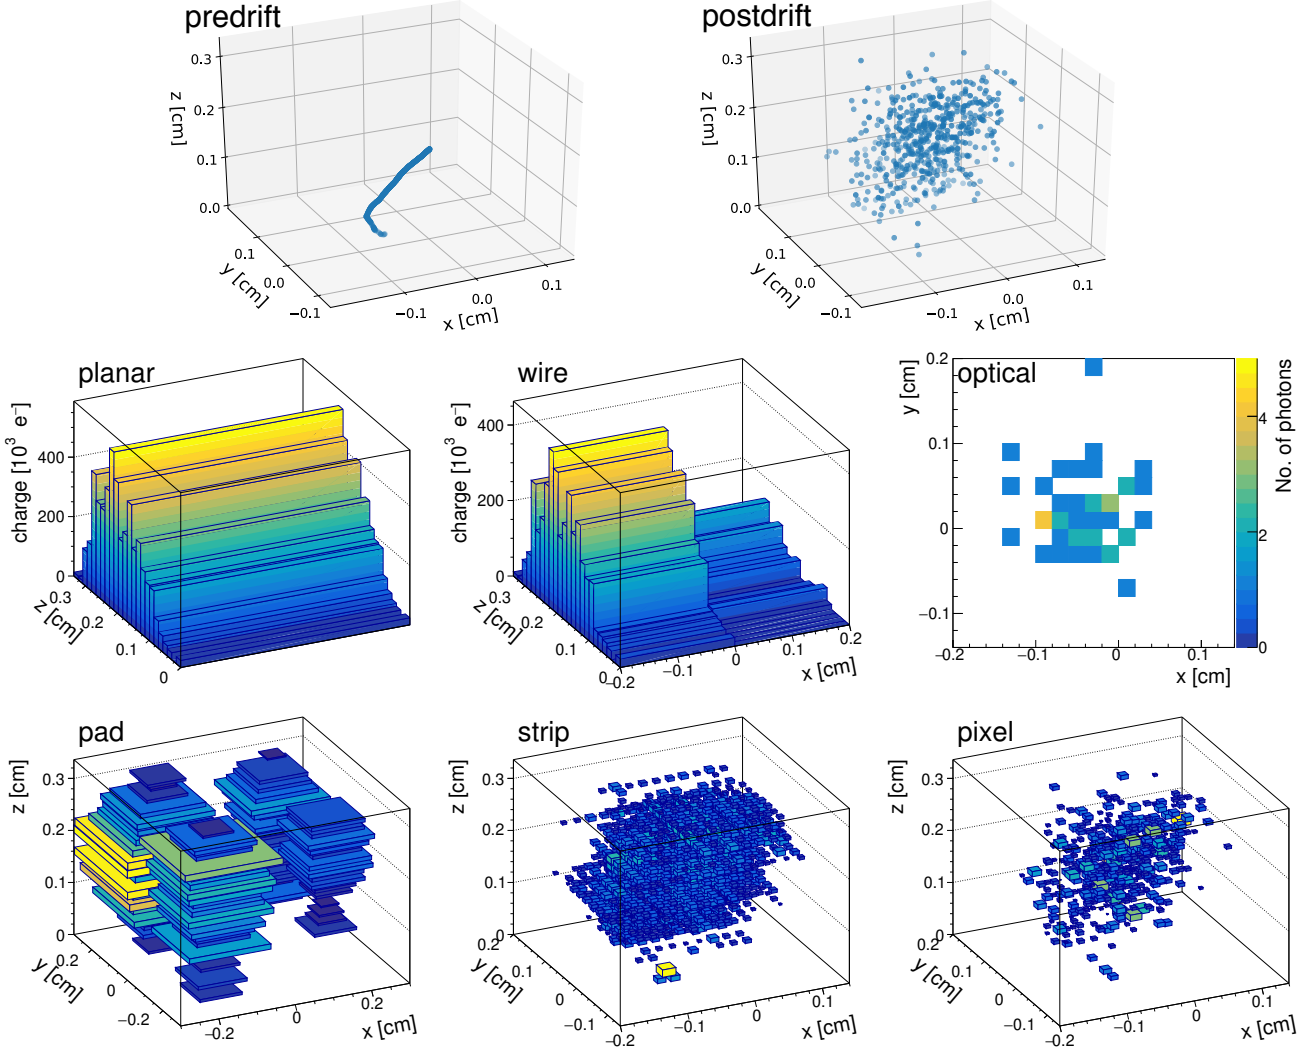


FIG. 9. Simulated 25 keV_r helium recoil event in 740:20 Torr He:SF₆ gas before drift (top left), after 25 cm of drift (top right), and as measured by six readout technologies (remaining plots as labeled). Readout noise and threshold effects have been disabled.

temporal resolution. As a result, such readouts integrate the signal in the drift direction, resulting in a 2d projection of the charge cloud. TPCs with optical readout are cost-effective in that they need only a few cameras per unit readout area, but the effective charge collection efficiency consequently suffers due to geometric photon acceptance losses which can be as severe as $\mathcal{O}(10^{-4})$.

Strip readouts, for example strip micromegas or μ PIC, use the coincidence of signals in orthogonal strips to create 3d hits. Such readouts are a good compromise between performance and cost.

Pixel readouts with application-specific integrated circuit (ASIC) chips generally have the lowest noise and highest performance, but are the most costly. Due to the relatively small $\mathcal{O}(\text{cm}^2)$ area of typical ASICs, the full readout plane for CYGNUS would require a substantial

number of chips and be quite labor-intensive to implement.

We note that readouts can be combined, for instance, Optical 3d readout can be achieved by combining the signal from 2d CCD or CMOS sensors with an independent 1d measurement of the time-dependence of the avalanche charge. See Section IV E.

Our aim here is to compare these quite different charge readout technologies in a fair and unified simulation framework, without invoking any technology-specific signal processing or reconstruction algorithms. In that spirit, we assume that the amount of charge arriving at each detection element can be recorded at 1 μ s intervals. For readout technologies that have already been constructed and characterized, we scale known readout noise levels to 1 μ s. For readouts that do not yet exist,

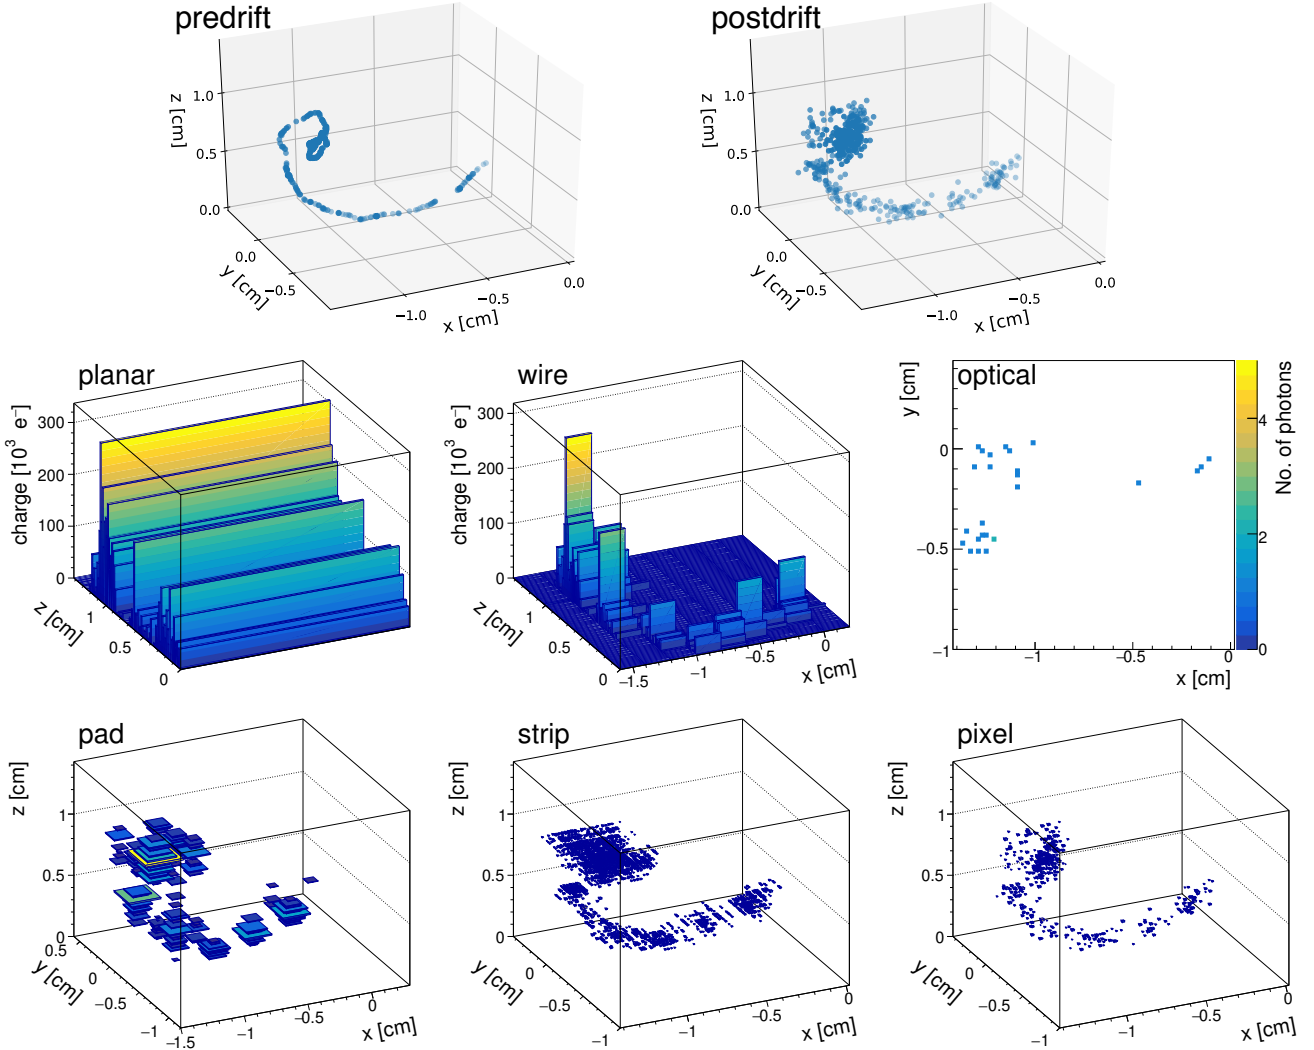


FIG. 10. Simulated 20 keV electron event in 740:20 Torr He:SF₆ gas before drift (top left), after 25 cm of drift (top right), and as measured by six readout technologies (remaining plots as labeled). Readout noise and threshold effects have been disabled.

such as the planar readout, we estimate the readout noise for this time interval from the typical capacitance of a detection element and the noise curves of a candidate charge sensitive preamplifier [245]. Capacitance values and the noise estimates used in the simulation are summarized in Table II. We assume that before event reconstruction a hard threshold is applied to the recorded signal in each channel. This threshold is set so that for each readout the expected rate of noise hits is equal to $10^{-5}/(\text{cm}^2\mu\text{s})$. As a result, the ratio of the threshold to the noise level, $\text{threshold}/\sigma_{\text{noise}}$, varies from approximately three for the planar readout to about six for the pixel ASIC readout. The $\text{threshold}/\sigma_{\text{noise}}$ values resulting from these noise requirements are larger for the more segmented readouts, which agrees with common practice in the field.

Since the optical readout cannot be sampled at the same speed as the other readouts, the same procedure

cannot be used to determine an analysis threshold. In Table II we instead set a requirement of $10^{-5}/\text{cm}^2$ noise hits per *event* on average, but the optical readout performance depends strongly on this threshold. For that reason, a fair comparison of optical and charge readout is not straightforward. In the comparisons that follows, we therefore focus only on charge readout, and have a short separate discussion of optical readout in Section IV E.

Figure 9 shows how a 25 keV_r helium recoil in 740:20 Torr He:SF₆ gas that has drifted 25 cm and is subsequently detected by each of the six readout types. The events are shown after simulation of gain, gain resolution, and spatial quantization into detection elements, using the parameters from Table II. Noise has been disabled for this particular visualization, otherwise noise hits obscure the 3d event displays. Similarly, Fig. 10 shows a 20 keV_{ee} electron in the same gas mixture and

with the same detector simulation. In all cases, the detected event topology or waveform differs visibly from those of the typical nuclear recoil in Figure 9. For a 755:5 Torr He:SF₆ mixture, the main change is that the recoil lengths approximately double, improving directionality and discrimination between nuclear recoils and electron recoils. This suggests electron rejection is achievable in the $\mathcal{O}(1\text{--}10\text{ keV})$ energy range, perhaps even with very simple readouts like the planar. We extend this discussion quantitatively in Sec. IV H.

C. Analysis methodology

In order to analyze the data from the different readout technologies in a uniform way, we turn the simulated output of each readout into 3d spatial points as follows. After the gain and noise simulation, charge above threshold in each detector element and temporal sampling interval is quantized by dividing the charge by the mean gain and turning the results into an integer. This quantized charge is then assigned an absolute (x, y) (transverse to drift) coordinate based on the location of the center of the detector element, and a relative z (parallel to drift) coordinate based in the detection time relative to the first charge above threshold in the event.

Once the reconstructed spatial distribution of the charge cloud has been obtained, its primary axis is found via singular value decomposition (SVD). At this stage the reconstructed primary axis vector has an arbitrary sign, so we only have axial-vector information. The sign of the vector is reconstructed based on the projection of the positions of the quantized charges onto the primary axis. Each charge is assigned to one half of the track, based on whether the charge is closer to the charge with the minimum or maximum projection onto the reconstructed primary axis. We assume that the kinetic energy of the recoiling nucleus being reconstructed is sufficiently low so as to be below the Bragg peak; so we take the track end with the largest assigned charge as the tail. Our analysis methodology neglects all charge below threshold, but it could be partially recovered with a more sophisticated noise suppression algorithm. This would improve reconstruction performance further, therefore all results on the directional performance do stand to improve. However such an algorithm would likely need to be readout-specific, while here we aim to be readout-general to first establish the optimum technology for CYGNUS. The implementation and optimization of a bespoke track reconstruction algorithm would be the next step once this technology is chosen.

D. Low-level performance of electronic readouts

We first look at the efficiency, energy resolution, and directional performance of each readout type, when detecting individual nuclear recoil events. All performance

estimates are made at the end of the simulation chain, after the energy thresholds in Table II have been applied, data have been digitized, and 3d space points have been reconstructed as discussed above. We evaluate the performance in 1-keV steps, and fit the results with analytical functions. These functions will allow us to perform high-statistics simulations of WIMP and neutrino detection scenarios, where the energy spectrum is continuous, without re-running the computationally expensive detector simulation. The parameterized detector performance versus energy may be useful for studies by the wider community, and the fit parameters are available from the authors by request. As mentioned earlier, a fair comparison of optical and electronic charge readout is not straightforward. We therefore compare only electronic readouts. Optical readout is discussed separately in Section IV E.

1. Detection efficiency

Figure 11 shows the event-level detection efficiency of each readout. Here, an event is counted as detected if the reconstructed avalanche charge above threshold is equivalent to at least one electron of ionization before gain, and if at least three space points have been reconstructed. Such criteria should ensure that some minimal directional information can be inferred. A very loose requirement, but one that is simply intended to give an indication of the lowest possible recoil energy sensitivity that could be achieved with each technology at the simulated gain. The energy dependence is obtained by fitting the function $a/(1 + e^{-b(E-c)})$ to the data points via χ^2 minimization. Here, a , b , and c are floating fit parameters and E is the recoil energy.

We find that all readouts except planar have an efficiency higher than 50% above 3 keV_T , even for the low gain of 9000 simulated (with electron drift, gains of order $10^5 - 10^6$ and drastically lower energy thresholds are expected.) The efficiency is primarily determined by the noise level of each readout, so that the pixel readout ends up being the most sensitive at low recoil energies, followed by the pad, wire, strip, and planar readouts. This begins to inform us of the hierarchy in our collection of readouts, but electron background rejection and obtaining good event-level directional information will typically require significantly more detected charge per event, as we discuss below.

2. Energy resolution

Figure 12 shows the energy resolution for each readout. Since each readout has been simulated with the same gain stage, the energy resolution differences seen are due to the charge readout. The energy dependence is obtained by fitting the functional form $\sigma_E/E = \sqrt{(a^2/E^2 + b^2/E + c)}$ to the data points via χ^2 minimization. Here, a , b , and c are floating fit parameters




FIG. 11. Event-level detection efficiency of each readout for fluorine (left) and helium (right) recoils in a 755:5 Torr gas mixture of He:SF₆. Error bars are binominal standard deviations, assuming the efficiency measured is the true efficiency.

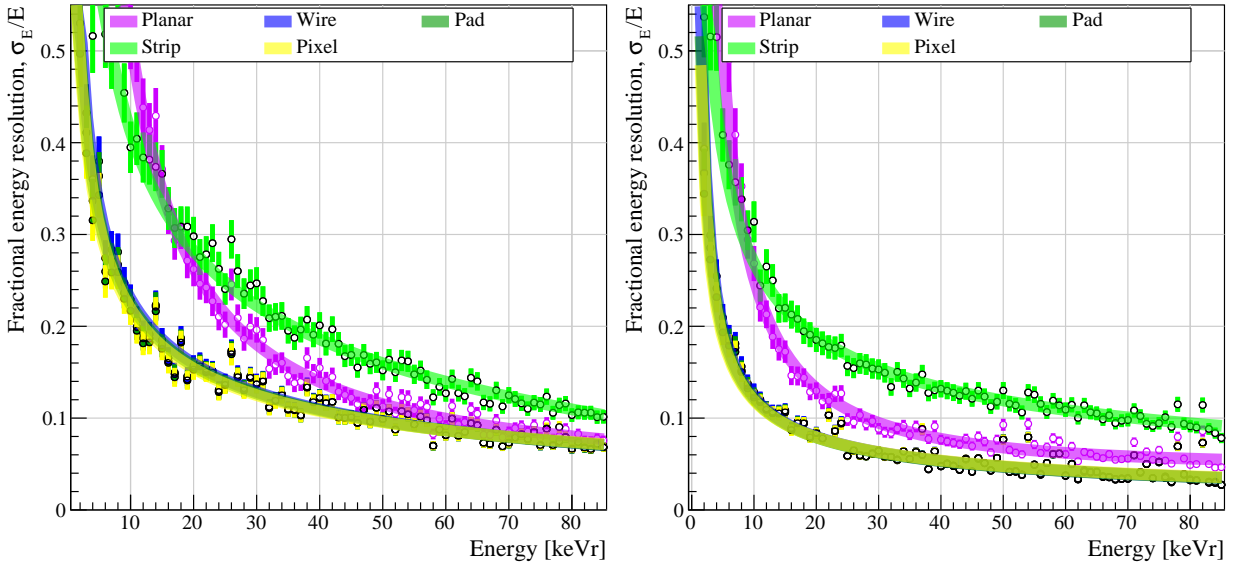


FIG. 12. Fractional energy resolution of each readout for fluorine (left) and helium (right) recoils in a 755:5 Torr gas mixture of He:SF₆. Error bars indicate the RMS variation in the recoil sample analyzed. Note that the curves for the pad readout are very close to the pixel and wire cases so are not easily visible here.

and E is the recoil energy. For pixel, wire, and pad readout, the low readout noise does not noticeably affect the energy resolution. As a result, these three readouts have identical energy resolution, limited by the resolution of the gain stage. For strip and planar readout, the higher noise floor increases energy resolution above the gain resolution. It is possible that this could be improved with a more clever energy determination at the algorithm level. If not, somewhat higher gain than simulated would be

beneficial for these two readouts.

3. Angular resolution and head/tail recognition

We quantify and compare the directional performance of each readout by computing their angular resolutions and head/tail recognition efficiencies.

We define angular resolution to be the mean difference




FIG. 13. Angular resolution of each readout for fluorine (left) and helium (right) recoils in a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean mismeasurement of the recoil axis direction determined from two hundred pseudo-experiments. Error bars show the statistical uncertainty of this mean. Lines are analytical functions fit to the data points, used to parameterize performance versus energy. The resolution that can be obtained with a perfect readout before (black) and after (red) charge diffusion during drift is also shown for comparison.

between the initial recoil axis and the reconstructed recoil axis. Two randomly oriented 3d axial vectors will on average be separated by an angle of 1 radian or 57.3° . Hence this value corresponds to no angular sensitivity. The energy dependence is obtained by fitting the functional form $a/\sqrt{(b + E^c)} + d$ to the data points via χ^2 minimization. Here, a , b , c , and d are floating fit parameters and E is the recoil energy. The angular resolution (Fig. 13) of the pixel readout is superior. For other readouts, the angular resolution gradually gets worse with increasing feature size and noise floor, as expected from the performance ordering in Table II. The performance of the pixel readout is essentially as good as the performance obtained by analyzing the primary charge distribution immediately after diffusion (indicated as postdrift in Fig. 13). This implies that there is no need for any finer segmentation of this readout.

At low energies, the effective resolution obtained by fitting the charge cloud *before* diffusion (pre drift) is substantially better than the post drift performance. This tells us that the performance at low energies becomes limited by the track length being short compared to the diffusion scale. This occurs below ~ 20 keV_r for helium recoils and below ~ 50 keV_r for fluorine. The difference arises because helium recoils are longer than fluorine recoils of the same energy. In this comparison, helium recoils also benefit from two additional effects: they deposit more ionization due to a higher quenching factor, and they exhibit less straggling. The latter limits the angular resolution before diffusion (pre drift). Further improvements may be achievable with a curved track fit

to account for straggling, but this is not explored here.

Figure 14 shows the head/tail recognition efficiency of each simulated readout type versus recoil energy. A random head/tail assignment will be correct in 50% of the cases, so a value of 0.5 corresponds to no sensitivity. The energy dependence is obtained by fitting the functional form $a/(1 + e^{-b(E-c)}) + d$ to the data points via χ^2 minimization. Here, a , b , c , and d are floating fit parameters and E is the recoil energy. In the case of fluorine, the efficiency is quite poor for all readouts, even at the higher recoil energies. We see that before diffusion, some head/tail information is still present in the charge cloud, but most of this information appears to be lost after diffusion. In contrast for helium recoils, the head/tail information at a given energy is much higher both before and after diffusion, and the readouts are able to measure head/tail with some significance down to energies of order 5 keV_r.

Because head/tail recognition and angular resolution at low energies are critical for detecting a WIMP signal, the findings here support utilizing helium as a target nucleus. Even with helium, the directionality gradually becomes diffusion-limited at the lowest recoil energies. This effect starts below 20 keV_r for the angular resolution, and below approximately 35 keV_r for the head/tail recognition. This suggests that the design simulated here is not yet fully optimized. We anticipate further improvements in performance at low recoil energies by lowering the diffusion and/or increasing the recoil length further. Lowering the gas pressure, lowering the drift length, implementing readout-specific reconstruction algorithms, and using hydrogen target nuclei are examples of possible

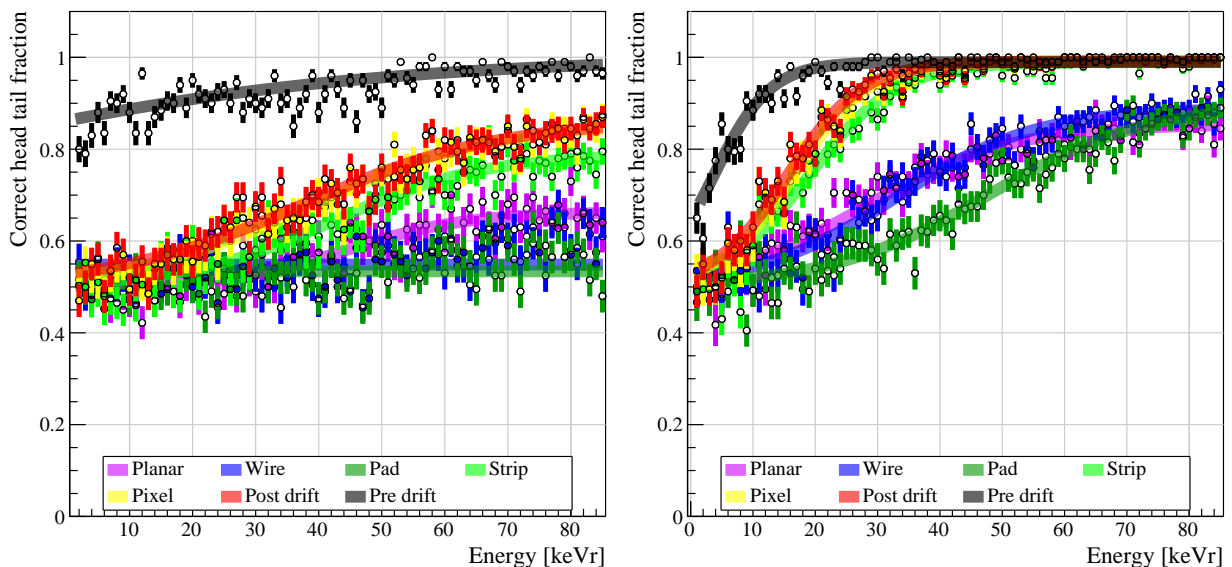


FIG. 14. Head/tail recognition efficiency versus recoil energy and readout type, for fluorine (left) and helium (right) recoils in a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the efficiency determined from two hundred pseudo-experiments. Error bars are binomial standard deviations, assuming the efficiency determined is the true efficiency. Lines are analytical functions fit to the data points, used to parameterize performance versus energy. The performance that can be obtained with a perfect readout before (black) and after (red) charge diffusion during drift is also shown, for comparison.

strategies that should be pursued. Improving this low energy performance will be crucial for maximizing the WIMP/solar neutrino discrimination.

E. Optical readout

For optical readout, we find that the simulated performance depends strongly on the exact value of the optical loss factor, the noise rejection signal threshold, and the reconstruction algorithm. Our overall finding is that optical readout looks challenging with the specific assumptions given in Table II. Optical readout can be competitive, however, if the yield of detected photons per primary electron is increased. This increase can be achieved with improved amplification structures for negative ion drift gases, by utilizing an electron drift gas, or by bringing the optical sensor closer to the amplification plane, thereby reducing the optical loss factor. It is also possible to increase the photon yield via electroluminescence [314].

To get a feel for the capabilities of optical readout with more favorable assumptions, Fig. 15 shows the simulated performance when the loss factor and noise are both disabled in the simulation. Then, 2d optical readout by itself has similar head/tail recognition as charge readout, but reduced angular resolution when compared to three-dimensional charge readout. This is as expected due to the 2d nature of the sensor. The angular resolution can be improved by combining the 2d optical signal with a

1d signal from a PMT. If we assume the PMT perfectly measures the true z -width and the true sign of the z component of the recoil vector (green data points and lines in Fig 15), this 2d+1d optical reconstruction performs quite well in terms of angular resolution, though not quite as well as true 3d reconstruction. Note that the head/tail performance in unrealistically good when assuming this perfect z -measurement. If we simulate a realistic PMT, with performance loosely based on Ref. [275], the performance somewhat degrades, but is still substantially better than optical 2d by itself. It may be possible to improve the optical 2d+1d readout performance substantially with more sophisticated algorithms, which would utilize not only the PMT signal width, but its detailed time structure. This has already been demonstrated by the CYGNO collaboration for electron tracks [275].

If we turn on either the optical loss factor or the signal noise, performance does not degrade substantially. However, if we turn on both effects, the optical 2d signal suffers strongly. For such a scenario, a readout specific noise-rejection algorithm will be required. This is currently being investigated by the CYGNO collaboration [315].

In summary, we find that the optical readout performance is strongly dependent on the algorithm that combines the PMT and optical 2d signal, and on the algorithm that rejects the noise. Optical readout also favors substantially higher gain than what is simulated here. For these reasons we conclude that a fair comparison against charge readout is not easily possible, and leave

the evaluation of optical readout for future work.

F. Directionality threshold

We have seen that there is substantial variation in the energy-dependent angular resolutions and head/tail recognition efficiencies of the different readouts considered. The optimal readout, however, is not the one with the best performance, but rather the one with the best compromise between performance and cost, so that the physics reach per unit cost is maximized. Different signals will have different ranges of relevant recoil energies, so the best compromise may depend on the specific physics goal.

As a first step towards a final cost-performance analysis, we quantify the directional performance versus recoil energy in a way that combines both the angular resolution and head/tail recognition efficiency. This is done by estimating how many recoils each readout needs to detect in order to discriminate between a monodirectional (“signal”) delta function (all recoils have initial momentum vectors in the same 3d direction) and an isotropic (“background”) recoil distribution. The goal here is to provide an intuitive result that clearly shows the recoil energy range where each technology has good directionality. Note that a 3d delta function is a hypothetical scenario with maximum directionality. In any real physical scenario, the recoil distribution is always broadened by the non-zero scattering angles.

Several discovery variables and test statistics have been proposed in the literature for detecting anisotropic recoil distributions for WIMP searches [14]. The most powerful variables are those which depend on the angle between the recoil direction and the direction of Cygnus (or potentially the direction of the Sun when the principal background is neutrinos). Not all of our simulated readouts, however, have a dimensionality which allows these angles to be measured at all times. Our 3d delta function points along the negative z direction, which is the direction that drift charge is traveling in the detector. This is the most sensitive direction for the less segmented readout technologies. Our discovery variable in this case is the polar angle, θ , with respect to the vertical z -axis. Our generated signal peaks at $\cos\theta = -1$, while our background is flat in $\cos\theta$. The *detected* $\cos\theta$ distribution for the pixel readout is shown in Fig. 16 for fluorine and helium recoils. As discussed earlier, the helium recoils are more strongly directional than fluorine recoils of the same energy. This is a consequence of the performance difference we saw for the two recoils species in Fig. 13, and primarily caused by helium recoils producing longer tracks than fluorine.

Also note that the helium mono-directional signal recoil distribution in Fig. 16 is significantly more asymmetric than the equivalent fluorine distribution, because the head/tail is assigned correctly more often for helium, as previously shown in Fig. 14: at 20 keV_r, the simulation predicts a correct head/tail reconstruction of only $\sim 55\%$

in fluorine, while it is $\sim 80\%$ for helium.

In general the reconstructed signal will differ even more from the background at higher recoil energies when the directional performance is better. The $\cos\theta$ signal distribution is also more asymmetric at higher energies. For lower energies than 20 keV_r the contrast between the signal and background distributions becomes gradually less defined, and at the lowest energies simulated here they are essentially indistinguishable. We saw this in Figs. 13 and 14 where all readout curves converged to 1 rad angular resolution and 0.5 head/tail efficiency respectively, corresponding to no directional sensitivity.

To quantify the physics reach corresponding to the directional performance of each readout, we perform a Kolmogorov-Smirnov test [316] comparing each signal and background $\cos\theta$ distribution. Figure 17 shows how many monodirectional signal events are required to reject the isotropic background hypothesis at the 90% confidence level (CL) when there are zero background events present. To fairly compare readouts, we count the number of *interacting* events, meaning that we also count those that go undetected due to finite reconstruction efficiency. For higher recoil energies, as little as two-three events are sufficient to exclude isotropy for the highest-performing readouts.

For lower recoil energies however, the number of required events diverges as the directional sensitivity and event-level detection efficiency deteriorates in all readouts. To quantify these observations with a single number, we define a *directionality threshold*: the recoil energy at which ten or more interacting events are required to exclude isotropy. Note that the choice of ten events is an arbitrary choice made to allow an easy comparison of readout — therefore the directional threshold is not to be interpreted as a hard threshold, and there is still directional sensitivity below this threshold.

With those disclaimers made, the directionality thresholds for pixel, strip, pad, wire, and planar readouts are approximately 4 keV_r, 6 keV_r, 17 keV_r, 22 keV_r, and 42 keV_r, respectively, for helium recoils. We note that to get such good performance at low recoil energies, the low mass density of the 755:5 Torr He:SF₆ gas mixture was required. For a 740:20 mixture (not shown here), the directional thresholds were approximately a factor of two higher.

Comparing Fig. 17 left and right, we see that for fluorine recoils, the directional thresholds are more than a factor of three higher than for helium recoils. Since our “signal” is maximally anisotropic here, these thresholds estimate the limit for good directional sensitivity. It should be noted however we have used a very simple measure of directionality. Directional sensitivity can be improved by performing statistical tests that use all recoil angles and their correlation with recoil energy, for example in a likelihood-based analysis. Furthermore, as noted earlier, the pixel and to some extent the strip readout, are diffusion-limited below 20 keV_r for helium, and 50 keV_r for fluorine recoil. Hence the directional-

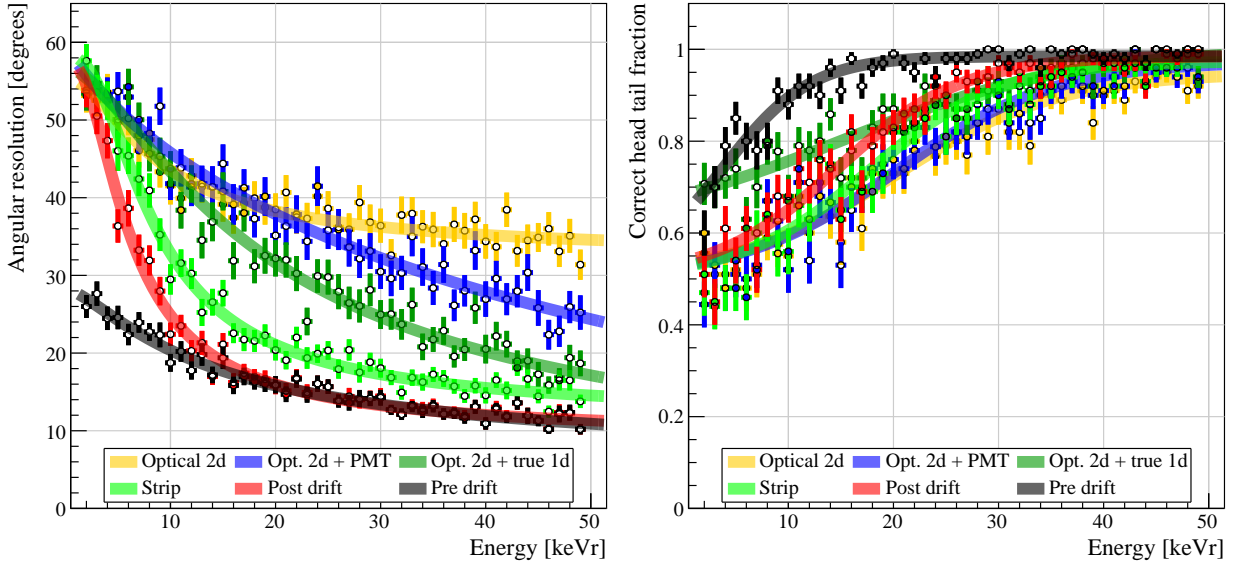


FIG. 15. Angular resolution (left) and head/tail recognition efficiency (right) for helium recoils in 755:5 Torr He:SF₆ gas, detected with various optical readout technologies, with optical loss factors and detector noise disabled in the simulation. Strip readout and maximum achievable performance before and after diffusion are also shows, for comparison. See text for discussion.




FIG. 16. Reconstructed polar angle distribution for 20 keV_r fluorine (left) and helium (right) recoils detected with a pixel readout for a hypothetical mono-directional recoil signal (red) and isotropic recoil distribution background (blue).

ity threshold and WIMP sensitivity should improve even further with a detector operating point (gas pressure and drift field) fully optimized for these high-resolution readouts. That optimization is beyond our scope, but will be an important part of the next stage, a technical design for a large directional detector.

G. Directional WIMP/neutrino discrimination and energy threshold requirements

Given that WIMP and neutrino recoils produce keV-scale energies and steeply falling energy spectra, it is clear that the directionality thresholds of different readouts will have a major effect on the WIMP and neutrino sensitivity. Figure 18 (left) shows the mean number of WIMP-helium recoil events required to exclude a back-

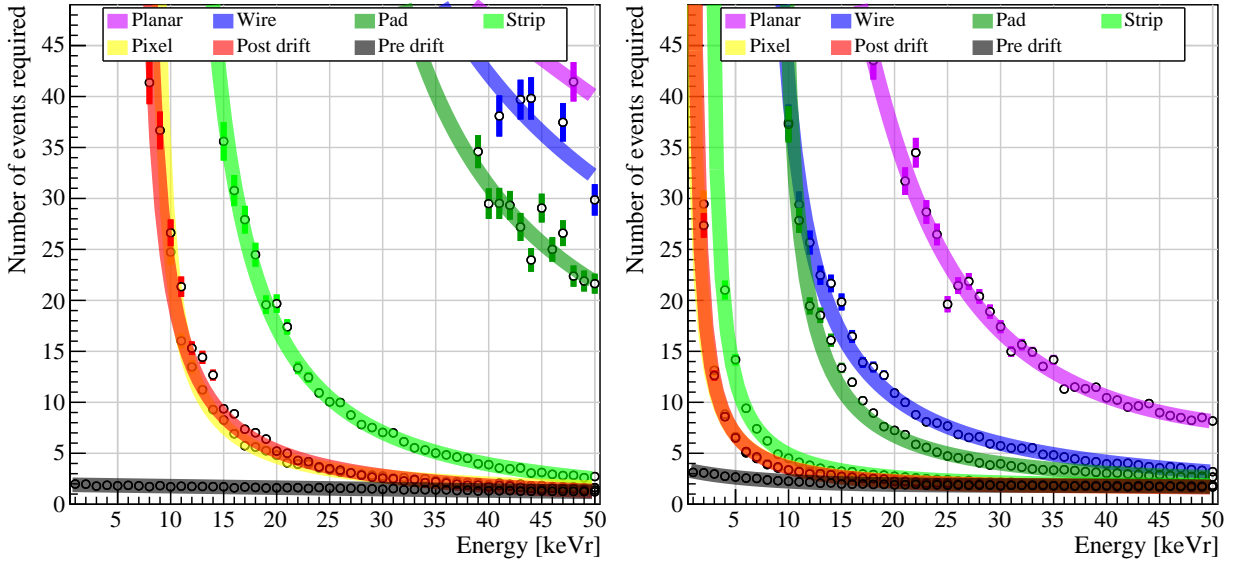


FIG. 17. Mean number of mono-directional fluorine (left) or helium (right) recoil events required to reject an isotropic recoil hypothesis at 90% CL, in the case of zero background events, for a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics. The performance that can be obtained with a perfect readout before (black) and after (red) charge diffusion during drift is also shown, for comparison.

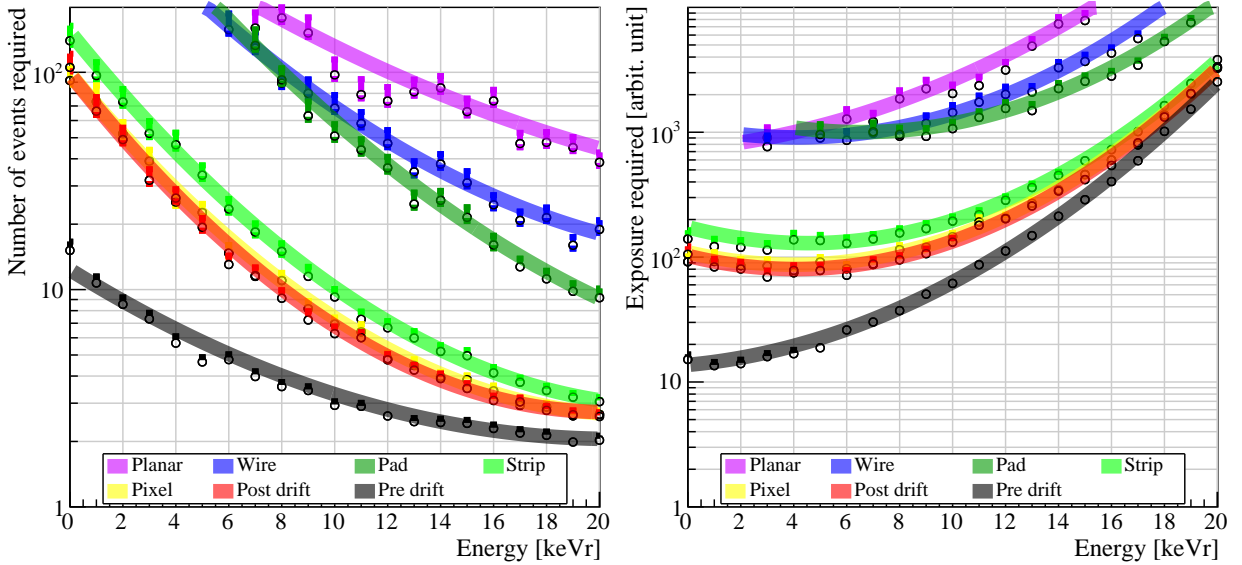


FIG. 18. Mean number of 10-GeV-WIMP-helium recoils (left) and exposure (right) required, in order to exclude isotropy in galactic coordinates at 90% CL, versus energy threshold, for a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics. See text for further details on this and following plots, including why some data points are not drawn.

ground recoil distribution that is isotropic in galactic coordinates, versus recoil energy threshold. The number of events required for exclusion falls with energy, because the angular resolution and head/tail efficiency both improve with higher energy, so that the WIMP recoil distri-

butions become more asymmetric. Because the WIMP recoil rate falls with energy, however, the best strategy for an experiment is not simply to raise the energy threshold. To get a feel for the energy threshold that minimizes the exposure need to exclude isotropy, we also show the

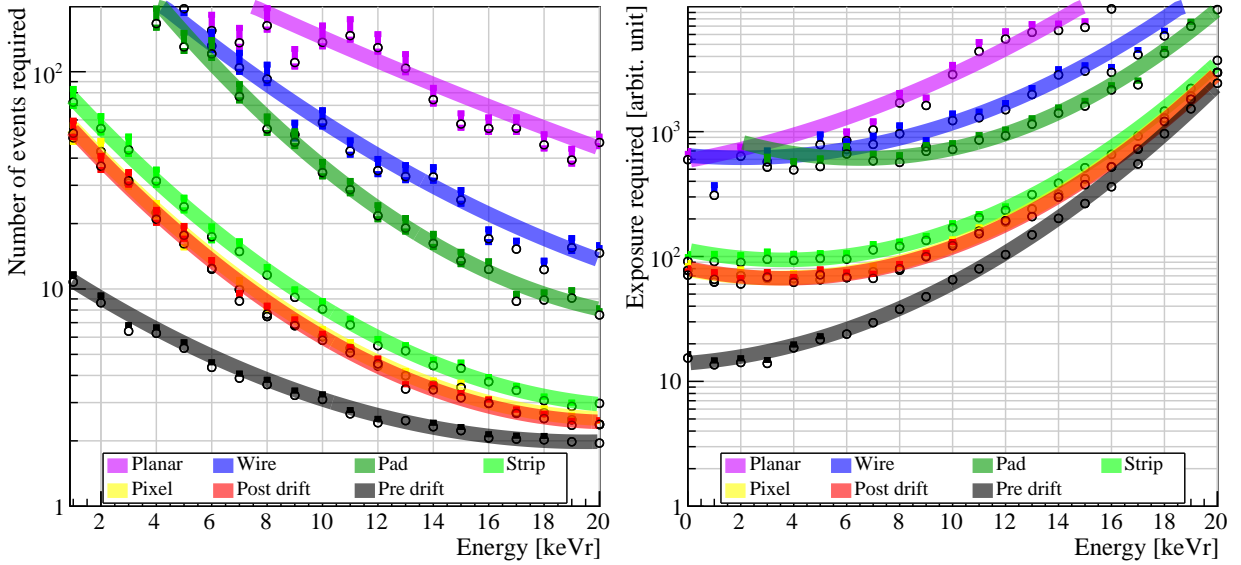


FIG. 19. Mean number of 10-GeV-WIMP-helium recoils (left) and exposure (right) required to exclude a neutrino background hypothesis at 90% CL, versus energy threshold, in the case of no detected neutrino background events, for a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics.




FIG. 20. Mean number of 100-GeV-WIMP-fluorine recoils (left) and exposure (right) required, in order to exclude isotropy in galactic coordinates at 90% CL, versus energy threshold, for a 755:5 Torr gas mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics.

exposure required for exclusion versus energy threshold in Fig. 18 (right). Note that the unit we use as a proxy for exposure is actually the number of recoil events before the energy threshold - so while these exposure units can be used to compare readouts, they are arbitrary in the sense that they cannot be used to compare scenarios with different target nuclei or different WIMP masses,

i.e. different figures in this article.

An important result in Fig. 18 (left) is that only three to four recoils above 20 keV_r are sufficient to exclude isotropy for the two highest-performing readouts, pixels and strips, respectively. This appears consistent with previous idealized studies such as Ref. [43], which found that 5 to 9 events are required for exclusion of isotropy

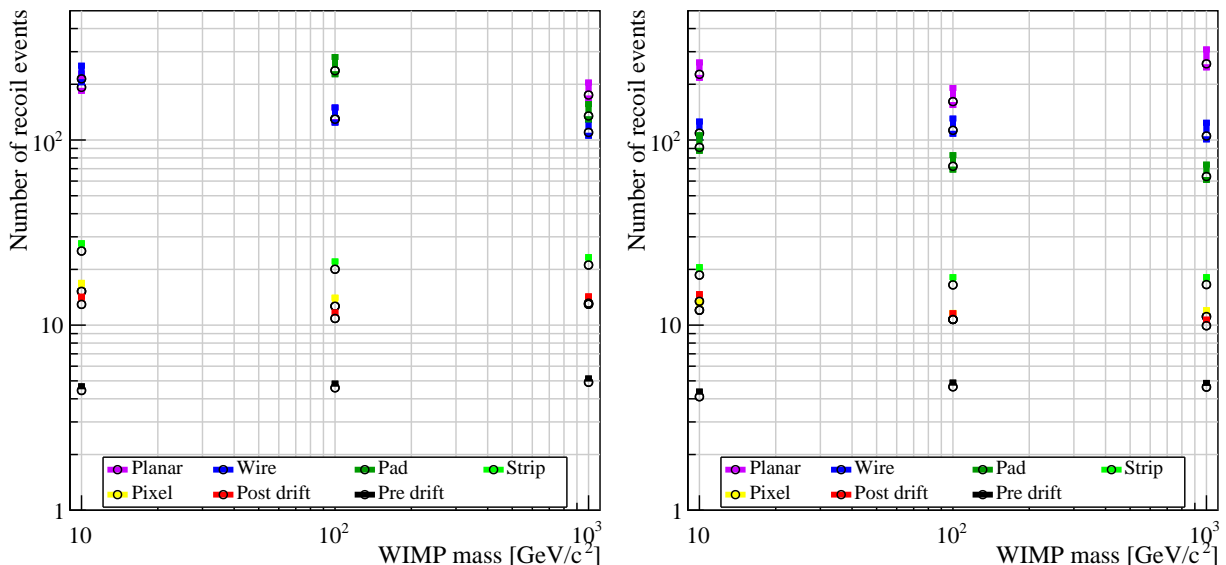


FIG. 21. Mean number of WIMP-helium recoil events with energy greater than 6 keV_r required to exclude an isotropic (left) and neutrino (right) background hypothesis at 90% CL, versus WIMP mass, in the case of no detected background events. The target gas simulated is a 755:5 Torr mixture of He:SF₆. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics.

at 90% CL in 90% of the cases, which is a stricter statistical requirement than ours. Fig. 18 (right) shows that the exposure required to exclude of isotropy plateaus below 6 keV_r for all readouts, but the required exposure is one order of magnitude lower for pixels and strips than the other readout technologies. This leads to two major conclusions: First, directional detectors with strip or pixel readout have about one order of magnitude higher directional WIMP sensitivity per unit of exposure than other readouts studies here. Second, to maximize directional sensitivity to 10-GeV c^{-2} WIMPs, we should aim to achieve an energy threshold of 6 keV_r or lower.

Before moving on, we give more detail on exactly how Fig. 18 is produced, as the same procedure is also used to produce Figs. 19 through 22. We perform 400 pseudo-experiments for each readout and energy threshold combination. For each experiment, we record the number of events required to achieve exclusion at 90% CL, using only the observed recoil direction $\cos\theta$ as a discriminant. The number of events required to achieve exclusion at 90% CL is a highly asymmetric distribution, with a significantly larger high-side than low-side tail. In Fig. 18 we conservatively plot the mean, which tends to be significantly larger than the median. The error bars shown are the asymmetric statistical uncertainties on that mean. For cases with limited directionality and low energy thresholds, sometimes a very large number of events is required for exclusion, requiring impractically long computations. To reduce computational time, we skip these scenarios: if any of the 400 pseudoexperiments require more than 5000 detected events to achieve

exclusion, we abort that particular simulation and do not draw a data point. This is the reason, for example, why there are no data points for the less directional readouts in Fig. 18 (right) for the lowest energies. In this simulation, we orient the detector so that the WIMP wind is aligned with the TPC drift axis (z). This is equivalent to putting the detector on an equatorial mount, and continuously rotating the drift axis towards the expected WIMP wind direction. For the 3d pixel and strip readouts, this subtlety is unimportant as these readouts have close to isotropic performance. However this is a choice that will artificially improve the sensitivity of less segmented readouts with anisotropic performance or lower dimensionality, though only mildly.

Figure 19 shows the analogous analysis, but now with ^8B neutrino recoils as the background hypothesis. We maintain the recoil angle distribution in galactic coordinates as the WIMP/neutrino discriminant. Compared to the isotropic background case, slightly fewer events are required to exclude the neutrino background hypothesis. Improved discrimination is possible by including the energy spectrum, event time, and signal normalization as discriminants, to be included in Sec. IV I.

Figure 20 shows the case of 100 GeV c^{-2} WIMPs. For the readouts with the best performance, five or fewer fluorine recoils at high energy are sufficient to exclude an isotropic background hypothesis at 90% CL. Note that in this case the plateau in the exposure required for exclusion occurs around 30 keV_r , while for the other scenarios studies here this plateau occurred around 6 keV_r . This means that for a directional detector optimized specifi-

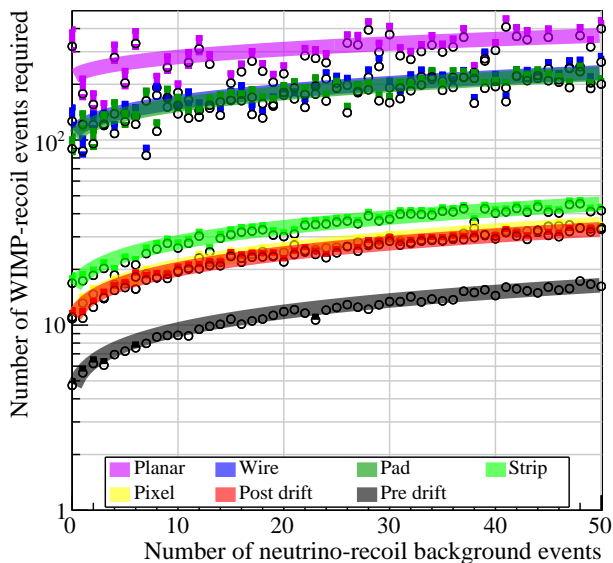


FIG. 22. Mean number of $10 \text{ GeV } c^{-2}$ WIMP helium-recoil events with recoil energy greater than 6 keV_r required to exclude a neutrino background hypothesis at 90% CL, versus number of neutrino background events in the detected event sample. The target gas simulated is a 755:5 Torr mixture of $\text{He}:\text{SF}_6$. Each color denotes a different readout as shown in the legend. Data points show the mean number of recoils required, based on pseudo-experiments, with errors indicating the uncertainty due to finite simulation statistics. The lines show χ^2 -minimizations of the functional form $a + b\sqrt{n}$ to the data points, with a and b being floating fit parameters and n being the number of neutrino background events.

cally for high-mass WIMPs, a higher energy threshold is acceptable.

In summary, we have found that a 3d readout with good vector tracking capability, *e.g.* pixel or strip readout, can rule out an isotropic background or neutrinos with as few as three to ten recoils above an energy threshold of 20 to 50 keV_r , depending on the exact scenario. The required exposure for exclusion is typically minimized, however, with an energy threshold of about 6 keV_r . Figure 21 provides a compact summary of the number of helium recoils required to exclude the isotropic or neutrino background hypotheses for 10, 100, and 1000 $\text{GeV } c^{-2}$ WIMPs, using this 6 keV_r energy threshold. In that case typically between ten and twenty recoil events are required for exclusion via directionality only. The findings here are consistent with past theoretical works, summarized in Ref. [14].

Finally, Fig. 22 shows how the number of WIMP events required to exclude the neutrino-only hypothesis increases with neutrino background events present in the detected event sample. For a $10 \text{ GeV } c^{-2}$ WIMP, the number WIMP-recoils required (in this case for an energy threshold of 6 keV_r) grows only slowly with the number of neutrino events, illustrating the power of directionality in discriminating against even large a neutrino background.

H. Electron background rejection

Electron backgrounds (*i.e.* gamma-recoils) are a key issue for CYGNUS in that they will effectively determine the energy threshold. In our background simulations of a strawman CYGNUS-1000 design (Section V) we find that the electron energy spectrum is dominated by Compton scattering, with a flat energy spectrum in the $\mathcal{O}(\text{keV})$ region where we expect the WIMP recoil signal. Our design goal, which we show can be met for the lowest-background readouts, is an electron background rate of $10^4 \text{ keV}_{ee}^{-1} \text{ year}^{-1}$. We will quantify the ability to reject electrons via the electron rejection factor, R , defined as

$$R = N_{all}/N_{surv}, \quad (14)$$

where N_{all} is the number of electron background detected in a given energy range, and N_{surv} is the subset of those events that survive an electron-veto algorithm that retains 50% of nuclear recoils. R is thus the factor by which we expect to suppress detected electron background via offline selections. We will see that the electron rejection factor in CYGNUS is high and rises exponentially with energy. Assuming a six-year exposure and a flat electron background energy spectrum at a rate of $10^4 \text{ keV}_{ee}^{-1} \text{ year}^{-1}$, if we assume a hard energy threshold at the center of the first 1- keV_{ee} energy bin where $R > 6 \times 10^4$, then (due to the steep rise of R with energy) CYGNUS will essentially be free of electron backgrounds. This is somewhat conservative; the electron rejection factor also rises strongly if we accept lower nuclear recoil efficiency, so that a real experiment can probably remain background free at even lower energies by accepting some efficiency loss. To quantify this lower-energy performance, we also define the *electron rejection turn-on* as the lowest energy where $>90\%$ of electrons can be rejected while retaining half of the nuclear recoils. This corresponds to the energy where $R = 10$.

The charge density distributions of electronic and nuclear tracks in gas are very different in both shape and scale, as demonstrated by Figs. 6 and 7. This suggests that good electron discrimination should be possible. As a first step, the linear charge density is a simple, yet powerful discriminant than can be obtained from the charge distributions. Lines of constant charge density correspond to lines of constant slopes through the origin in Fig. 8 (right). Furthermore, we can expect the discrimination between electrons and fluorine recoils to be better than between electrons and helium recoils. Fluorine recoils have higher specific ionization and therefore differ more from electron recoils than helium recoils do. In Fig. 23, we show the electron rejection factor R that can be achieved using only this simplest linear charge density. Before diffusion (data points with blue error bars), we obtain significant discrimination down to the lowest energy studied (1 keV_{ee}) for fluorine. The electron rejection factor rises quickly with energy, and exceeds 6×10^4




FIG. 23. Electron rejection factors for fluorine recoils (left) and helium recoils (right) in 755:5 Torr of He:SF₆. Only the simplest observable, the fitted-track length at a given ionization energy, is used to discriminate between electrons and nuclear recoils. Within each energy bin we apply a variable minimum track length that retains 50% of the nuclear recoils. Error bars show combined uncertainties due to finite simulated electron and nuclear recoil statistics. The electron rejection rises exponentially with energy. For the higher energy bins, where no data point are shown, *all* simulated electron events were rejected, and the electron rejection factor exceeds 10^6 .

at approximately 3.5 and 7.5 keV_{ee} for fluorine and helium recoils respectively. We also see that the electron rejection turn-on occurs at order 1 keV_{ee} and 2.5 keV_{ee} for fluorine and helium, respectively.

The data points with red errors bars in Fig. 23 show the result after diffusion corresponding to 25 cm of drift. We see that the discrimination power remains high even after this amount of diffusion, and that the electron-rejection turn-on moves up to about 4 keV_{ee}. It should be noted that the diffusion simulated here is *more* than the average for a 50 cm long maximum drift length (due to the $\sqrt{\text{distance}}$ scaling), and that for recoils with short drift length the electron rejection will be closer to the pre-drift case. This result is remarkable because conventional direct detection experiments based on ionization signals typically run out of discrimination power at these low energies.

An important limitation of this study worth remarking upon can be seen in the absence of pre-diffusion data points above 4 keV_{ee} (fluorine) and 9 keV_{ee} (helium). This is an artificial effect due to the finite Monte Carlo samples; in fact *all* simulated electron events were rejected for these bins. In reality we expect the exponential increase of the electron discrimination factors to continue. It is clear from Fig. 8 (right) that the discrimination should continue to rapidly improve with energy, but studying this further will require substantial CPU time. For similar reasons, we have not studied the readout dependence of electron discrimination, unlike previous results in this section. Nevertheless, from results such as

Fig. 13, we can expect that the most segmented readouts (pixel and strip) will have rejection factors quite close to what we report here.

While already promising, we expect these electron rejection factors to only improve from here. Preliminary experimental and simulation work by some of the authors has already achieved improved discrimination with more sophisticated shape analysis of the detected recoils [317, 318]. For example, deep learning neural net studies at the University of Hawaii achieved an impressive electron rejection factor of $R = 10^6$ at 3 keV_{ee} for 740:20 He:SF₆ mixtures and helium recoils with 100 μm of diffusion included. We expect this to improve by one additional order of magnitude in the lower-density 755:55 He:SF₆ mixture. In these preliminary results, the neural networks are not fully optimized, but already we see background rejection better than chance at 1 keV_{ee}. With additional tuning, we expect to see good discrimination at energies less than 1 keV_{ee} [319]. Putting all this together, we expect to achieve R of order 10^7 at 3 keV_{ee} for helium recoils, which corresponds to 7.5 keV_r, for a 50% helium recoil efficiency. Based on this estimate, we should remain background free at substantially lower energies, but more work is required to establish the exact threshold electron-background-free operation. In our reach curves, we draw 8 keV_r as the highest possible threshold line. Because we expect high electron rejection at substantially lower recoil energies, and some electron rejection even at 1 keV_r, we also show a number of lower energy thresholds.

Future work will therefore be devoted to experimental verification of the electron rejection factors reported here, followed by the implementation of discrimination methods using the detailed charge density distribution.

I. WIMP sensitivity below the neutrino floor

Following our discussion of the directional performance of each readout technology, we now examine how this performance translates into sensitivity to WIMP cross sections. To do this we build a signal and background model for the WIMP and neutrino recoil distributions based on the theoretical input described in Sec. II. Then, all energy-dependent angular resolutions and head/tail efficiencies from Figs. 13 and 14 are applied, as well as the event-level energy resolutions and efficiencies from Figs. 12 and 11. For the threshold, a hard cut on recoil energy is imposed at a range of values up to 8 keV_r , but the Gaussian energy resolution is applied to the underlying dR/dE_r before this cut is made. A harder cut of 0.25 keV_r is applied first however, to address the uncertainty in determining the readout performance for energies below the single electron level.

Then we use our signal and background models to calculate median 90% CL exclusion limits using the standard profile likelihood ratio test (see *e.g.* Ref. [320] for the statistical formalism, but the WIMP+neutrino analysis here closely follows Ref. [36]). Each neutrino background flux normalization is accounted for as a nuisance parameter using a Gaussian parameterization. No other uncertainties or backgrounds are accounted for, so our hard threshold effectively enforces the assumption that perfect electron discrimination can be achieved above it.

The resulting limits on the SI and SD parameter spaces are shown in Fig. 24 for a 1000 m^3 experiment and for each readout. We also include several additional theoretical limits with improved directional sensitivity to highlight the room for improvement. The limits labeled ‘predrift’ are similar to those introduced earlier. We assume the energy resolution and efficiency of the pixel readout, but the angular resolution and head/tail efficiency is as though we were able to measure the track topology before diffusion (this emphasizes the impact of diffusion on the sensitivity).

We also introduce in this plot a limiting case labeled ‘ideal’, which takes the predrift sensitivity and simply sets the angular resolution to zero and the head/tail recognition to 100%. This is a theoretical ideal, not possible in practice, but our fully simulated readouts are promisingly close already. This limit corresponds to almost perfect discrimination between WIMPs and neutrinos, a feat that should be possible in principle based on their angular distributions (see Fig. 5). We ignore the ‘postdrift’ curve here (which was shown in previous results), since its corresponding limits are nearly identical to the pixel readout. As demonstrated by the fact that most of the curves converge at high masses, the di-

rectional information is playing a subdominant role in setting limits when compared with statistics alone. Interestingly the poor energy resolution of several readouts does enable them to set limits at lower masses than those with superior resolution, but at the cost of poorer sensitivity over larger masses due to the lack of ability to perform signal characterization. This demonstrates that for a CYGNUS-1000 experiment the sensitivity is much more sensitive to the number of events than the overlap between the signal and background distributions. Therefore the hard threshold is something that should be considered carefully.

Fig. 25 shows in more detail how the limits of two separate cases, predrift (top) and strip (bottom), depend on the imposed electron discrimination threshold. Lowering the threshold down to 3–4 keV, can claim a generous portion of unconstrained low-mass WIMP parameter space. Attempting to lower this threshold even a small amount is therefore highly motivated, even though the directionality at low energies is poorer. Unfortunately, for 1000 m^3 of He:SF₆ at atmospheric pressure, the low target mass only allows the experiment to just reach the neutrino floor, although there is an expected neutrino event rate greater than one. To see this explicitly we show 1-neutrino event lines rather than the neutrino floor (which typically corresponds to $\mathcal{O}(100)$ expected neutrino events). Limits intersecting these boundaries (for the relevant target) observe more than one neutrino event. The current electron discrimination threshold of 8 keV_r in fluorine is extremely close to the tail of the solar CE ν NS recoil spectrum; anything slightly larger would not observe the background and not cut into the floor. In fact the improvement brought about by directionality, comparing the top and bottom row, is relatively minor in comparison to the substantial benefit at low masses when lowering the threshold. In other words, the 1000 m^3 exposure is still statistics limited. For larger experiments this conclusion will change, as we show next.

In Fig. 26 we show two discovery limits for individual WIMP masses as a function of the number of expected neutrino background events. The readout assumed in each case is the same (pixel) but we compare several analyses which use all, or a subset of the available information. The region outlined by dotted lines assumes only the recoil event time information is used, so that only event counts and the annual modulation enable sensitivity. The dashed lines assume that recoil energy information is also used. The dot-dashed lines assume that the directional information is used with directional performance curves simulated for the pixel readout. The lowest, solid lines correspond to the pixel readout if it were able to measure the track before diffusion, so as to highlight how much of the sensitivity is limited by diffusion.

Until now we have considered a CYGNUS-1000 benchmark volume. Comparing the limits at this point in Fig. 26 we can see that, as hinted by the previous figures, we are still in a regime where the sensitivity gains

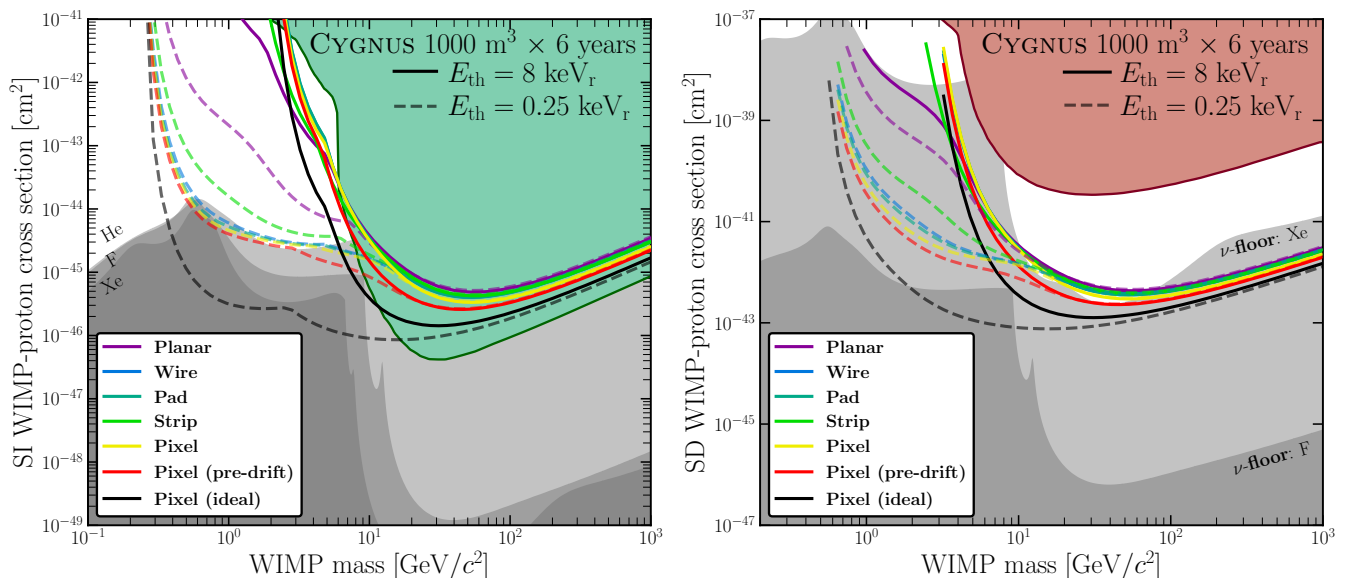


FIG. 24. Expected SI-nucleon (left) and SD-proton (right) 90% CL exclusion limits for CYGNUS-1000 using He:SF₆ at 755:5 Torr over an exposure of six years (~ 1 ton-year exposure). The solid lines all correspond to limits assuming an 8 keV_r threshold whereas the dashed are for 0.25 keV_r. We compare all readouts that have been simulated in this section, as well as two theoretical limits labeled “pre-drift” and “ideal” which are included to highlight the role of the directional performance in limiting the sensitivity. The two additional examples both assume the charge detection efficiency and resolution of the pixel readout, but have their directional performance modified. The “pre-drift” limits correspond to cases already shown in previous results, *i.e.* the angular resolution and head/tail efficiency of the track as if it were measured before diffusion. The “ideal” limit corresponds to an angular resolution of 0° and a head/tail efficiency of 100%. The combination of all existing SI and SD exclusion limits are shown as a green and red filled region respectively. We also show the neutrino floors for helium, fluorine and xenon targets in grey.

more from the improved threshold via statistics, than from directionality. For larger experiments however the directional sensitivity adds a considerable benefit. In an experiment without directional information that also observes a sizable neutrino background, the sensitivity scales with detector volume at a rate much slower than the standard Poissonian background subtraction of $V^{-1/2}$. However, directional limits scale much faster. This information can improve limits by almost an order of magnitude when the directionality is good (see the pre-drift solid lines). This is also the case for the more realistic directional performance of the full readout accounting for diffusion (dot-dashed lines), however the conclusion again seems to be that diffusion is the dominating factor in restricting better sensitivity. We reiterate that the only difference between the pre-drift and post-drift lines is an improved angular resolution and head/tail recognition. This figure makes clear how much our sensitivity could stand to improve.

J. Summary of readout performance and conclusion on cost-optimal readout choice

Table III summarizes the main detector performance parameters resulting from the detector and readout simulation. We define each of these parameters in turn.

We define the *event detection threshold* as the recoil energy at which the event-level detection efficiency exceeds 50%. This threshold is mainly determined by the avalanche gain and the readout noise floor, and the noise floor is in turn determined by the capacitance of the readout elements. Even with the modest gain of 9000 assumed in this study, the readouts with the lowest capacitance per readout element (pixels, pads, wires) have high efficiency at 1 keV_r, the lowest recoil energy simulated. The strip and planar readouts have larger detector elements with larger capacitance and noise, and require higher thresholds to compensate. The threshold are slightly lower for fluorine than helium recoils, because the former have higher charge density.

The *directionality threshold*, defined as the recoil energy where ten events are sufficient to identify a maximally directional source, is determined by the recoil track length, energy threshold, readout segmentation as well as the nuclear straggling and diffusion of drift charge. We note that without diffusion or readout effects simulated, this directionality threshold is below 1 keV_r. After diffusion but before detection, the directionality threshold is about 15 keV_r for fluorine and 4 keV_r for helium, mainly because the helium recoils are longer, which helps overcome the diffusion. Comparing the different readouts, we see that the pixel readout directional threshold is nearly identical to what is seen after diffusion. This means that

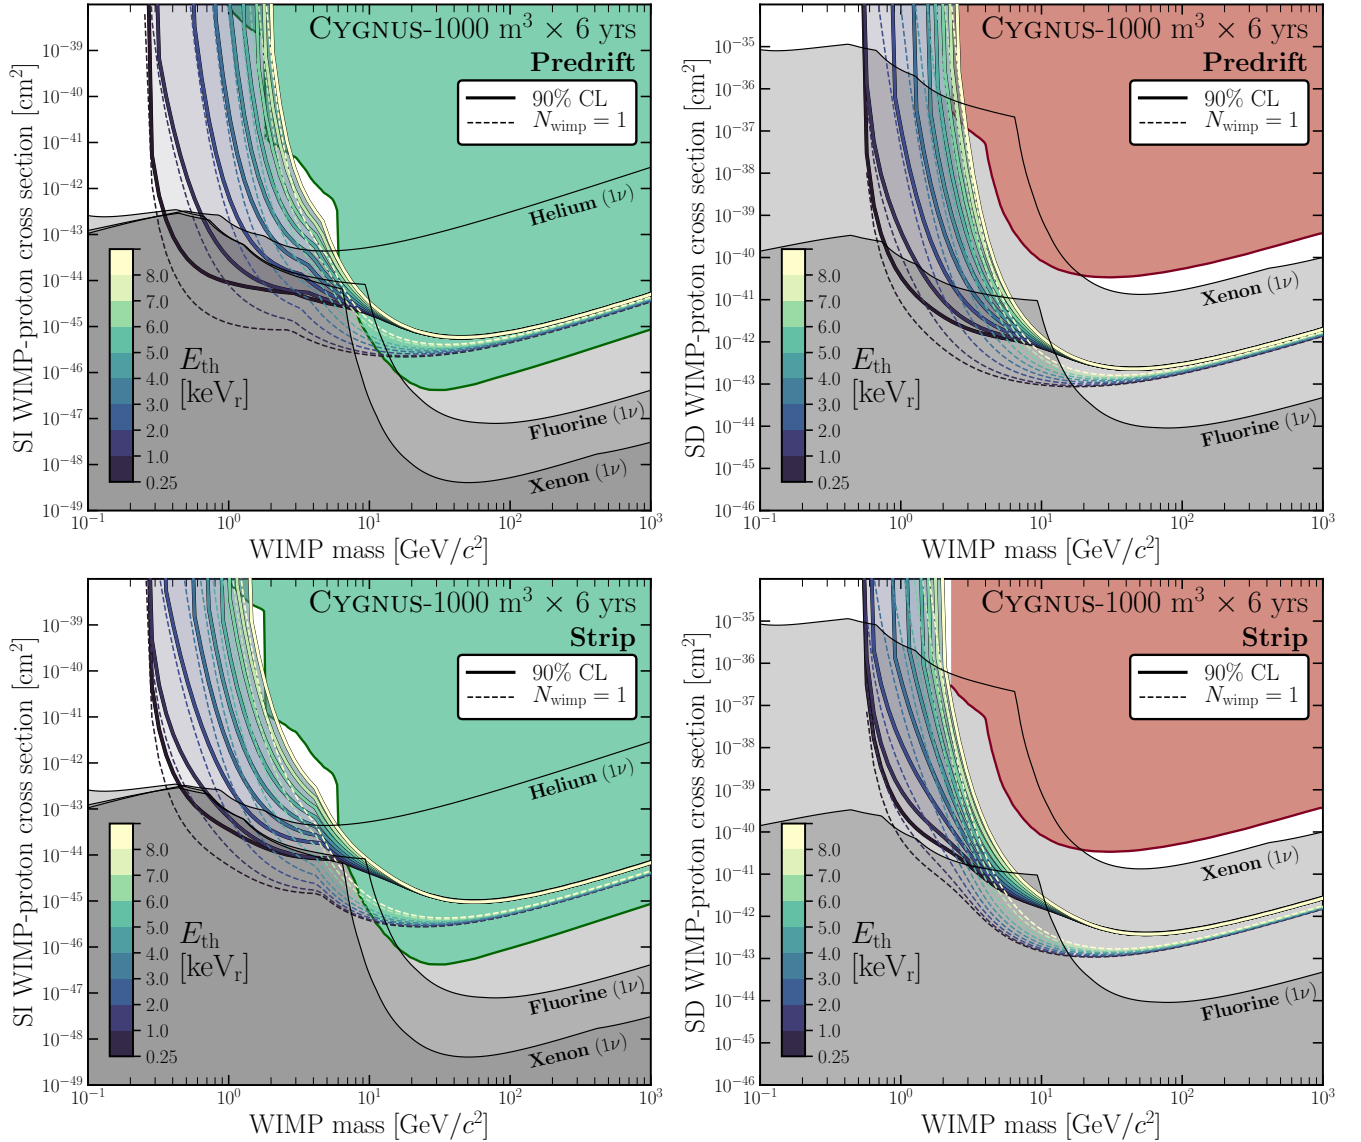


FIG. 25. Expected SI-nucleon (left column) and SD-proton (right column) 90% CL exclusion limits for CYGNUS-1000 as a function of the hard recoil energy threshold. We focus on the analyses corresponding to the predrift directional sensitivity (top row) and the strip readout (bottom row). We vary the threshold again from 0.25 keV_r (single electron) to 8 keV_r (electron discrimination). In contrast to the previous figure we now show the 1-neutrino event lines for helium, fluorine and xenon targets in grey, as opposed to the neutrino floor. This allows us to see that the 8 keV_r threshold only just observes a non-zero neutrino background—highlighting the importance of improving electron discrimination.

the pixel readout essentially extracts all relevant information from the diffused ionization distribution. The segmentation of this readout is thus sufficient, and finer segmentation would not improve the directional performance further, unless we also achieve lower diffusion. This may be possible with NID at larger electric field strengths, which should be explored. Note that the performance ordering of the detectors is different for directionality threshold than for energy threshold, especially for strip readout. The reason is that strip readout has high capacitance per readout element, which raises the energy threshold, but excellent segmentation once x/y

strip coincidence is utilized, which lowers the directionality threshold.

We define the *electron rejection turn-on* threshold as the detected energy at which at least 90% of electron background events can be rejected while maintaining a 50% nuclear recoil detection efficiency. This threshold is less than 1 keV_{ee} (fluorine) and 2.5 keV_{ee} (helium) before diffusion, and approximately 4 keV_{ee} after diffusion for the 755:5 Torr He:SF₆ gas mixture where we have simulated electrons. At higher energies, the electron rejection power improves exponentially. We take these electron rejection turn-on thresholds as upper limits, as we expect

Charge readout	predrift	postdrift	pixels	strips	pads	wires	planar
Event detection threshold (F) [keV _r]	n/a	n/a	< 1	1.5	< 1	< 1	4
Event detection threshold (He) [keV _r]	n/a	n/a	< 1	3	< 1	< 1	7
Directionality threshold (F) [keV _r]	< 1	< 1	14	25	> 50	> 50	> 50
Directionality threshold (He) [keV _r]	< 1	< 1	4	6	17	22	42
Electron rejection turn-on (F)[keV _{ee}]	< 1	< 4	-	-	-	-	-
Electron rejection turn-on (He)[keV _{ee}]	< 2.5	< 4	-	-	-	-	-
Exp. penalty, exclude isotropy, 10-1000 GeV c^{-2} WIMPs (F)	4.7	12	14	22	186	151	166
Exp. penalty, exclude neutrinos, 10-1000 GeV c^{-2} WIMPs (He)	4.5	11	11	17	76	109	215
Average relative exposure penalty factor	n/a	n/a	1	1.6	10	10	15
Approx. cost per unit readout area [US \$/m ²]	n/a	n/a	400k	22.5k	5k	5k	0.050k
Total readout cost (US \$)	n/a	n/a	800M	71M	105M	104M	2M
Total volume cost (US \$)	n/a	n/a	25M	39M	262M	261M	382M
total detector cost, constant WIMP sensitivity (US \$)	n/a	n/a	825M	111M	367M	365M	384M
total detector cost, 1000 m ³ volume (US \$)	n/a	n/a	825M	70M	35M	35M	25M

TABLE III. Summary of main performance parameters and estimated detector cost at equal directional sensitivity, for the simulated TPC charge readout technologies. Results assume a 755:5 He:SF₆ gas mixture at atmospheric pressure, room temperature operation, a gain of 9000, charge diffusion of 78.6 $\mu\text{m}/\sqrt{\text{cm}}$, and maximum drift length of 50 cm.

improved electron rejection, and consequently even lower energy thresholds, are possible with more sophisticated selections. Preliminary work on this is highly encouraging, see Section IV H. A full analysis of electron rejection capabilities of different readouts, with more sophisticated discriminants, is important for future experiments, and should be followed up in future work.

We have seen that ruling out an isotropic background (in galactic coordinates) or neutrino background hypothesis using only the directional distribution of detected events, is possible at 90% CL with as few as three recoils with energy E greater than 20 keV_r, or of order ten recoils with $E > 6$ keV_r, for the highest performance readouts (pixels and strips), for WIMP masses of 10, 100, and 1000 GeV c^{-2} . To fairly compare readouts with different energy thresholds and directional performance, we report in Table III the *exposure penalty*, defined as the number of WIMP interactions with $E > 6$ keV_r that must take place, in order to detect sufficient events to rule out the background hypothesis. This would be the first step towards confirming the cosmological origin of a tentative WIMP dark matter signal. In order to reach identical WIMP sensitivity in this context, detectors with different readouts need to have an exposure proportional to their respective exposure penalty. We find that this penalty does not depend very strongly on whether the background is isotropic or due to neutrinos. The penalty is also roughly the same for 10, 100, and 1000 GeV c^{-2} WIMPs.

We see that it instead increases with the directionality threshold, which makes sense, given the falling energy spectrum of WIMP recoils. Quantitatively, we find that a pixel readout needs the lowest exposure to rule out a background via directional measurements. Hence we average the penalty factors for the six scenarios (three WIMP masses and two background hypotheses for each), and normalize them to the pixel readout, to arrive at the *average relative exposure penalty* factor. We find that

strip, pad, wire, and planar readouts require on average 1.6, 10, 10, and 15 times higher exposures, respectively, to reach the same directional WIMP sensitivity as the pixel readout.

Next, we estimate cost. Given that pixel readout has the lowest exposure penalty, we take a detector with a volume of 1000 m³ and with pixel charge readout as a baseline scenario. For each of the other readout technologies, we set the volume equal to the product of 1000 m³ and the average relative exposure penalty factor for that particular technology. This way, we are comparing designs with identical expected directional sensitivity. For all designs, we assume a drift length of 50 cm, so that there are two readout planes per 1 m³ unit cell of the detector.

We consider two cost components: The first component is *total readout cost*, which we take to be the product of three factors: 2000 m², the *readout cost per square meter*, and the average relative exposure penalty factor. The two latter factors both depend on readout choice, and vary substantially. Interestingly, the two effects cancel to some degree, so that the variation in total readout cost is not as large as one might have expected. Strip readout and planar readout have the lowest total readout cost. For the strip readout, the cost includes \$12,500/m² for strip micromegas or similar charge-readout and \$10,000/m² for readout ASICs (10,000 readout channels per m² at a cost of \$1/channel).

The second cost component considered is cost per unit volume, which is independent of readout type and assumed to be US \$ 5000 (\$5k). This cost term accounts for all volume-dependent cost other than the charge readout itself: the vacuum vessel, field cages, shielding, gas, the charge avalanche device, and downstream data acquisition. For the pixel readout, we base the cost on a volume of 1000 m³, resulting in a total volume cost of US \$ 5 million (\$5M). For the other readouts, we increase the volume (and thus the volume cost) by the relative

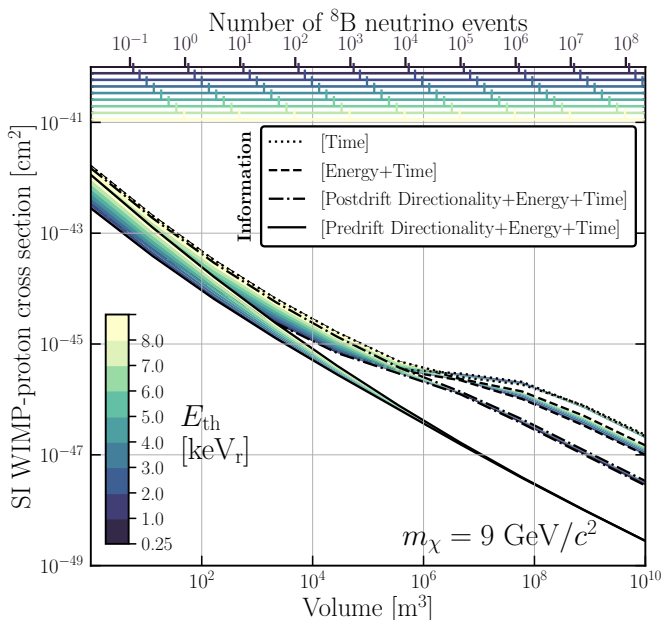


FIG. 26. Minimum excludable SI WIMP-nucleon cross section at 90% CL, for a $9 \text{ GeV}c^{-2}$ WIMP, as a function of detector volume. The colored bands correspond to energy thresholds in the same range used in the previous two figures. The number of ^8B neutrino events observed for each threshold as a function of volume is displayed in the upper colored horizontal axes. The various regions correspond to the performance of the pixel readout when different levels of information are used in setting the limits. The four sets of limits display various levels of information used from the least powerful to the most powerful. The solid lines correspond to the theoretical optimum for a gas detector whereas the dotted lines use only the minimum amount of available information (event times).

exposure penalty factor. The resulting estimated total volume cost is quite high for the less directional readout choices, such as pad, wire and especially planar readout. While cost-effective per unit readout area, these readout choices would require expensive and large detector volumes to reach the desired directional sensitivity.

Strip readout emerges as the technology with the lowest cost at equal directional sensitivity, or equivalently, as the technology with the highest directional sensitivity per unit cost. The latter is shown graphically in Fig. 27, normalized to pixel readout. A strip readout TPC with a fiducial volume of about 1000 m^3 , would have total cost of about \$70M, with about two thirds of that cost being associated with the specific charge readout. This total cost is similar to the cost of a typical HEP detector, but much less than the cost scale of a particle collider. This seems like a reasonable scientific investment compared to others in the field, given that it would allow us to make unique experimental progress on one of the most significant unsolved problems in physics (dark matter). Such a detector could rule out isotropy of a detected recoil distribution from 10 to 1000 $\text{GeV} c^{-2}$ WIMPs at 90% CL with 22 WIMP-helium-recoils above 6 keV_r , or with only

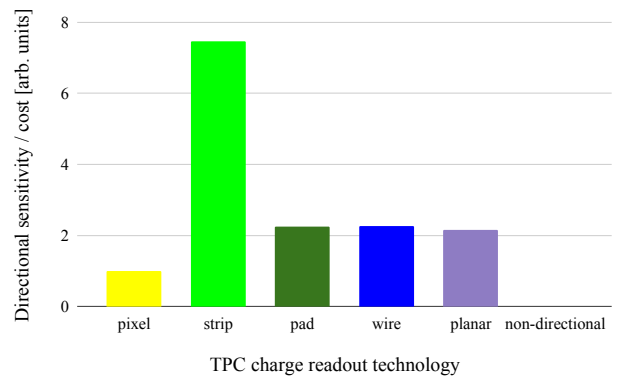


FIG. 27. Estimated directional sensitivity per unit cost for different TPC readout technologies. Pixel readout, which has the highest performance, is taken as the reference, with directional sensitivity per unit cost equal to unity by definition. We find that strip readout provides the optimal tradeoff between cost and performance. Note that non-directional nuclear recoil detectors score zero on this performance metric.

3-4 WIMP-helium-recoils above 20 keV_r .

A detector based on pixels would be an order of magnitude more expensive, but have better performance. With the lowest exposure penalty, it would require the smallest detector for a given physics reach, but the cost is currently dominated by the relatively expensive pixel ASICs. The number of pixel chips, and hence cost of pixel chip readout planes, could potentially be greatly reduced via so-called charge focusing – electrostatic focusing of the drifting ionization before detection [321]. Alternatively, the cost of pixel chips may be considerably lower with future semiconductor manufacturing processes or chip designs. In either case, pixel chip readouts, or similar ultra-high-resolution charge readouts could then become the most cost-effective readout options. With further reductions in diffusion, which would need to be demonstrated, a pixel-based design is likely to be the most competitive options at WIMP masses of order $1 \text{ GeV} c^{-2}$ and lower.

As it stands, pad, wire, and planar score similarly in terms of directional sensitivity per unit cost. These readouts are performance-limited due to relatively coarse segmentation. For planar readout, large segmentation is inherent to the design and cannot be reduced. Planar readout may still prove somewhat more competitive by utilizing a dedicated discriminant, as opposed to the general event reconstruction utilized in the comparison here. A detector based on wires or pad technologies could become substantially more cost competitive if a reduction in feature size is possible, which would lower the directional threshold and exposure penalty, while maintaining the relatively low readout cost. Here, pads have the benefit of offering full 3d directionality, while wires have the benefit of lower backgrounds.

This brings us to our next topic: we note that the

cost/performance analysis here does not directly consider the effect of electron discrimination. As discussed above, it appears likely that pixel and strip readout will have the best electron background rejection capabilities due to their high segmentation. However the wire readout has less intrinsic radioactive backgrounds, so that a smaller electron rejection factor will be required.

Taking a step back, we note in closing that a higher-gain TPC with electron drift gas could also be competitive. That option should be studied further, but is beyond the scope of this article.

V. ZERO BACKGROUND FEASIBILITY

Direct DM search experiments typically strive to control backgrounds to less than one expected event within the fiducial volume over the anticipated exposure time and energy range (so-called zero background). As we have discussed in previous sections, directional detectors—more-so than non-directional detectors—are in principle able to tolerate a non-zero level of nuclear recoil background, for example from neutron or neutrino interactions, while still being able to positively identify a DM signal. However, any backgrounds will negatively impact the sensitivity, so this level of tolerance will depend on the capabilities of the detector technology.

For CYGNUS we aim for zero electron background from radioactivity of the detector and surrounding materials, and neutron recoil background that is a factor of four smaller than the expected nuclear recoil count from neutrinos. This section describes the feasibility of this for a 1000 m^3 TPC. The simulations and preliminary result with machine learning previously discussed suggest that an electron background of order $10^4 \text{ yr}^{-1} \text{ keV}^{-1}$ could be tolerated even for energies below 10 keV. Given the quick rise of electron rejection with energy, this means that for a six year exposure, CYGNUS will be background free above the energy at which the electron rejection crosses 6×10^4 . We also saw in earlier sections that we expect approximately 40 nuclear recoils above a 1 keV_r threshold over a six year exposure, or about 4 per year (see Fig. 4).

We therefore set the following, preliminary, background design goals for CYGNUS: (1) an intrinsic electron background rate of $10^4 \text{ yr}^{-1} \text{ keV}^{-1}$, (2) one nuclear recoil above 1 keV_r from neutrons per year, and (3) any other backgrounds from other sources, notably radon, should be controlled to at least this level. This then leaves neutrino induced recoils as the only non-negligible nuclear recoil background. In what follows we must verify if these design goals are achievable. To streamline a detailed discussion of all anticipated backgrounds for CYGNUS we set a benchmark nuclear recoil threshold of 1 keV_r and calculate background rates above this value. This will inform the amount of shielding required and the level of material radiopurity needed to reach the design goals.

As discussed in Section IV A, we have been iterating on the CYGNUS target gas during the conceptual design pro-

cess. While the final results presented here are for a 755:5 Torr He:SF₆ mixture, the discussion on backgrounds below will also iterate through a number of gases, and scale backgrounds from one to the next.

To start out, we make a few simplifying assumptions. For the vessel and gas we assume that the TPC is a $10 \times 10 \times 10 \text{ m}$ cube filled with 20 Torr of SF₆. This choice allows detailed exploration of a basic, large, low-density negative ion gas configuration, but one for which the data can reasonably be extrapolated to estimate backgrounds in other gas mixtures and pressures, including those used elsewhere in this paper. We assume a single average thickness for the vessel wall, rather than individually simulating the support structures that would be needed in a real vessel. Then, for the readout we assume it is configured as in the TPC design of Fig. 1, with a 50 cm drift distance (following the previous section). The total readout area is therefore 2000 m^2 . The laboratory environment is based on the salt rock found at Boulby, UK. This site has certain advantages such as a low rock gamma flux, but we leave the discussion of the various site considerations to Sec. VI.

Sections VA–VD detail all neutron, gamma, radon-related and cosmogenic backgrounds respectively. We then describe in Sec. VE how the results can be extrapolated to the atmospheric pressure He:SF₆ gas mixture currently envisaged for CYGNUS. We give our final background summary and recommendations in Sec. VF.

A. Neutron backgrounds

Neutrons produce nuclear recoil events in the same range of energies as those expected from WIMPs. The neutron background is therefore one of the most problematic for all direct detection experiments. However, mitigating against it in a low-density gas TPC is distinct from the methods used in solid or liquid-based detectors. The low density of gases means that neutrons are less likely to undergo double or multiple scattering, reducing the potential to veto on this basis. Estimating neutron backgrounds by extrapolation from existing background simulations for massive xenon or bolometric detectors [322] for example is therefore not appropriate. Instead, we must perform a dedicated Monte Carlo simulation for CYGNUS, for which we use Geant4 [323]. We simulate the neutron background from the rock, shielding, vessel, and internal readout components using neutron energy spectra obtained from SOURCES [324]. The spectra include neutrons originating from the ²³⁸U and ²³²Th decay chains due to spontaneous fission and (α , n) reactions. For neutrons due to cosmic ray muons interacting with the detector or surrounding materials, we use the simulation code MUSUN [325]. We now detail the rate calculation for each of these components in turn.

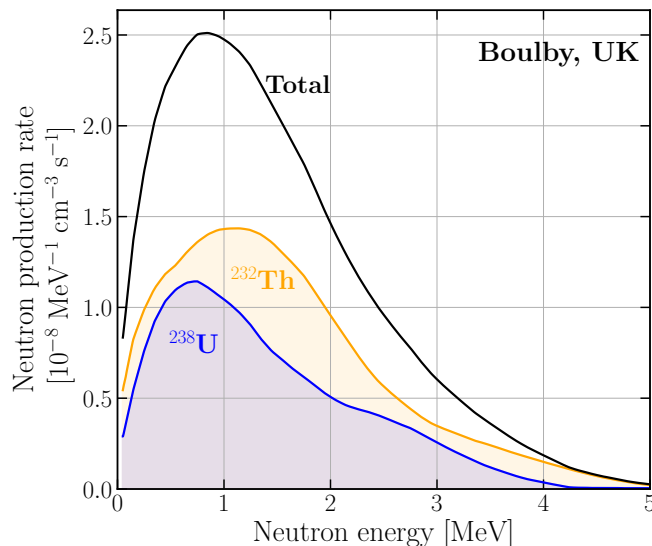


FIG. 28. Neutron energy spectrum from the salt rock cavern surface at Boulby. The blue and orange curves show the spectra due to ^{238}U and ^{232}Th respectively, and the black line shows the sum of the two.

1. Rock neutrons and passive shielding

Concerning the salt rock at the Boulby site we use the measured values of 70 ± 10 ppb of ^{238}U and 125 ± 10 ppb of ^{232}Th [326]. Figure 28 shows the energy spectrum of the neutron flux from isotopic decays of these nuclei. Since the flux saturates beyond 3 m, it is sufficient to only simulate the rock to this depth to accurately replicate the laboratory neutron flux. We then simulate the resulting nuclear recoils in the gas volume. To reduce the computational burden the simulations assumed 600 Torr of gas and then used the linear relationship between the nuclear recoil rate and the gas pressure to scale each result to 20 Torr. This should hold well assuming that the double scattering of neutrons in the gas remains negligible. Nevertheless, we have crosschecked that the rates at the nominal and increased pressure are consistent with this scaling. Water was selected for the neutron shielding and is simulated at different thicknesses surrounding the CYGNUS gas volume. The rock, water and gas material is simulated using a cubic geometry enclosing each respective volume. No space is simulated between the volumes as this was seen to have a negligible effect on the results. We find that a 75 cm thickness reduces the rate to an acceptable $\lesssim 0.93 \text{ yr}^{-1}$ (90% CL).⁶

⁶ Possible contamination of the water by radionuclides is not accounted for here as the background rate is likely to be subdominant. See for example Ref. [326], which lists water radioactivity measured as low as 10^{-4} ppb for both ^{238}U and ^{232}Th .

2. Vessel neutrons

The vacuum vessel—the TPC component with by far the largest mass—has the potential to dominate the neutron background. Here we consider the possibilities of vessels constructed of steel, titanium, copper or acrylic. For each material, we consider a thickness, T , ranging from 5–30 cm, and compute the neutron recoil rate, R_n , in the gas volume using the formula,

$$R_n = \frac{N_r}{N_{\text{tot}}} \Phi_n V a. \quad (15)$$

Where N_r/N_{tot} is the ratio between the number of neutron recoils that occurred within the gas volume and the total number simulated. The rate is proportional to the material volume, $V(T)$, and the neutron production rate within the vessel material, Φ_n , which is expressed as $V^{-1} s^{-1} a^{-1}$, where a is the contamination of the vessel material by radioisotopes. The values of a used here for ^{238}U or ^{232}Th isotopes within each material under consideration are listed in Table IV.

Material	^{238}U	^{232}Th	^{40}K	Reference
Steel	0.27	0.49	0.40	LZ [327]
Titanium	<0.09	0.23	< 0.54	LZ [327]
Copper	< 0.012	< 0.0041	0.061	NEXT-100 [328]
Acrylic	0.029	0.039	2.1	SNO+ [329]
Silicon	< 12.35	< 4.07	< 6.81	UKDMC [326]
Aluminum	< 0.52	1.94	< 6	[330]
Polyimide	< 36.79	< 27.52	< 410	NEWAGE [331, 332]
Kapton	< 98.77	< 36.59	58.82	UKDMC [326]

TABLE IV. Published material contamination levels and upper limits (in mBq kg^{-1}), used for our background simulation.

In some cases the published upper limits on a for certain materials will likely give background rates exceeding our requirements. In these cases we rearrange Eq.(15) to find the value that would be required to give $R_n < 1 \text{ yr}^{-1}$. The values of R_n and required purities are summarized in Table V. Note that in this and subsequent tables the upper limit is inherited from the isotope measurement and the error is the statistical error from the Monte Carlo simulation.

Some initial conclusions can be extracted from the results in Table V. Based on the most recent low background limits reported, both acrylic and copper produce significantly fewer neutron recoils than either steel or titanium of the same width. Acrylic in particular has a neutron self-shielding property due to its relatively high hydrogen content. However, while the construction of a 5–10 cm thick copper vessel or a 5–30 cm thick acrylic vessel would keep the neutron recoil rate within 1 per year, this may not be practical. In reality, a steel or titanium frame could be needed to support the vessel. The exact structure of such a frame is not explored here but from a background perspective, if this frame were to average 5 cm in width, then the ^{238}U and ^{232}Th content

Material	Thickness (cm)	Rate (yr ⁻¹)	²³⁸ U limit (mBq kg ⁻¹)	²³² Th limit (mBq kg ⁻¹)
Steel	5	21±4	0.016	0.019
	10	50±9	5.5×10 ⁻³	9.4×10 ⁻³
	20	177±25	1.8×10 ⁻³	2.3×10 ⁻³
	30	242±36	1.5×10 ⁻³	1.5×10 ⁻³
Titanium	5	< 11±3	0.013	0.015
	10	< 45±8	4.2×10 ⁻³	3.2×10 ⁻³
	20	< 88±15	2.9×10 ⁻³	1.5×10 ⁻³
	30	< 200±28	1.2×10 ⁻³	6.9×10 ⁻⁴
Copper	5	< 0.39±0.07	✓	✓
	10	< 1.0±0.16	✓	✓
	20	< 2.0±0.3	5.6×10 ⁻³	✓
	30	< 2.6±0.4	4.1×10 ⁻³	✓
Acrylic	5	0.11±0.01	✓	✓
	10	0.17±0.02	✓	✓
	20	0.21±0.04	✓	✓
	30	0.17±0.04	✓	✓

TABLE V. Neutron recoil rate for different vessel materials at different thicknesses. We display a checkmark where the existing ²³⁸U and ²³²Th limits give rates of less than one recoil per year. In cases where they do not, we instead give the value of radiopurity that would.

would need to be reduced by a factor of 17 and 26, respectively for steel, or 7 and 15 for titanium. It is worth remarking that the ²³⁸U activities for titanium and copper are upper limits, so further screening tests may reveal them to be within the limit.

An alternative solution for improving the radiopurity with steel or titanium would be to use an internal acrylic shield of around 10 cm. This width would not contribute significantly towards the total neutron background but it would help to shield against neutrons produced by a steel or titanium support frame. Running at atmospheric pressure may also reduce the vessel neutron background since it would require thinner walls and less support structure, as we explain in Sec. V E. Overall we can summarize that there is a good prospect for reaching a design of TPC vessel that can achieve the required level of neutron-induced background from that source.

3. Internal TPC neutron background

We now consider the unshieldable neutron backgrounds from readout components within the TPC. We simulate a 1 m² sheet readout positioned centrally inside the TPC gas volume, and the simulated rates are then scaled up to the required 2000 m². We focus on THGEM and GEM amplification stages, with readout stages comprising a generic future strip readout, wire-based readout or pixel chips.

There are some technologies which are neglected in this analysis due to the relatively high background of current devices, such as the micro-RWELL—which combines a GEM-like structure and printed circuit board readout—

and the micromegas, see Refs. [326] and [333] respectively. These are of high interest for CYGNUS, but would both need improved material radiopurities to be incorporated in a zero-background CYGNUS-1000. Similarly, current μ -PIC readout backgrounds exceed what is acceptable for a 1000 m³ detector, but for illustration we do include this readout in our simulation. Extensive work by authors here is underway to reduce μ -PIC backgrounds. We point out that even the current version of these higher-background readouts are still of interest for near-term smaller-exposure detectors. We also do not simulate readouts based on optical technology such as CCD or CMOS cameras. These would need to be placed outside of the vessel with low-activity transparent windows between the cameras and the TPC volume, thus requiring a major rethinking of the vessel design. For these reasons we focus only on the technologies listed in Table VI. We briefly describe each below.

Technology	Material	Thickness (mm)	Total mass (tons)
Gain stages:			
THGEM	Acrylic	1.0	2.36
THGEM	Copper	0.1	3.6
GEM	Kapton	(0.05)	0.0142
Readout stages:			
Strip	Acrylic-Cu	1.05	4.16
MWPC (wires)	Steel	0.05	1.94×10 ⁻³
MWPC (frame)	Acrylic	10	0.236
Pixel chip	Silicon	0.400	1.86
Pixel chip	Copper	3.9×10 ⁻³	0.07
Pixel chip	Aluminum	4.5×10 ³	0.024
Other:			
μ -PIC	Polyimide	1.0	2.84

TABLE VI. The simulated readouts, material components, thickness and total mass.

Starting with the amplification stage, a typical GEM has a 50 μ m-thick Kapton layer coated with a thin, \sim 5 μ m copper layer on either side. We only simulate the Kapton layer here as it is the main mass contributor and has a relatively high background compared to copper. In order to achieve high gas gains of \sim 10⁴–10⁵, the stacking of GEMs may be required [334], doubling or even tripling the Kapton background. A Thick GEM (THGEM) on the other hand has a thicker Kapton or acrylic layer, usually between 0.4 and 1 mm thick. We use the upper end of 1 mm for these simulations as the THGEM can then be made of low background acrylic at this thickness [335]. We also simulate 0.1 mm of copper on either side of the acrylic because the copper layers of the THGEM are significantly thicker than that of an ordinary GEM.

Turning to the readout technologies, for the strip readout we assume a single layer of x and y copper strips on a low background acrylic substrate and use the same radiopurity values as those adopted for the THGEM components detailed above. Regarding the wire-based readout (MWPC), this has two main sources of background:

the acrylic frame supporting the wires, and the wires themselves. The DRIFT-II detector consists of two MWPCs each with three planes of 552 steel wires, where the middle plane is made from 20 μm thick wire and the two outer planes 100 μm thick wire [242]. DRIFT-II uses the field between the middle plane, which is grounded, and the two outer planes, which are set to around -3000 V, to provide the signal amplification. If used in CYGNUS however, the MWPC readout would be grounded and a GEM or THGEM would provide the amplification stage. Therefore, with no voltage being applied, the wires do not need to be 100 μm thick. However, to ensure greater stability at larger scales, they do need to be slightly thicker than 20 μm . We simulate 50 μm thick wires with a total steel mass of 1.94 kg (based on the total steel wire mass of 1.94 g that makes up the two MWPCs in the 1 m^3 DRIFT-II detector). We assume two 1 m^2 wire planes separated by 1 cm with a 2 cm thick border as the acrylic frame. The pixel chip readout is based on the ATLAS FE-I4 pixel chip [260], made of silicon with metal and dielectric layers. Copper and aluminum make up the bulk of the metal layers, in addition to small amounts of tantalum, chromium and titanium. Here we consider a simplified model made from a 400 μm thick block of 98% silicon, 1% copper and 1% aluminum by mass. Finally, the $\mu\text{-PIC}$ readout as used by NEWAGE [336] is composed of a double-sided circuit board separated by a 1 mm-thick polyimide substrate, the latter being the main background contributor.

A large-scale TPC readout would require structural support, most likely made of acrylic. We do not explore the potential design and the impact on the background contribution because the focus here is on the base materials. We remark however that the inclusion of a large scale support structure would further increase the background inside the vessel for all readouts and that the exact amount of support required would depend on the readout requirements. For example, with an MWPC the forces and torques exerted due to the tension in the wires may mean that a more rigid support structure would be required than for other readouts. The basic acrylic MWPC frame investigated in this section is not for support but to simply hold the wires in position. For the same reasons, the background contribution from readout electronics downstream of the primary charge sensitive device is not studied here. However, in general this could be reduced by placing the electronics outside of the vessel or by using internal shielding. For the pixel chip readout, backgrounds from front end electronics are already accounted for, as they reside on the chip itself.

For the field cage and cathode, the acrylic frame from both of these structures provides the dominant mass, conservatively estimated to be 86.4 tons ($1000\times\text{DRIFT-II}$). However, the highest source of background is from alumina found inside resistors that make up the resistive divider used to define the drift field between the cathode and readout. Around 700 resistors would be required in a 10 m long CYGNUS vessel, based on ex-

trapolation from the DRIFT-II field cage. We assume radiopurity levels of 0.4 ± 0.2 and < 0.023 mBq/pc for ^{238}U and ^{232}Th respectively, following the levels reported by TREX-DM [252] for their field cage SM5D resistors. These resistors have dimensions $6.4 \times 3.2 \times 0.55$ mm [337]. A block of alumina with these dimensions is simulated at the center of the TPC gas volume and the resulting neutron background rate is scaled up to the required number of resistors.

Technology	Material	Neutron recoils (yr^{-1})	U/Th limit (mBq kg^{-1})
Gain stages:			
THGEM	Acrylic	0.122 ± 0.002	✓
THGEM	Copper	$2.4 \pm 1.6 \times 10^{-3}$	✓
GEM	Kapton	$< 9 \pm 1$	13.6/18.7
Readout stages:			
Strip	Acrylic-Cu	0.123 ± 0.002	✓
MWPC (wires)	Steel	$4 \pm 0.5 \times 10^{-4}$	✓
MWPC (frame)	Acrylic	0.048 ± 0.004	✓
Pixel chip	Silicon	25 ± 3	0.31/0.40
Pixel chip	Copper	$2.9 \pm 0.3 \times 10^{-4}$	✓
Pixel chip	Aluminum	0.29 ± 0.03	✓
Other:			
$\mu\text{-PIC}$	Polyimide	$< 160 \pm 16$	0.185/0.20
Resistors	Ceramic	$< 0.35 \pm 0.23$	✓

TABLE VII. Table of neutron rates and required radiopurities for the readout background. In the final right-hand column, a tick indicates that the U and Th levels are indeed seen to be sufficiently low for that material to reach the $< 1 \text{ yr}^{-1}$ goal. Where this is not true we give an estimate of the maximum allowed ^{238}U and ^{232}Th radiopurity levels (termed U/Th limit) that must be reached to obtain the $< 1 \text{ yr}^{-1}$ goal.

The predicted rates and required radiopurities for all the generic internal TPC components studied in this section are listed in Table VII. A full readout plane requires both a gain stage and a readout stage, with background rates from both parts contributing to the total background. An example would be a THGEM gain stage combined with a strip readout. In Table VII (and similarly Table IX) the first three rows refer to components of possible charge gain elements. For a single THGEM gain stage, values in the first two rows need to be summed. The various components for strip, wire and pixel readouts, are given in the following rows of the table. Values for the strip readout assume a 2d copper strip format on an acrylic support. The final rows give typical values for the necessary field cage resistors and an example $\mu\text{-PIC}$ readout [336]. Current values for Micromegas strip readout [338] are similarly high.

Assuming the radiopurity measurements of acrylic and steel, made by SNO+ and LZ respectively (see Table IV), it appears that only the THGEM amplification stage combined with either wire readout or future acrylic-based strip readout can likely satisfy the CYGNUS background criteria without significant improvements to intrinsic radiopurity. Although recent $\mu\text{-PIC}$ design developments

have allowed for a significant decrease in their alpha emission [331], a better understanding of the polyimide radiopurity will be required to estimate the true neutron-induced recoil rate more accurately. This is also the case for the GEM amplification stage and the pixel chip readout, where the activity levels of Kapton and silicon respectively are still only upper limits. For all these cases, further study could possibly reveal a lower neutron background, but if the true levels are close to the upper limits alternative materials would need to be found. Finally, readouts aside, the field cage fortunately has a neutron background within limit, based on using the TREX-DM resistor values. New resistive sheet technology [339], may also provide an alternative to conventional field cage structures in the future and possibly contribute a total lower background. Note also that in this study no allowance has been made for additional internal detector support structures.

These results are useful benchmark values. Ultimately however an important developmental step for CYGNUS will be to perform extensive material screening to resolve these uncertainties further and to develop new ways to install the readout planes that minimize material use. The background from all readouts could be lowered by using more radiopure materials or by attempting to improve the radiopurity of the same materials. This is particularly relevant to the μ -PIC readout and GEM amplification. For instance, better-refined glass polyimides are being developed for the μ -PIC, and a Kapton replacement for the GEM made from high purity G10 insulator is already being investigated by some of the authors. The background for pixel chips could also be reduced if charge focusing [321] becomes a possibility, as this would reduce the number of units required to fill the readout plane.

4. Muon-induced neutrons and active vetoing

Neutrons produced by muon interactions in and around the detector are another source of background. To compute the energy spectrum, flux and angular distribution of incoming muons we use the MUSUN code [325]. The simulation takes into account the angular profile and the composition of the rock overburden for the transportation of cosmic-ray muons. The muon energies, positions, and momenta are then put through Geant4, which transports them through 20 m of rock before reaching the detector. More than 200 million muons are generated at the surface of the rock volume. Those reaching the rock-cavern boundary are then normalized to the measured muon flux at Boulby: $(4.09 \pm 0.15) \times 10^{-8} \text{ cm}^{-2}\text{s}^{-1}$ [340] (with a vertical rock overburden of $2805 \pm 45 \text{ m.w.e.}$). Over a period of $(2.8 \pm 0.1) \times 10^7 \text{ s}$, no events are recorded within the fiducial volume, providing us an upper limit on the rate of muon-induced neutron nuclear recoils of 2.75 yr^{-1} . More extensive simulations will be required to see if the actual rate might be lower. If not, these remaining few events could be identified and tagged using




FIG. 29. Gamma-induced electron recoil background contribution observed in the 20 Torr SF_6 gas volume, originating from a 30 cm thick copper vessel. We see clearly that each nuclide has an essentially flat recoil spectrum that lies below $\lesssim 10^4 \text{ keV}^{-1} \text{ yr}^{-1}$.

either an active muon veto, such as a plastic scintillator, or by looking for coincident electromagnetic events in the gas. This could be made possible by installing PMTs in the water shielding described earlier.

B. Gamma backgrounds

The challenges in performing electron/nuclear recoil discrimination below 10 keV_{ee} were discussed in Sec. IV H. The electron rejection rises exponentially with energy, and is expected to exceed 10^4 between 1 and 10 keV_{ee} , even after diffusion. The exact threshold where this occurs will be readout-dependent, and must be determined in the future. We focus therefore on the gamma-electron recoils in this range, aiming to find vessel and readout materials with intrinsic radiopurity levels that can allow the detected background rate to be limited to $\lesssim 10^4 \text{ keV}^{-1} \text{ yr}^{-1}$.

Using Geant4, we homogeneously populate the rock, vessel, and sheets of readout materials with ^{238}U , ^{232}Th , and ^{40}K , assuming secular equilibrium. The 75 cm water shield is also included for these simulations (see Sec. V A). The corresponding gamma fluxes are then computed following the simulation of the decay chains of each isotope. For materials with relatively small thicknesses, such as the readouts, the homogeneous distribution of the isotopes allows for the simultaneous simulation of the bulk and surface background. For each isotope, the electron

recoil rate due to gammas was computed as,

$$R_\gamma = \frac{N_r}{N_{\text{tot}}} M a. \quad (16)$$

Again N_r is the observed number of recoils and N_{tot} is the total number of decays simulated for the isotope in question. As before, a is the material radioactivity per unit mass for that isotope and M is the total material mass. The recoil spectrum is flat to a good approximation over the 1–10 keV range because the majority of gamma recoils are due to Compton scattering. Figure 29, shows an example of this: the recoil energy spectrum for gammas originating from a 30 cm copper vessel and scattering within the fiducial volume. This example is also below our electron recoil target rate, as we discuss next.

1. Rock, vessel and shielding gammas

As in our discussion of the rock neutron background, we assume the measured contamination levels of the salt rock at Boulby, now with the inclusion of 1130 ± 200 ppm ^{40}K [326]. For rock gammas, an even smaller 25 cm depth is all that contributes to the background, beyond which the rock is self-shielding. Since we found that 75 cm of water shielding was needed to shield the rock neutrons, we now include this also in our rate calculation.

The rock and vessel gamma-induced electron recoil rates for each material and thickness are listed in Table VIII. As before for the neutron results (see Table VII) in the right-hand columns a tick indicates where we find the levels sufficiently low for that material to reach the $< 10^4$ keV $^{-1}$ yr $^{-1}$ goal for the listed isotope. Where this is not true we give an estimate of the maximum allowed radiopurity levels (termed U/Th/K limit) that must be reached to obtain this goal.

For metal vessel materials, rock gamma recoils are reduced by ~ 3 orders of magnitude as the material thickness is increased from 5 to 30 cm. For acrylic, the corresponding reduction is less than 1 order of magnitude, so that acrylic by itself is insufficient to shield the rock gammas to acceptable levels. With the exception of acrylic, the radioactivity from the vessel materials does not vary significantly with material thickness. This is due to the gamma self-shielding properties of metals which cause the gamma flux to saturate at relatively low thicknesses. Therefore, there is an unavoidable gamma flux in the metals that cannot be mitigated by increasing the thickness. This occurs at a flux of $\sim 7 \times 10^5$ keV $^{-1}$ yr $^{-1}$ at 20 cm, $\sim 5 \times 10^5$ at 30 cm, and $\sim 10^4$ at 30 cm for steel, titanium, and copper, respectively. This effect does not occur in acrylic. Indeed, Table VIII shows that acrylic is not comparatively effective as gamma shielding and as such would not be an appropriate vessel material, at least not by itself. The only way the saturated gamma flux from the metals can be addressed further is if the

radioactivity levels are lowered from what we are assuming here. Since some of the values from Table IV are in fact only upper limits, further screening tests may reveal certain materials to be within limit in addition to those already shown. For example, with 30 cm thickness, copper gives a rate only slightly above 10^4 keV $^{-1}$ yr $^{-1}$ and a ^{238}U activity of around one half of the upper limit assumed here would be acceptable. A purely copper vessel could be a viable option but as we discussed in Sec. V A 2 if a steel or titanium frame were needed in addition to the internal acrylic shield then every material would need a reduction in radioactivity, if actual rates are close to the limits seen.

2. Internal TPC gamma background

Next, we consider the gamma background due to the same readout systems and field cage resistors listed in Sec. V A 3. For the latter, we assume the same radiopurity levels of ^{238}U and ^{232}Th from before, but we now also include the 0.17 ± 0.7 mBq pc $^{-1}$ of ^{40}K [252].

Table IX shows the gamma-induced electron recoil rates, given here in the same manner as Table VII, split between gain stage and readout components. We can see clearly that the ^{40}K radioactivity levels of acrylic causes the dominant gamma background for the THGEM and wire readout, both of which produce a background rate only slightly above limit. Cutting the current ^{40}K activity of acrylic in half would bring the wire frame gamma background to within 10^4 keV $^{-1}$ yr $^{-1}$. Similarly, a factor of four would be needed for a THGEM gain stage and for the generic strip readout. For a single GEM gain stage however one would need an order of magnitude reduction; as with acrylic, the ^{40}K activity of Kapton is the dominant factor here. Additionally, based on the most current measurements of micromegas readouts [338] and of μ -PIC readouts [331], we confirm that current typical radioactivity levels are, respectively, a factor of 10^6 and 10^3 away from matching the activity produced by the 1 mm thick acrylic, as used here for the conceptual strip readout example. For the polyimide in the μ -PIC readout, further screening is underway to better understand its radiopurity levels. In the pixel chip, silicon is the main source of gammas, since they only contain a small quantity of copper and aluminum. Again, since the silicon radiopurity is an upper limit, true levels could be lower, and more sensitive tests are required. Also, charge focusing [321]—as mentioned earlier—could help to reduce the number of pixels needed and therefore the background. Finally, the gamma background caused by the field cage resistor chain is twice the tolerable limit and a reduction in the ^{238}U activity of ceramic by a factor of four is needed. Again, resistive sheets [339] should be evaluated as a potential resistor replacement.

In summary, the general conclusion for the gamma-induced backgrounds internal to the TPC is that some existing technologies are already close to meeting the de-

Material	Width (cm)	Rock γ recoils (keV ⁻¹ yr ⁻¹)	Vessel γ recoils (keV ⁻¹ yr ⁻¹)	Total γ recoils (keV ⁻¹ yr ⁻¹)	²³⁸ U limit (mBq kg ⁻¹)	²³² Th Limit (mBq kg ⁻¹)	⁴⁰ K Limit (mBq kg ⁻¹)
Steel	5	3.8±0.3×10 ⁶	6.6±0.6×10 ⁵	4.4±0.4×10 ⁶	0.003	0.0045	0.08
	10	6.0±1.0×10 ⁵	7.2±0.9×10 ⁵	1.32±0.19×10 ⁶	0.003	0.004	0.06
	20	2.1±0.6×10 ⁴	7.3±1.4×10 ⁵	7.5±1.5×10 ⁵	0.0027	0.0042	0.075
	30	4.6±3.0×10 ³	6.3±1.5×10 ⁵	6.3±1.5×10 ⁵	0.003	0.0053	0.053
Titanium	5	1.0±0.2×10 ⁷	< 2.9±0.2×10 ⁵	< 1.0±0.2×10 ⁷	0.003	0.0046	0.06
	10	3.8±0.9×10 ⁶	< 4.13±0.36×10 ⁵	< 4.2±0.9×10 ⁶	0.0022	0.0031	0.05
	20	6.6±1.1×10 ⁵	< 4.17±0.53×10 ⁵	1.08±0.16×10 ⁶	0.002	0.0035	0.041
	30	< 4.8±3.1×10 ⁴	< 5.11±0.71×10 ⁵	< 5.6±1.0×10 ⁵	0.0017	0.0027	0.041
Copper	5	2.3±0.2×10 ⁶	< 1.57±0.17×10 ⁴	2.3±0.2×10 ⁶	0.0057	✓	✓
	10	4.0±0.9×10 ⁵	< 1.60±0.24×10 ⁴	4.1±0.9×10 ⁵	0.0058	✓	✓
	20	9.5±4.0×10 ³	< 1.58±0.33×10 ⁴	< 2.53±0.73×10 ⁴	0.0056	✓	✓
	30	5.1±3.3×10 ²	< 1.58±0.43×10 ⁴	< 1.6±0.5×10 ⁴	0.0053	✓	✓
Acrylic	5	2.5±0.3×10 ⁸	3.44±0.32×10 ⁵	2.5±0.3×10 ⁸	0.0002	0.0017	0.037
	10	1.90±0.19×10 ⁸	5.97±0.57×10 ⁵	1.90±0.19×10 ⁸	5.7×10 ⁻⁴	9.3×10 ⁻⁴	0.024
	20	9.7±1.4×10 ⁷	1.14±0.12×10 ⁶	9.8±1.4×10 ⁷	3.4×10 ⁻⁴	5.4×10 ⁻⁴	0.011
	30	4.1±0.9×10 ⁷	1.12±0.14×10 ⁶	4.2±0.9×10 ⁷	3.2×10 ⁻⁴	4.9×10 ⁻⁴	0.013

TABLE VIII. Gamma recoil rates in a 20 Torr SF₆ target volume, for vessel wall materials of varying thicknesses, and the required radiopurities to obtain a rate of 10⁴ keV⁻¹ yr⁻¹.

Readout	Material (width)	γ recoils (keV ⁻¹ yr ⁻¹)	U limit (mBq kg ⁻¹)	Th limit (mBq kg ⁻¹)	K limit (mBq kg ⁻¹)
Gain stages:					
THGEM	Acrylic (1 mm)	3.3±0.7×10 ⁴	✓	✓	0.54
	Copper (0.1 mm ×2)	< 1.5±0.3×10 ³	✓	✓	✓
GEM	Kapton (50 microns)	1.57±0.02×10 ⁵	✓	✓	3.65
Readout stages:					
Strip	Acrylic-Cu (1 mm)	3.4±0.7×10 ⁴	✓	✓	0.54
Wires	Steel (50 μ m)	1.8±0.3	✓	✓	✓
	Acrylic (2 cm × 1 cm)	2.4±0.1×10 ⁴	✓	✓	0.88
Pixel chip	Silicon (400 μ m)	< 2.55±0.19×10 ⁵	0.26	0.29	0.46
	Copper (3.9 μ m)	< 24±2	✓	✓	✓
	Aluminum (4.5 μ m)	< 937±77	✓	✓	✓
Other:					
μ -PIC	Polyimide (1 mm)	< 1.3±0.2×10 ⁷	0.12	0.09	0.12
Resistors	Ceramic	2.5±1.3×10 ⁴	0.13	✓	✓

TABLE IX. Readout gamma recoil rates in 20 Torr SF₆ for different readout materials and the ²³⁸U, ²³²Th, and ⁴⁰K radiopurity required to achieve 10⁴ recoils keV⁻¹ yr⁻¹ (as well as the materials which already satisfy this requirement).

sign goal of 10⁴ keV⁻¹ yr⁻¹. This appears particularly true for the THGEM-wire combination and generic strip readout. Most other technologies and materials will require further screening and/or further reductions in background levels before they can be shown to be viable for CYGNUS.

C. Radon and radon progeny backgrounds

Radon gas emanating from materials is a further source of background for rare-event experiments. In particu-

lar, the low energy (~ 100 keV) of radon progeny recoils (RPRs) can mimic WIMP interactions [16]. Being a noble gas, ²²²Rn has a low chemical reactivity making it particularly difficult to filter. Moreover, its 3.8 day half-life allows enough time for the radon gas to spread from the materials in which it is produced, resulting in a rather widespread source of background.

Alpha particles from radon decays are relatively easy to identify. For a back-to-back TPC (such as that shown in Fig. 1 of Sec. I), events that trigger on both the left and right readout planes of the detector can be spotted via their coincidence in time. The only events with tracks




FIG. 30. Amplitude as a function of time of the main signal peak as well as several minority peaks from the DRIFT-IIId gas mixture, adapted with permission from Battat et al. [241]. The minority peaks are labelled D, P and S by convention.

long enough to achieve this are alpha particles originating from an RPR event fully contained within the fiducialized gas volume. These types of events can be used to estimate the level of radon present in the detector [16, 253]. However, not all RPR events originate within the fiducial volume as the positively charged radon daughter nuclei drift towards the central cathode and become stuck within the material. A recoil event caused by the decay of RPRs can prove difficult to reject as the associated alpha particle may not be identified if it remains trapped. One method of reducing the number of trapped particles is to have smaller cathode widths, thus increasing the transparency. This approach was successfully implemented in DRIFT with the use of a submicron thick cathode. This helped to reduce the RPR rate from approximately 130 events per day to less than one [341]. In DRIFT-IIId, RPRs originating from the cathode surface without a detected alpha track are vetoed by defining the fiducial region such that any recorded signal within 2 cm of the cathode is rejected [242]. A similar technique is also used by MIMAC [253] where an analysis cut based on the z positions of events is used to reject RPRs originating from the cathode. Ideally, a similar z -cut to veto RPRs would be used in CYGNUS.

Many techniques for dealing with RPRs rely on fiducialization: the ability to fully locate in three dimensions the source of ionization within the TPC gas volume. This is of vital importance for rare event search TPCs as it allows for background events from material surfaces to be distinguished from potential signal events occurring within the target volume. Fiducialization in the TPC readout plane (x, y) can be done at the reconstruction stage for the more segmented readouts.

For less segmented readouts one can use a surrounding veto area that triggers when events enter from outside of the fiducial volume. Fiducializing parallel to the drift field (z) presents more of a challenge as no active veto can be placed in front of either the cathode or readout. However, z -fiducialization has been achieved in previous work [264, 298, 301]. DRIFT-IIId uses a gas mixture of CS_2 , CF_4 and O_2 at a ratio of 30:10:1 Torr to produce distinct groups of “minority carriers” within the gas that drift at different speeds. The signal then has a main peak adjacent to three minority peaks as shown in Fig. 30. The separation in time between these peaks is proportional to the drift distance traveled and therefore the z position at which the primary event occurred. Reference [299] showed that the same fiducialization method is possible with SF_6 due to the presence of an SF_5^- minority peak; so we believe fiducialization should be achievable in CYGNUS. There are other options available, such as using timing differences between the signals at the cathode and induced signals at the anode as in MIMAC [342], or by measuring diffusion of drift charge transverse to the recoil axis as demonstrated by D³ [264, 301]. Several of these techniques may also be combined to further improve performance. Note that fiducialization via minority carriers is only possible with negative ion drift gases, as proposed here, so that we expect the best fiducialization in such gases.

D. Cosmogenic activation

When materials are present at the surface they are subject to a much higher flux of cosmic rays than when underground. If materials spend a significant time above ground cosmic ray interactions can activate different long-lived isotopes, which when part of the detector become an additional source of background. As rare event searches increase in both size and sensitivity, this cosmogenic activation is going to have an increasing influence on the ultimate background rate. See for example Ref. [343] for further details.

The activity of a cosmogenically activated isotope is expressed as,

$$a = R[1 - e^{-\lambda t_{\text{exp}}}]e^{-\lambda t_{\text{cool}}}, \quad (17)$$

where R is the production rate, λ is the decay constant, t_{exp} is the time the material spent exposed to the cosmic ray flux and t_{cool} is the ‘cooling off’ time the material spent underground.

Activation can occur in all materials, but we use copper here as an example since it is our best candidate vessel material so far. Reference [344] showed that cosmic ray interactions with copper at sea level can produce multiple unstable isotopes. The longest-lived is ^{60}Co which originates from ^{59}Co contamination and has a half-life of 5.26 years. The measurement from Ref. [344] was taken after a relatively long $t_{\text{exp}} = 345$ days compared to the much shorter $t_{\text{cool}} = 14.8$ days. The saturation activity

for ^{60}Co was $340 \pm_{68}^{82} \mu\text{Bq kg}^{-1}$. Taking the upper end of this measurement, we estimate that the rate of gamma-induced electron recoils for a 5 cm thick copper vessel will be $< 9.8 \pm 0.6 \times 10^4 \text{ keV}^{-1} \text{ yr}^{-1}$ in the 1–10 keV energy region. This rate is ~ 6 times larger than the recoil rate due to ^{238}U , ^{232}Th , and ^{40}K radioactivity in the same thickness of copper (see Sec. VB1), assuming the upper limits are the actual radiopurity values.

Since the measurement and simulation of the cosmogenic activation in copper have shown good agreement [344], for a more detailed study we can make use of the simulation tool ACTIVIA [345]. This enables us to predict the minimum amount of time a material should spend above ground and cooling off underground in order to limit activation.

For SF_6 , the longest-lived isotope of fluorine, ^{18}F (with a half-life of 109.77 minutes), is still too short-lived to produce a lasting background contribution. Nearly all isotopes of sulfur have half-lives on the scale of seconds to minutes with the exception of ^{35}S , which has a half-life of 87.5 days. We simulate 90 days of surface time with ACTIVIA, during which the gas is exposed to cosmic rays sampled from an energy spectrum ranging between 10–10000 MeV. After this we assume a 180 day cooling-off period underground, during which time the gas is not exposed to cosmic rays. This resulted in a ^{35}S production rate of $R = 0.021 \text{ kg}^{-1} \text{ day}^{-1}$, giving a decay rate of 2.61 mBq/kg. A 1000 m^3 vessel filled with SF_6 at 20 Torr will hold ~ 160 kg of gas, of which ~ 35 kg is sulfur, so the total decay rate for the detector is 91.35 mBq. We simulate 10^5 ^{35}S decays in Geant4, originating from the center of the 20 Torr SF_6 gas volume. Only three gamma induced electron recoil events between 1–10 keV were observed in the simulation time of ~ 12.5 days. We estimate therefore a rate of $90 \pm 50 \text{ yr}^{-1}$, too small to be a significant electron recoil background.

In general the amount of cosmogenic activation is reduced by limiting t_{exp} and extending t_{cool} . In addition one should limit the amount of time the material spends as air-shipment and therefore subject to a higher cosmic ray flux. One might also want to electroform the metallic components underground, at the construction site, where possible.

E. Backgrounds at atmospheric pressure

A 1000 m^3 vessel filled with low-pressure target gas would need to withstand outside pressures of ~ 760 Torr, requiring larger amounts of structural support than a vessel filled with atmospheric pressure gas. Sections VA2 and VB1 found that a steel or titanium support structure would increase the total vessel background. Therefore, running at atmospheric pressure rather than at 20 Torr SF_6 would be advantageous. An addition of 740 Torr He allows for TPC operations at atmospheric pressure, while increasing the density only by a factor of ~ 2 when compared to 20 Torr SF_6 alone (see Table I).

Adapting our previous simulation to this 740:20 Torr mixture, we find that an increase in density corresponds to an increase in both the gamma and neutron background by around the same factor of ~ 2 . The results presented thus far throughout Sec. V can scale straightforwardly between low pressure and atmospheric pressure by this factor. In addition, the results can be scaled to give an approximation of the background rates at other ratios of He: SF_6 . For example, a ratio of 755:5 Torr would incur a density increase and, therefore, a background rate increase of around 25 percent compared to pure SF_6 at 20 Torr.

A major advantage of atmospheric pressure operation is not only the reduction in vessel support needed, but also that the vessel itself could be built with a lower average thickness. A 1000 m^3 CYGNUS vessel operating at 760 Torr would need an average thickness of around 30, 12, 10 or 8 mm of acrylic, copper, steel or titanium respectively [346]. Due to the strength of titanium it may even be possible to construct a vessel with as little as 6 mm average thickness. The acrylic estimate does not include additional structural support which would also need to be taken into consideration, but the copper, steel and titanium thicknesses are averaged over both the vessel walls and structural support. We employ the simulation described in Secs. VA2 and VB1 to re-estimate the neutron and gamma recoil rate due to vessel materials for the new 740:20 Torr He: SF_6 gas mixture. Table X lists these estimates. Note that the rock and water shielding are not included now, but expected background changes due to these will be commented on below.

Material	Thickness (mm)	Gamma recoil rate (yr^{-1})	Neutron recoil rate (yr^{-1})
Acrylic	30	$3.57 \pm 0.03 \times 10^5$	0.40 ± 0.02
Copper	12	$< 1.50 \pm 0.02 \times 10^4$	$< 0.12 \pm 0.02$
Steel	10	$5.75 \pm 0.08 \times 10^5$	5.6 ± 1.0
Titanium	8	$< 2.21 \pm 0.02 \times 10^5$	$< 4.0 \pm 0.7$
Titanium	6	$< 1.88 \pm 0.02 \times 10^5$	$< 2.4 \pm 0.4$

TABLE X. Vessel backgrounds for different vessel materials at reduced thickness, for an atmospheric-pressure 740:20 Torr He: SF_6 gas mixture scenario.

We can compare the 740:20 Torr mixture results of Table X with the 20 Torr SF_6 results at the thinnest previously simulated vessel thickness, 5 cm, shown in Table VIII. While the gamma recoil rates are similar, the neutron rates from the metallic vessel walls are all lower with the atmospheric pressure. However our general conclusion from before—that the steel and titanium vessel neutron background exceed the design target—still persists at atmospheric pressure.

It seems that the increase in background due to the increased gas density is approximately compensated by a decrease in background due to the reduction in required vessel material. So at the very least we have shown that the vessel background does not *increase* for the atmo-

spheric gas mixture. However the rock-gamma shielding efficiency would suffer with a thinner vessel. If an additional copper shield were to be implemented then this too would contribute to the total background, requiring potentially a thickened water shielding to counterbalance and keep the total background acceptable.

We have seen that for a fixed detector, background rates are proportional to gas density. Hence can already conclude that a 755:5 Torr He:SF₆ target gas, which has only 25% higher mass density than 20 Torr SF₆, is already close to acceptable from a background point of view, provided we do not change the vessel design. It seems quite likely, that with thinner vessel material that can be considered for atmospheric pressure, backgrounds can be reduced further, probably below the 20 Torr SF₆ scenario that has been studied in detail here. We therefore are cautiously optimistic, but also note that detailed, dedicated simulations of complete atmospheric pressure scenarios should be carried out in future work with high priority.

F. Conclusions

In this section we have discussed in detail the intrinsic backgrounds for a 1000 m³ CYGNUS TPC. We have taken into account neutron and gamma backgrounds originating radioisotopes in the rock, vessel and detector, as well those originating from cosmic ray muons, radon emissions and cosmogenic activation. Following Sec. IV C, we have aimed to determine the feasibility of limiting electron recoils to a rate of 10⁴ recoils keV⁻¹ yr⁻¹ in the energy range 1–10 keV_{ee}, and nuclear recoils to a rate of < 1 yr⁻¹ between 1–100 keV. These numbers are compatible with zero electron background after offline electron rejection and limit nuclear recoils from neutrons to less than a sixth of the solar neutrino signal. The feasibility of achieving the electron background limit is essential as it will ultimately control the threshold of the experiment and the sensitivity to low WIMP masses, whereas meeting the nuclear recoil rate target is important for optimizing sensitivity to WIMPs and solar neutrinos.

While we have assumed a fill gas of SF₆ at 20 Torr in much of our discussion, we have also shown how the results scale to atmospheric pressure He:SF₆ mixtures, which appear most desirable for CYGNUS. From a background perspective, SF₆ is advantageous due to its minority carriers which enable *z*-fiducialization. A vessel running at atmospheric pressure also requires less structural support and thinner walls than a vessel running at near-vacuum pressure and therefore produces less intrinsic background.

Summarizing the results we can state first that to shield rock neutrons the vessel would need to be surrounded by 75 cm of water shielding to bring the nuclear recoil rate, above 1 keV_r, down to less than one event per year. Next, for the vessel material itself, we determined that 20–30 cm of copper is optimum. The

electron recoils from the rock gamma background are already shielded to within our limit using this thickness, whereas the neutron and gamma backgrounds from the vessel fall only slightly above it. If it can be demonstrated in further screening tests that copper’s ²³⁸U radioactivity is around half the upper limit used here (< 0.012 mBq kg⁻¹ [328]) then it will be an attractive candidate for the CYGNUS TPC. Although such a vessel would likely need a steel support frame—introducing a further background—conceptual designs are feasible in which the frame is embedded in the copper, thereby providing the necessary additional shielding. Ideally this vessel would spend only a short time above and below ground, by limiting air-shipment of materials, procuring the materials and, if possible, electroforming, as close to the site as possible; all in the aim of reducing further cosmogenic activation.

In simulating the various readout technologies and amplification stages we conclude that while some readout-generated backgrounds can be kept within the limit, further R&D is clearly still needed to establish the optimum choice for CYGNUS from a background standpoint. A few issues that would need to be addressed include: (1) for strip micromegas, GEMs and μ -PIC, a better understanding of the radiopurity, as well as an investigation into the use of lower background material alternatives such as acrylic; (2) for the pixel chip readout, measurements of the radioactivity levels of silicon, combined with an investigation of charge focusing to reduce the required number of chips; and (3) an improved field cage design using a smaller number of resistors, which ideally would contain lower levels of ²³⁸U than the ceramic resistors we assumed here. Resistive sheets [339] should be evaluated as a potential resistor replacement.

The issue of the readout background may also be partially addressed on the analysis side. For instance, improving the electron/nuclear recoil discrimination by a factor of 10-100 would both widen the available choices of readout material, while also permitting background-free operation with a lower energy threshold. This may be possible by using a larger set of recoil discrimination parameters than already discussed in Sec. IV C, and/or by decreasing the drift distance to limit diffusion.

As long as the total neutron recoil background remains subdominant to solar neutrinos, then the WIMP cross section sensitivities from Sec. IV I would be largely unaffected. We reiterate that the most important impact on our final physics reach at this stage comes from the electron discrimination threshold. Hence the electron background and electron discrimination should take priority in future work. Ultimately, while we conclude that while there is more work still needed to prove that CYGNUS can sustain completely background-free operation, we do not yet consider there to be any immediate showstoppers, and have identified steps that should be taken to address our unresolved issues.

VI. UNDERGROUND SITES AND ENGINEERING

We displayed the locations of several candidate underground sites that are currently under consideration for CYGNUS in Fig. 2. This section provides a summary of the significant logistical and engineering requirements that must be considered for these sites to host directional dark matter detectors. Principal of these for gas TPCs are space requirements and the need to provide handling of the gas target underground. Both of these factors depend on the choice of gas and operating pressure, so we center the discussion around a first stage CYGNUS TPC of order 10 m^3 operating at one atmosphere.

A. Site requirements

While a TPC with 1000 m^3 fiducial volume is rather large in total volume, a corresponding advantage of the concept is that there is no particular restriction on the shape of that volume. For instance, one could consider an elongated ‘worm-like’ sectional design where one dimension is substantially longer than the other two. This is not only feasible but actually has certain advantages for operations and maintenance. CYGNUS is therefore well-suited to sites where long tunnels and restrictions on available height are normal, *e.g.* mine sites like Stawell, Australia or Boulby, UK. The relative simplicity of the services needed—in particular that there are no cryogenics involved—is also an advantage here. Additionally, the experiment can be built as separate detectors, either in one site, or in multiple sites. This also allows cross-correlation of data at different latitudes which could aid in the control of systematics. Such geometric features will not necessarily be required for a first 10 m^3 stage but will become important in 1000 m^3 designs aiming to go below the neutrino floor.

Due to finite door, elevator, and cavern sizes, certain sites will have more stringent constraints than others on the maximum size of an object able to be brought in, *e.g.* $2 \times 2 \text{ m}$ at Boulby. This would need to be accounted for in the vessel design. A baseline single $10 \times 10 \times 10 \text{ m}^3$ detector could be engineered in sites such as Gran Sasso, where there is available head-room and good access. A different, segment-based design could be used for a detector of the same total volume distributed in different sites, as well as for an elongated version of, say, $5 \times 5 \times 40 \text{ m}^3$ that would conceivably fit in an existing facility at Boulby.

The site overburden requirement is not foreseen to be as important a consideration here since the candidate sites are all $> 1 \text{ km}$ in depth. The muon-induced neutron event rate at this depth (calculated in Sec. V A 4) can be rejected with a combination of an active external muon veto and the inherent charged particle tracking capability of TPCs, depending on the adopted TPC readout technology. The need for an external veto would have only a

modest impact on the overall dimensions of the experiment, probably a $< 1 \text{ m}$ increase in linear dimension for the case of plastic scintillator with support structure.

B. Engineering requirements

One of the most important engineering constraints for gas TPCs is that they typically need to be operated at low pressure, requiring a vacuum vessel. We have proposed here however that the addition of helium—such that the experiment operates at atmospheric pressure—alleviates this constraint whilst simultaneously gaining the experiment access to low WIMP masses and a larger rate of solar and reactor neutrinos. The impact on directional readout sensitivity and the intrinsic background were discussed in Secs. IV and V respectively, and the results were promising. From an engineering perspective however, it is possible that the TPC would still need significant certification as a pressure vessel if contamination is an issue. Purification is conventionally achieved via outgassing under vacuum, so it would need to be proven that flushing gas through the vessel prior to operation can sufficiently reduce contamination. If possible this would obviate the need for the vessel to be of vacuum standard, greatly reducing the mass of steel, titanium or copper needed. This option is under investigation. We note that the NID gas CS_2 is known to be highly tolerant of impurities so that the DRIFT experiment operates well with up to 1% gas impurity, but SF_6 needs more study in this regard.

Figure 31 shows an example design for an 8 m^3 CYGNUS vessel based on a steel construction. An alternative is to use a cylindrical design to reduce the mass by a factor of about two, but at the cost of additional complexity in the detector and shielding.

For the design of Fig. 31, it is feasible to obtain the required ~ 4 tons of steel with sufficient radiopurity to reach the goal of < 1 background neutron per year, however the intrinsic gamma background would still be too high (see Table VIII). An alternative would be copper, as proposed for NEXT-100 [347] whose screening process proved ultra low-background levels of $< 0.012 \text{ mBq/kg U}$, $< 0.0041 \text{ mBq/kg Th}$, and 0.061 mBq/kg K . Although it may not be possible to build a large-scale pure copper vessel, a hybrid design in which copper is used for the main vessel panels with steel or titanium support structures could be engineered instead. Ultimately, we envisage that the scale and engineering requirements for a 10 m^3 CYGNUS vessel are unlikely to be challenging for any of the possible underground sites.

In addition to the $\mathcal{O}(10 \text{ m}^3)$ vessel, external neutron and gamma shielding would likely contribute an additional $\sim 1 \text{ m}$ surrounding the detector. A variety of designs are possible for this. The use of interlocking plastic water containers is one possibility, as adopted for the DRIFT upgrade shielding shown in Fig. 32. Another is to use a purpose-built water container. Though more ex-




FIG. 31. Design for a first stage 8 m^3 steel vessel (courtesy of D. Warner, Colorado State University).



FIG. 32. DRIFT-IIId interlocking acrylic block shielding surrounding the detector and containing water, located at Boulby underground laboratory.

pensive, it would have a potential advantage in that it could be instrumented with photomultipliers to provide a Cherenkov muon veto as well.

While there is no need for complex cryogenics engineering in CYGNUS, consideration is still needed with regards to the engineering aspects of the gas supply. In most current generation directional detectors the target gas is flowed and disposed of through filters to the atmosphere. However, for CYGNUS, recirculation will be needed, both for cost and environmental reasons: SF_6 is a powerful greenhouse gas. Recirculation also provides a potential means for the reduction of radon and water from the target. Purification of SF_6 is well known in industry [348] so

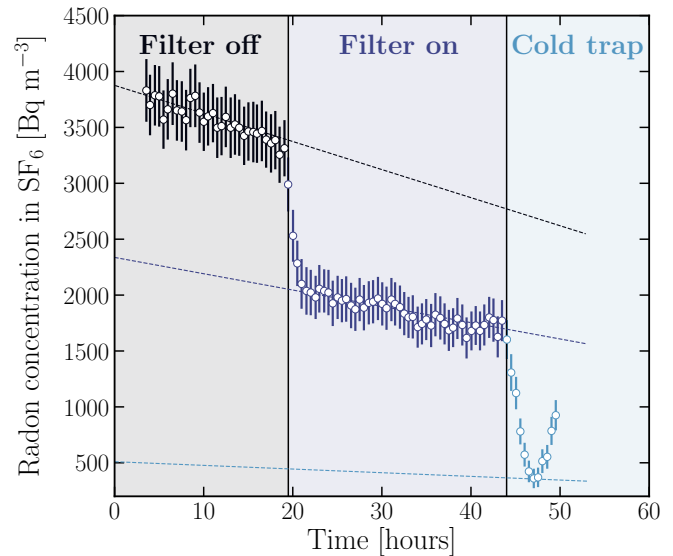


FIG. 33. Experimental results for the radon in SF_6 as a function of time when using a 5 \AA molecular sieve. The black points show the radon concentration prior to being diverted through the sieve. Then dark blue points show the reduction when the filter is turned on. The filter is then cooled with dry ice resulting in a further reduction, shown in light blue.

therefore is not seen as a major impediment here. Meanwhile, recent experiments by the Sheffield group have investigated active radon removal in SF_6 [349]. Figure 33 shows the results from an experiment in which SF_6 was circulated through a vessel with a known level of radon added via a sealed radon source. When the gas is diverted through a 5 \AA molecular sieve, the radon is seen to reduce. When the filter is cooled with dry ice a further reduction is seen. The final up-turn in this data is believed

to be due to the gradual temperature increase of the dry ice over time. An earlier test also checked that the SF₆ itself was not absorbed by the filter. These experiments show for the first time that active radon removal in SF₆ is possible.

Summarizing this brief overview of the CYGNUS site and engineering issues we conclude that, while the required vessels do present challenges, they are surmountable. One of the major benefits of the CYGNUS experiment is the significant flexibility in the approach to its shape and modularization. This opens up prospects for construction in sections underground and at multiple locations so that various constraints imposed by each site can be met. It is clear that building CYGNUS at the sufficient scale to reach below the neutrino floor is a challenge. Nevertheless, it is worth noting that a 1000 m³ vessel represents approximately 1/50th of the *cryogenic* internal vessel volume proposed for DUNE [199] at Sandford Underground Laboratory, for which rock excavation is now underway.

VII. SUMMARY AND RECOMMENDATIONS

We have outlined the physics case, technology choices, and design of the first large-scale directional nuclear recoil observatory. We name the project CYGNUS after the constellation from which the wind of dark matter originates.

While previous work on the subject of directional detection (summarized in Ref. [14]) was often based on idealized theoretical concepts, here we have calculated the sensitivity of a directional detector, incorporating experimental realities. These include energy-dependent performance parameters such as detection efficiency, energy resolution, head/tail recognition, and angular resolution; as well as background considerations such as radioactive impurities in the vessel, readout and environment, and the ability to discriminate between nuclear and electron recoils. This article is therefore the first to bridge the gap between the theoretical and experimental literature. Our principal conclusion is that, despite some remaining technical challenges, a CYGNUS nuclear recoil observatory is indeed feasible in line with the sensitivities displayed in Fig. 3.

A ton-scale ‘CYGNUS-1000’ detector, with a fiducial target volume of 1000 m³, filled with a He:SF₆ gas mixture at room temperature and atmospheric pressure, and with 1–3 keV_r event detection thresholds, would have a non-directional sensitivity to WIMP-nucleon cross sections extending significantly beyond existing limits. For SI interactions this sensitivity could extend into presently unexplored sub-10 GeV c^{-2} parameter space, whereas for SD interactions, even a 10 m³-scale experiment would be sufficient to compete with generation-two (G2) detectors currently under construction, such as LZ and SuperCDMS.

As we discussed in Sec. II, the physics motivation for

CYGNUS is substantial, prior to, or even in the absence of, a positive detection of DM. Even for the worst-case scenario of an 8 keV_r nuclear recoil threshold, CYGNUS-1000 would observe around 13 CE ν NS events over six years from ⁸B and *hep* solar neutrinos. This would be a significant achievement, given that CE ν NS will not become an appreciable signal in conventional direct detection experiments until LZ or XenonNT have taken data. For a threshold of 1 keV_r this number increases to 37 which would already be enough to begin to characterize the neutrino spectrum.

In addition to solar neutrinos, supernovae occurring at distances closer than 3 kpc would be sufficient to produce a measurable number of highly energetic nuclear recoil events in CYGNUS-1000. For even larger volumes, the possibilities open up for the detection of geological and reactor neutrinos. Another interesting signal, but one that is not explored in detail here, are the signals from neutrino electron scattering events (see Fig. 4). A large number of high energy electron recoils would be expected in CYGNUS-1000, primarily from the lower energy components of the flux like *pp* or ⁷Be neutrinos. CYGNUS could have a lower threshold for electron events than even Borexino, so the prospects to contribute to the physics of solar neutrinos may even extend beyond what we have discussed here. Just as the excellent nuclear/electron recoil discrimination permits the identification of a WIMP signal, this same capability also implies excellent potential background rejection if solar neutrino-electron recoils are the signal of interest. An investigation of both solar and geoneutrino-electron recoils in directional experiments is therefore worthy of a detailed follow-up study.

We expect CYGNUS to be capable of full 3d event reconstruction, good directional sensitivity and excellent electron/nuclear recoil discrimination even at energies well below 10 keV_r. This will require a highly segmented charge readout and the limiting of the diffusion of ionization in the target gas. High electron rejection will enable the detector to operate free of internal background, with the capability of distinguishing WIMP-nucleus scattering from CE ν NS or unexpected nuclear recoil backgrounds. For example, with a high-resolution pixel or strip readout, on average 4-5 detected 100-GeV c^{-2} WIMP-fluorine recoils above 50 keV_r are sufficient to rule out an isotropic recoil distribution at 90% C.L. For a 10-GeV c^{-2} WIMP, less than 20 detected helium-recoil events above 6 keV_r or 3-4 events above 20 keV_r would suffice to rule out a neutrino-induced nuclear recoil distribution at 90% C.L. This would constitute the first step in firmly establishing the galactic origin of a tentative dark matter signal.

We estimate that the most cost-effective way to achieve the desired high readout segmentation, low diffusion, and good directional performance at low recoil energies, is via strip readout technologies and negative ion drift. CYGNUS-1000 would then require 2000 m² of strip readout planes. Thanks to LHC technology development, large strip micromegas planes from CERN that would meet our technical requirements are already available at

a cost of order \$12,500 /m². If a radiopure version of these detector as well as preamps with integration time appropriate for NID are developed, then CYGNUS-1000 could be constructed relatively soon and at quite reasonable cost. Assuming 20 million readout channels at an electronics cost of US \$1/channel for mass production, the total charge readout cost of CYGNUS-1000 would then amount to US \$45 million. Downstream DAQ, gas vessels and shielding would add to the cost, but due to the ability of CYGNUS to operate with low noise at room temperature and atmospheric pressure, these costs will be kept reasonable.

We do not anticipate CYGNUS to be a monolithic experiment like most detectors currently in operation. The 1000 m³ target volume would be best achieved using multiple smaller detectors (see Fig. 1), say of $5 \times 5 \times 10$ m³, potentially located at multiple underground sites (see Fig. 2). As well as simply facilitating our the sensitivity goals while maintaining low pressure operation, the modularity and distribution of the experiment would allow for the control of systematics by comparison between detectors, and importantly, leaves room for future expansion at each site. Utilizing multiple detectors would allow also the use of multiple target gases and pressures, but the currently favored scenario is to include a negative-ion drift gas, preferably SF₆ mixed with helium, to arrive at atmospheric total pressure. This mixture has multiple advantages: it improves the directionality of all recoil species, permits *z*-fiducialization via minority carriers, and extends the sensitivity to both low WIMP masses and neutrinos. Atmospheric pressure also avoids the need for a vacuum vessel.

While a low-density gas such as 755:5 torr He:SF₆ is essential to maintain low-mass WIMP and neutrino sensitivity with directionality, the planned segmentation of CYGNUS naturally enables operation of parts of the detector with a higher-density "search mode" gas. If the final CYGNUS design does not require a vacuum vessel, then for example half the fiducial volume could run with 760 of SF₆ gas, which would boost the exposure by a factor of 150. If we choose a vacuum-capable gas vessel design, then this would be capable of withstanding a 1 atmosphere pressure differential. In that case the search mode could utilize 1520 torr of SF₆ for a factor 300 boost in exposure. The beauty of CYGNUS is that the exact partitioning of the target volume into low-density and search mode running can be optimized and varied even after construction, and be responsive to new developments in the field. This flexibility may prove particularly important for larger volume detectors, *e.g.* a CYGNUS-100k, which will required a substantial investment of time and funding.

The new simulations presented here, validated in part by experiment, suggest that electron rejection is feasible down to 1 keV_{ee} in the atmospheric pressure He:SF₆ mixture. The discrimination power depends on the sophistication of the readout but rises exponentially with energy. With a simple discriminant the expected electron

rejection factor to exceed 10⁶ at 5 keV_{ee} for fluorine, and 10 keV_{ee} for helium. Our first investigations with deep learning neutral networks suggest this can be improved upon by several orders of magnitude[319]. Machine learning efforts by the MIMAC collaboration [350] similarly suggest great improvements are possible. The optimization of the electron rejection, and finding the optimal tradeoff of intrinsic background versus the sophistication of readout is the subject of future work.

To limit the background from the CYGNUS vessel, we envisage a skeletal design based on titanium, copper and acrylic, surrounded by an inner layered structure of gamma shielding, incorporating copper and steel. Shielding from neutron backgrounds can be achieved by using water blocks outside the inner shield (see Fig. 32), and a conventional muon veto depending on site depth. Some improvement in copper radiopurity will be needed for the shielding but a back-up to this would be to use entirely water shielding.

Final recommendations

To clarify the route forward, we have five main recommendations: First, the strawman design discussed here should be advanced to a properly optimized technical design. We expect that the low-mass reach of a fully optimized strip readout detector will exceed the physics reach estimated here, as we found that our strawman design is diffusion-limited. Second, all energy-dependent performance metrics introduced here, such as the recoil efficiency, angular resolution, head-tail efficiency, and electron rejection, must be demonstrated experimentally in a small prototype with full drift length and high readout resolution: a 'CYGNUS HD1 Demonstrator'. Third, the intrinsic radioactivities of the components of the strip readout must be reduced to where they are consistent with the measured electron rejection capabilities. Fourth, inexpensive, scaleable readout electronics for large-area strip readout planes should be designed. Lastly, the expected physics sensitivity of CYGNUS should be more broadly investigated. For example, we expect that the 3d ionization measurements should be uniquely capable of detecting and identifying non-conventional DM-nucleon scattering final states; and as mentioned above, the prospects for novel physics based on *electron* recoils should also be investigated, such as in the context of neutrinos or axion-like particles.

Independently of this main development path, we also recommend the pursuit of alternative design approaches, based on electron drift. Electron drift gases will allow much higher avalanche gains than NID gases, but at the cost of increased diffusion. For electron drift, fiducialization would be performed via measurements of diffusion, rather than via observation of minority carriers. The higher gain could be a good match for optical readout, and hence the CYGNO experiment in Italy [278] will pursue that option. Another interesting configuration is a

pixel or strip readout detector with electron drift and low threshold. This is being pursued by US CYGNUS collaborators. We leave the evaluation of these options for future work.

VIII. CONCLUSION

An exciting physics program will be possible with the anticipated network of CYGNUS detectors, as illustrated in Fig. 34. The initial 1 m^3 CYGNUS HD1 demonstrator, or a slight scale-up, could be used to demonstrate directional sensitivity to $\text{CE}\nu\text{NS}$ with reactor or spallation sources. This would guarantee sensitivity to a directional solar neutrino signal in the subsequent detectors. Then, for every factor ten increase in exposure, interesting new measurements are possible. CYGNUS-1000 would detect between 13 and 37 $\text{CE}\nu\text{NS}$ events over six years, depending on the exact energy threshold for background free operation. An ambitious CYGNUS-100k detector, with volume similar to that of DUNE [199], would have non-directional WIMP sensitivity in excess of any proposed experiment, and would, in addition, allow us to utilize directionality to penetrate deep into the neutrino floor. Finally, if a dark matter signal is observed, this would mark the beginning of a new era in physics. The community would seek to firmly establish the galactic origin of the signal, which is only possible by observing some form of direction-dependence. After that, a worldwide effort would ensure to map the local velocity distribution

and explore the particle phenomenology of dark matter. A large directional detector such as CYGNUS-100k would be essential to enable this.

ACKNOWLEDGMENTS

SEV acknowledges support from the U.S. Department of Energy (DOE) via Award Numbers DE-SC0018034 and DE-SC0010504. CAJO is supported by the University of Sydney and the Australian Research Council. EB is supported by the European Research Council (ERC) under the European Union Horizon 2020 programme (grant agreement No. 818744) KS is supported by the U.S. Department of Energy and the National Science Foundation. WL, NJCS, CE, ACE and FMM would like to thank STFC for continued support under grant ST/S000747/1 KM is supported by the Japanese Ministry of Education, Culture, Sports, Science and Technology, a Grant-in-Aid for Scientific Research, ICRP Joint-Usage, Japan Society for the Promotion of Science (JSPS) KAKENHI Grant Numbers 16H02189, 26104005, 26104009, 19H05806 and the JSPS Bilateral Collaborations (Joint Research Projects and Seminars) program and Program for Advancing Strategic International Networks to Accelerate the Circulation of Talented Researchers, JSPS, Japan (R2607). Computing resources were provided by the University of Chicago Research Computing Center.

-
- [1] G. Bertone and D. Hooper, *Rev. Mod. Phys.* **90**, 045002 (2018), [arXiv:1605.04909 \[astro-ph.CO\]](#).
 - [2] M. Battaglieri *et al.*, in *U.S. Cosmic Visions: New Ideas in Dark Matter College Park, MD, USA, March 23-25, 2017* (2017) [arXiv:1707.04591 \[hep-ph\]](#).
 - [3] G. Angloher *et al.* (CRESST), *Eur. Phys. J. C* **76**, 25 (2016), [arXiv:1509.01515 \[astro-ph.CO\]](#).
 - [4] R. Agnese *et al.* (SuperCDMS), *Phys. Rev. D* **97**, 022002 (2018), [arXiv:1707.01632 \[astro-ph.CO\]](#).
 - [5] L. Hehn *et al.* (EDELWEISS), *Eur. Phys. J. C* **76**, 548 (2016), [arXiv:1607.03367 \[astro-ph.CO\]](#).
 - [6] C. Amole *et al.* (PICO), *Phys. Rev. Lett.* **114**, 231302 (2015), [arXiv:1503.00008 \[astro-ph.CO\]](#).
 - [7] C. Amole *et al.* (PICO), *Phys. Rev. D* **93**, 052014 (2016), [arXiv:1510.07754 \[hep-ex\]](#).
 - [8] P. Agnes *et al.* (DarkSide), *Phys. Rev. D* **93**, 081101 (2016), [arXiv:1510.00702 \[astro-ph.CO\]](#).
 - [9] A. Tan *et al.* (PandaX-II), *Phys. Rev. Lett.* **117**, 121303 (2016), [arXiv:1607.07400 \[hep-ex\]](#).
 - [10] D. S. Akerib *et al.* (LUX), *Phys. Rev. Lett.* **118**, 021303 (2017), [arXiv:1608.07648 \[astro-ph.CO\]](#).
 - [11] E. Aprile *et al.* (XENON), *Phys. Rev. Lett.* **119**, 181301 (2017), [arXiv:1705.06655 \[astro-ph.CO\]](#).
 - [12] G. Adhikari *et al.* (COSINE-100), *Phys. Rev. Lett.* **123**, 031302 (2019), [arXiv:1903.10098 \[astro-ph.IM\]](#).
 - [13] S. Ahlen *et al.*, *Int. J. Mod. Phys. A* **25**, 1 (2010), [arXiv:0911.0323 \[astro-ph.CO\]](#).
 - [14] F. Mayet *et al.*, *Phys. Rept.* **627**, 1 (2016), [arXiv:1602.03781 \[astro-ph.CO\]](#).
 - [15] J. B. R. Battat *et al.*, *Phys. Rept.* **662**, 1 (2016), [arXiv:1610.02396 \[physics.ins-det\]](#).
 - [16] J. B. R. Battat *et al.*, *J. Instrum.* **9**, P11004 (2014), [arXiv:1407.3938 \[physics.ins-det\]](#).
 - [17] D. Santos, J. Billard, G. Bosson, J. Bouly, O. Bourrion, *et al.*, *EAS Publ. Ser.* **53**, 25 (2012), [arXiv:1111.1566 \[astro-ph.IM\]](#).
 - [18] S. Vahsen, H. Feng, M. Garcia-Sciveres, I. Jaegle, J. Kadyk, *et al.*, *EAS Publ. Ser.* **53**, 43 (2012), [arXiv:1110.3401 \[astro-ph.IM\]](#).
 - [19] C. Deaconu, in *UCLA 11th Symposium on Sources and Detection of Dark Matter and Dark Energy in the Universe* (2014).
 - [20] K. Miuchi *et al.*, *Phys. Lett. B* **686**, 11 (2010), [arXiv:1002.1794 \[astro-ph.CO\]](#).
 - [21] T. Marrodán Undagoitia and L. Rauch, *J. Phys. G* **43**, 013001 (2016), [arXiv:1509.08767 \[physics.ins-det\]](#).
 - [22] M. Schumann, *J. Phys. G* **46**, 103003 (2019), [arXiv:1903.03026 \[astro-ph.CO\]](#).
 - [23] G. Arcadi, M. Dutra, P. Ghosh, M. Lindner, Y. Mambrini, M. Pierre, S. Profumo, and F. S. Queiroz, *Eur. Phys. J. C* **78**, 203 (2018), [arXiv:1703.07364 \[hep-ph\]](#).
 - [24] R. Bernabei *et al.*, *Eur. Phys. J. C* **73**, 2648 (2013), [arXiv:1308.5109 \[astro-ph.GA\]](#).

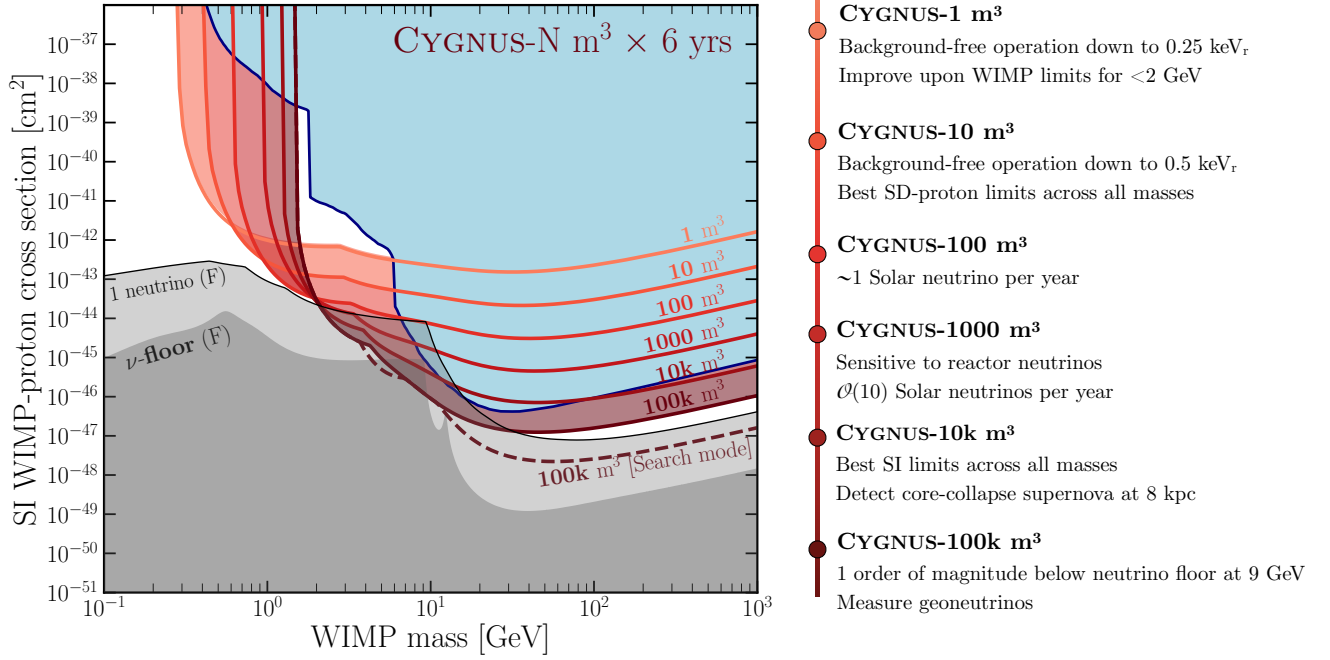


FIG. 34. Summary of the projected SI WIMP 90% CL exclusion limits as a function of the total fiducial volume of the network of detectors comprising the CYGNUS experiment. All experimental exposures are multiplied by a running time of 6 years such that the listed benchmarks are all approximately multiples of 1 ton year (1 ton year corresponds to 1000 m³). We shade in blue the currently excluded limits on the SI WIMP-nucleon cross section as of 2020. The thresholds are increased evenly between 0.25 and 8 keV_r and for each increasing volume to illustrate the range of possible thresholds envisaged for the final experiment. Below the final volume we also extend the reach by including the possibility of a “search mode” experiment which would have 200 Torr of SF₆ (as opposed to 5 Torr), but would have no directionality sensitivity.

- [25] R. Bernabei *et al.*, *Proceedings, 7th International Conference on New Frontiers in Physics (ICNFP 2018): Kolymbari, Crete, Greece, July 4-12, 2018*, *Universe* **4**, 116 (2018), arXiv:1805.10486 [hep-ex].
- [26] E. Barbosa de Souza *et al.* (DM-Ice), *Phys. Rev. D* **95**, 032006 (2017), arXiv:1602.05939 [physics.ins-det].
- [27] K. W. Kim *et al.* (KIMS), *JHEP* **03**, 194 (2019), arXiv:1806.06499 [hep-ex].
- [28] F. Froberg (SABRE), *Proceedings, 14th International Conference on Topics in Astroparticle and Underground Physics (TAUP 2015): Torino, Italy, September 7-11, 2015*, *J. Phys. Conf. Ser.* **718**, 042021 (2016), arXiv:1601.05307 [physics.ins-det].
- [29] G. Adhikari *et al.*, *Nature* **564**, 83 (2018).
- [30] J. Amaré *et al.*, *Phys. Rev. Lett.* **123**, 031301 (2019), arXiv:1903.03973 [astro-ph.IM].
- [31] J. Amaré *et al.*, *Eur. Phys. J. C* **79**, 228 (2019), arXiv:1812.01472 [astro-ph.IM].
- [32] G. Angloher *et al.* (COSINUS), *JINST* **12**, P11007 (2017), arXiv:1705.11028 [physics.ins-det].
- [33] G. Angloher *et al.*, *Eur. Phys. J. C* **76**, 441 (2016), arXiv:1603.02214 [physics.ins-det].
- [34] D. N. Spergel, *Phys. Rev. D* **37**, 1353 (1988).
- [35] P. Grothaus, M. Fairbairn, and J. Monroe, *Phys. Rev. D* **90**, 055018 (2014), arXiv:1406.5047 [hep-ph].
- [36] C. A. J. O’Hare, A. M. Green, J. Billard, E. Figueroa-Feliciano, and L. E. Strigari, *Phys. Rev. D* **92**, 063518 (2015), arXiv:1505.08061 [astro-ph.CO].
- [37] C. A. J. O’Hare, B. J. Kavanagh, and A. M. Green, *Phys. Rev. D* **96**, 083011 (2017), arXiv:1708.02959 [astro-ph.CO].
- [38] J. Billard, F. Mayet, and D. Santos, *Phys. Rev. D* **85**, 035006 (2012), arXiv:1110.6079 [astro-ph.CO].
- [39] G. Arcadi, M. Lindner, F. S. Queiroz, W. Rodejohann, and S. Vogl, *JCAP* **1803**, 042 (2018), arXiv:1711.02110 [hep-ph].
- [40] N. Linzer and K. Scholberg, *Phys. Rev. D* **100**, 103005 (2019), arXiv:1909.03151 [astro-ph.IM].
- [41] M. Leyton, S. Dye, and J. Monroe, *Nature Commun.* **8**, 15989 (2017), arXiv:1710.06724 [physics.geo-ph].
- [42] C. A. J. O’Hare, (2020), arXiv:2002.07499 [astro-ph.CO].
- [43] B. Morgan, A. M. Green, and N. J. C. Spooner, *Phys. Rev. D* **71**, 103507 (2005), arXiv:astro-ph/0408047 [astro-ph].
- [44] J. Billard, F. Mayet, J. F. Macias-Perez, and D. Santos, *Phys. Lett. B* **691**, 156 (2010), arXiv:0911.4086 [astro-ph.CO].
- [45] S. K. Lee and A. H. G. Peter, *J. Cosmol. Astropart. Phys.* **1204**, 029 (2012), arXiv:1202.5035 [astro-ph.CO].
- [46] C. A. J. O’Hare and A. M. Green, *Phys. Rev. D* **90**, 123511 (2014), arXiv:1410.2749 [astro-ph.CO].
- [47] B. J. Kavanagh and C. A. J. O’Hare, *Phys. Rev. D* **94**, 123009 (2016), arXiv:1609.08630 [astro-ph.CO].
- [48] P. Gondolo, *Phys. Rev. D* **66**, 103513 (2002), hep-ph/0209110.

- [49] S. Deans, *The Radon transform and some of its applications* (John Wiley and Sons, 1983).
- [50] A. M. Green and B. Morgan, *Phys. Rev. D* **81**, 061301 (2010), arXiv:1002.2717 [astro-ph.CO].
- [51] N. Bozorgnia, G. B. Gelmini, and P. Gondolo, *J. Cosmol. Astropart. Phys.* **1206**, 037 (2012), arXiv:1111.6361 [astro-ph.CO].
- [52] N. Bozorgnia, G. B. Gelmini, and P. Gondolo, *J. Cosmol. Astropart. Phys.* **1208**, 011 (2012), arXiv:1205.2333 [astro-ph.CO].
- [53] P. Agnes *et al.* (DarkSide), *Phys. Rev. Lett.* **121**, 111303 (2018), arXiv:1802.06998 [astro-ph.CO].
- [54] S. Archambault *et al.* (PICASSO), *Phys. Lett. B* **711**, 153 (2012), arXiv:1202.1240 [hep-ex].
- [55] S. C. Kim *et al.*, *Phys. Rev. Lett.* **108**, 181301 (2012), arXiv:1204.2646 [astro-ph.CO].
- [56] E. Behnke *et al.* (COUPP), *Science* **319**, 933 (2008), arXiv:0804.2886 [astro-ph].
- [57] K. Choi *et al.* (Super-Kamiokande), *Phys. Rev. Lett.* **114**, 141301 (2015), arXiv:1503.04858 [hep-ex].
- [58] M. G. Aartsen *et al.* (IceCube), *Eur. Phys. J. C* **77**, 146 (2017), arXiv:1612.05949 [astro-ph.HE].
- [59] B. Cabrera, L. M. Krauss, and F. Wilczek, *Phys. Rev. Lett.* **55**, 25 (1985).
- [60] J. Monroe and P. Fisher, *Phys. Rev. D* **76**, 033007 (2007), arXiv:0706.3019 [astro-ph].
- [61] L. E. Strigari, *New J. Phys.* **11**, 105011 (2009), arXiv:0903.3630 [astro-ph.CO].
- [62] A. Gutlein *et al.*, *Astropart. Phys.* **34**, 90 (2010), arXiv:1003.5530 [hep-ph].
- [63] J. Billard, L. Strigari, and E. Figueroa-Feliciano, *Phys. Rev. D* **89**, 023524 (2014), arXiv:1307.5458 [hep-ph].
- [64] R. Essig, M. Sholapurkar, and T.-T. Yu, *Phys. Rev. D* **97**, 095029 (2018), arXiv:1801.10159 [hep-ph].
- [65] J. F. Beacom, *Ann. Rev. Nucl. Part. Sci.* **60**, 439 (2010), arXiv:1004.3311 [astro-ph.HE].
- [66] G. Battistoni, A. Ferrari, T. Montaruli, and P. R. Sala, *Astropart. Phys.* **23**, 526 (2005).
- [67] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa, *Phys. Rev. D* **83**, 123001 (2011), arXiv:1102.2688 [astro-ph.HE].
- [68] C. A. J. O'Hare, *Phys. Rev. D* **94**, 063527 (2016), arXiv:1604.03858 [astro-ph.CO].
- [69] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari, *Phys. Rev. D* **90**, 083510 (2014), arXiv:1408.3581 [hep-ph].
- [70] J. H. Davis, *JCAP* **1503**, 012 (2015), arXiv:1412.1475 [hep-ph].
- [71] N. Agafonova *et al.* (NEWSdm), *Eur. Phys. J. C* **78**, 578 (2018), arXiv:1705.00613 [astro-ph.CO].
- [72] T. Franarin and M. Fairbairn, *Phys. Rev. D* **94**, 053004 (2016), arXiv:1605.08727 [hep-ph].
- [73] F. Kadribasic, N. Mirabolfathi, K. Nordlund, A. E. Sand, E. Holmström, and F. Djurabekova, *Phys. Rev. Lett.* **120**, 111301 (2018), arXiv:1703.05371 [physics.ins-det].
- [74] J. Bovy and S. Tremaine, *Astrophys. J.* **756**, 89 (2012), arXiv:1205.4033 [astro-ph.GA].
- [75] M. C. Smith, S. H. Whiteoak, and N. W. Evans, *Astrophys. J.* **746**, 181 (2012), arXiv:1111.6920 [astro-ph.GA].
- [76] O. Bienaymé *et al.*, *Astron. Astrophys.* **571**, A92 (2014), arXiv:1406.6896.
- [77] T. Piffl *et al.*, *Mon. Not. Roy. Astron. Soc.* **445**, 3133 (2014), arXiv:1406.4130.
- [78] S. Sivertsson, H. Silverwood, J. I. Read, G. Bertone, and P. Steger, *Mon. Not. Roy. Astron. Soc.* **478**, 1677 (2018), arXiv:1708.07836 [astro-ph.GA].
- [79] J. I. Read, *J. Phys. G* **41**, 063101 (2014), arXiv:1404.1938 [astro-ph.GA].
- [80] A. G. A. Brown *et al.* (Gaia), *Astron. Astrophys.* **616**, A1 (2018), arXiv:1804.09365 [astro-ph.GA].
- [81] B. J. Kavanagh, T. Emken, and R. Catena, (2020), arXiv:2004.01621 [hep-ph].
- [82] N. W. Evans, C. A. J. O'Hare, and C. McCabe, *Phys. Rev. D* **99**, 023012 (2019), arXiv:1810.11468 [astro-ph.GA].
- [83] T. Piffl *et al.*, *Astron. Astrophys.* **562**, A91 (2014), arXiv:1309.4293 [astro-ph.GA].
- [84] A. J. Deason, A. Fattahi, V. Belokurov, N. W. Evans, R. J. J. Grand, F. Marinacci, and R. Pakmor, *Mon. Not. Roy. Astron. Soc.* **485**, 3514 (2019), arXiv:1901.02016 [astro-ph.GA].
- [85] A. H. G. Peter, *Phys. Rev. D* **83**, 125029 (2011), arXiv:1103.5145 [astro-ph.CO].
- [86] B. J. Kavanagh and A. M. Green, *Phys. Rev. Lett.* **111**, 031302 (2013), arXiv:1303.6868 [astro-ph.CO].
- [87] C. McCabe, *Phys. Rev. D* **82**, 023530 (2010), arXiv:1005.0579 [hep-ph].
- [88] P. J. Fox, Y. Kahn, and M. McCullough, *JCAP* **1410**, 076 (2014), arXiv:1403.6830 [hep-ph].
- [89] E. Del Nobile, G. Gelmini, P. Gondolo, and J.-H. Huh, *JCAP* **1310**, 048 (2013), arXiv:1306.5273 [hep-ph].
- [90] B. Feldstein and F. Kahlhoefer, *JCAP* **1408**, 065 (2014), arXiv:1403.4606 [hep-ph].
- [91] G. B. Gelmini, J.-H. Huh, and S. J. Witte, *JCAP* **1610**, 029 (2016), arXiv:1607.02445 [hep-ph].
- [92] A. Ibarra and A. Rappelt, *JCAP* **1708**, 039 (2017), arXiv:1703.09168 [hep-ph].
- [93] A. Fowlie, *JCAP* **1710**, 002 (2017), arXiv:1708.00181 [hep-ph].
- [94] G. B. Gelmini, J.-H. Huh, and S. J. Witte, *JCAP* **1712**, 039 (2017), arXiv:1707.07019 [hep-ph].
- [95] B. J. Kavanagh, *JCAP* **1507**, 019 (2015), arXiv:1502.04224 [hep-ph].
- [96] C. A. J. O'Hare and A. M. Green, *Phys. Rev. D* **90**, 123511 (2014), arXiv:1410.2749 [astro-ph.CO].
- [97] M. Kuhlen, M. Lisanti, and D. N. Spergel, *Phys. Rev. D* **86**, 063505 (2012), arXiv:1202.0007 [astro-ph.GA].
- [98] G. Besla, A. Peter, and N. Garavito-Camargo, *JCAP* **11**, 013 (2019), arXiv:1909.04140 [astro-ph.GA].
- [99] J. M. D. Kruijssen, J. L. Pfeffer, M. Reina-Campos, R. A. Crain, and N. Bastian, *Mon. Not. R. Astron. Soc.* **486**, 3180 (2019), arXiv:1806.05680 [astro-ph.GA].
- [100] G. C. Myeong, N. W. Evans, V. Belokurov, J. L. Sanders, and S. E. Koposov, *The Astrophysical Journal Letters* **856**, L26 (2018), arXiv:1802.03351.
- [101] V. Belokurov, D. Erkal, N. W. Evans, S. E. Koposov, and A. J. Deason, *Mon. Not. Roy. Astron. Soc.* **478**, 611 (2018), arXiv:1802.03414.
- [102] J. T. Mackereth *et al.*, *Mon. Not. Roy. Astron. Soc.* **482**, 3426 (2019), arXiv:1808.00968 [astro-ph.GA].
- [103] N. Bozorgnia, A. Fattahi, C. S. Frenk, A. Cheek, D. G. Cerdeno, F. A. Gómez, R. J. J. Grand, and F. Marinacci, (2019), arXiv:1910.07536 [astro-ph.GA].
- [104] J. T. Mackereth and J. Bovy, (2019), arXiv:1910.03590 [astro-ph.GA].
- [105] Z. Yuan *et al.*, (2019), arXiv:1910.07538 [astro-ph.GA].

- [106] R. P. Naidu, C. Conroy, A. Bonaca, B. D. Johnson, Y.-S. Ting, N. Caldwell, D. Zaritsky, and P. A. Cargile, (2020), [Arxiv:2006.08625v1](#).
- [107] G. C. Myeong, N. W. Evans, V. Belokurov, J. L. Sanders, and S. E. Kopusov, *Mon. Not. Roy. Astron. Soc.* **478**, 5449 (2018).
- [108] G. C. Myeong, N. W. Evans, V. Belokurov, N. C. Amorisco, and S. Kopusov, *Mon. Not. Roy. Astron. Soc.* **475**, 1537 (2018), [arXiv:1712.04071 \[astro-ph.GA\]](#).
- [109] C. A. J. O'Hare, C. McCabe, N. W. Evans, G. Myeong, and V. Belokurov, *Phys. Rev. D* **98**, 103006 (2018), [arXiv:1807.09004 \[astro-ph.CO\]](#).
- [110] C. A. O'Hare, N. W. Evans, C. McCabe, G. Myeong, and V. Belokurov, *Phys. Rev. D* **101**, 023006 (2020), [arXiv:1909.04684 \[astro-ph.GA\]](#).
- [111] D. S. Aguado *et al.*, (2020), [arXiv:2007.11003 \[astro-ph.GA\]](#).
- [112] P. Adhikari *et al.* (DEAP), (2020), [arXiv:2005.14667 \[astro-ph.CO\]](#).
- [113] D. Tucker-Smith and N. Weiner, *Phys. Rev. D* **64**, 043502 (2001), [arXiv:hep-ph/0101138 \[hep-ph\]](#).
- [114] D. P. Finkbeiner, T. Lin, and N. Weiner, *Phys. Rev. D* **80**, 115008 (2009), [arXiv:0906.0002 \[astro-ph.CO\]](#).
- [115] M. Lisanti and J. G. Wacker, *Phys. Rev. D* **81**, 096005 (2010), [arXiv:0911.1997 \[hep-ph\]](#).
- [116] J. Eby, P. J. Fox, R. Harnik, and G. D. Kribs, *JHEP* **09**, 115 (2019), [arXiv:1904.09994 \[hep-ph\]](#).
- [117] J. Bramante, B. Broerman, R. F. Lang, and N. Raj, *Phys. Rev. D* **98**, 083516 (2018), [arXiv:1803.08044 \[hep-ph\]](#).
- [118] R. Laha, *Phys. Rev. D* **92**, 083509 (2015), [arXiv:1505.02772 \[hep-ph\]](#).
- [119] J. Fan, M. Reece, and L.-T. Wang, *JCAP* **1011**, 042 (2010), [arXiv:1008.1591 \[hep-ph\]](#).
- [120] B. J. Kavanagh, *Phys. Rev. D* **92**, 023513 (2015), [arXiv:1505.07406 \[hep-ph\]](#).
- [121] R. Catena, *JCAP* **1507**, 026 (2015), [arXiv:1505.06441 \[hep-ph\]](#).
- [122] R. Catena, J. Conrad, C. Döring, A. D. Ferella, and M. B. Krauss, *Phys. Rev. D* **97**, 023007 (2018), [arXiv:1706.09471 \[hep-ph\]](#).
- [123] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [124] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, (2020), [arXiv:2003.01100 \[hep-ph\]](#).
- [125] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *Phys. Rev. D* **81**, 123530 (2010), [arXiv:0905.4720 \[hep-th\]](#).
- [126] L. F. Abbott and P. Sikivie, *Phys. Lett. B* **120**, 133 (1983).
- [127] M. Dine and W. Fischler, *Phys. Lett. B* **120**, 137 (1983).
- [128] J. Preskill, M. B. Wise, and F. Wilczek, *Phys. Lett. B* **120**, 127 (1983).
- [129] S. Chang, C. Hagmann, and P. Sikivie, *Phys. Rev. D* **59**, 023505 (1999), [arXiv:hep-ph/9807374 \[hep-ph\]](#).
- [130] M. Kawasaki, K. Saikawa, and T. Sekiguchi, *Phys. Rev. D* **91**, 065014 (2015), [arXiv:1412.0789 \[hep-ph\]](#).
- [131] A. R. Liddle and M. S. Madsen, *Int. J. Mod. Phys. D* **1**, 101 (1992).
- [132] C. J. Hogan and M. J. Rees, *Phys. Lett. B* **205**, 228 (1988).
- [133] E. W. Kolb and I. I. Tkachev, *Phys. Rev. D* **50**, 769 (1994), [arXiv:astro-ph/9403011 \[astro-ph\]](#).
- [134] A. Vaquero, J. Redondo, and J. Stadler, *JCAP* **1904**, 012 (2018), [arXiv:1809.09241 \[astro-ph.CO\]](#).
- [135] D. J. E. Marsh, *Phys. Rept.* **643**, 1 (2016), [arXiv:1510.07633 \[astro-ph.CO\]](#).
- [136] I. G. Irastorza and J. Redondo, *Prog. Part. Nucl. Phys.* **102**, 89 (2018), [arXiv:1801.08127 \[hep-ph\]](#).
- [137] D. S. Akerib *et al.* (LUX), *Phys. Rev. Lett.* **118**, 261301 (2017), [arXiv:1704.02297 \[astro-ph.CO\]](#).
- [138] E. Aprile *et al.* (XENON100), *Phys. Rev. D* **90**, 062009 (2014), [arXiv:1404.1455 \[astro-ph.CO\]](#).
- [139] E. Aprile *et al.* (XENON), *Phys. Rev. Lett.* **123**, 251801 (2019), [arXiv:1907.11485 \[hep-ex\]](#).
- [140] E. Aprile *et al.* (XENON), (2020), [arXiv:2006.09721 \[hep-ex\]](#).
- [141] E. Armengaud *et al.*, *JCAP* **1311**, 067 (2013), [arXiv:1307.1488 \[astro-ph.CO\]](#).
- [142] Z. Ahmed *et al.* (CDMS), *Phys. Rev. Lett.* **103**, 141802 (2009), [arXiv:0902.4693 \[hep-ex\]](#).
- [143] C. Fu *et al.* (PandaX), *Phys. Rev. Lett.* **119**, 181806 (2017), [arXiv:1707.07921 \[hep-ex\]](#).
- [144] A. Derevianko, V. A. Dzuba, V. V. Flambaum, and M. Pospelov, *Phys. Rev. D* **82**, 065006 (2010), [arXiv:1007.1833 \[hep-ph\]](#).
- [145] A. V. Derbin, I. S. Drachnev, A. S. Kayunov, and V. N. Muratova, *JETP Lett.* **95**, 339 (2012), [*Pisma Zh. Eksp. Teor. Fiz.*95,379(2012)], [arXiv:1206.4142 \[hep-ex\]](#).
- [146] A. Ljubicic, D. Kekez, Z. Krecak, and T. Ljubicic, *Phys. Lett. B* **599**, 143 (2004), [arXiv:hep-ex/0403045 \[hep-ex\]](#).
- [147] S. J. Asztalos *et al.* (ADMX), *Phys. Rev. Lett.* **104**, 041301 (2010), [arXiv:0910.5914 \[astro-ph.CO\]](#).
- [148] K. Zioutas *et al.* (CAST), *Phys. Rev. Lett.* **94**, 121301 (2005), [arXiv:hep-ex/0411033](#).
- [149] H. An, M. Pospelov, J. Pradler, and A. Ritz, *Phys. Lett. B* **747**, 331 (2015), [arXiv:1412.8378 \[hep-ph\]](#).
- [150] J. Redondo, *JCAP* **1312**, 008 (2013), [arXiv:1310.0823 \[hep-ph\]](#).
- [151] F. Capozzi and G. Raffelt, (2020), [arXiv:2007.03694 \[astro-ph.SR\]](#).
- [152] P. Athron *et al.*, (2020), [arXiv:2007.05517 \[astro-ph.CO\]](#).
- [153] D. Croon, S. D. McDermott, and J. Sakstein, (2020), [arXiv:2007.00650 \[hep-ph\]](#).
- [154] A. Garcon *et al.*, (2019), [10.1126/sciadv.aax4539, arXiv:1902.04644 \[hep-ex\]](#).
- [155] S. Knirck, A. J. Millar, C. A. J. O'Hare, J. Redondo, and F. D. Steffen, *JCAP* **1811**, 051 (2018), [arXiv:1806.05927 \[astro-ph.CO\]](#).
- [156] D. Z. Freedman, *Phys. Rev. D* **9**, 1389 (1974).
- [157] D. Akimov *et al.* (COHERENT), *Science* **357**, 1123 (2017), [arXiv:1708.01294 \[nucl-ex\]](#).
- [158] D. Akimov *et al.* (COHERENT), (2020), [10.5281/zenodo.3903810, arXiv:2006.12659 \[nucl-ex\]](#).
- [159] D. Akimov *et al.* (COHERENT), (2020), [arXiv:2003.10630 \[nucl-ex\]](#).
- [160] A. Drukier and L. Stodolsky, *Phys. Rev. D* **30**, 2295 (1984), [[395\(1984\)](#)].
- [161] E. Vitagliano, I. Tamborra, and G. Raffelt, (2019), [arXiv:1910.11878 \[astro-ph.HE\]](#).
- [162] J. Monroe and the DM-TPC Collaboration, *Journal of Physics: Conference Series* **136**, 022037 (2008).
- [163] W. C. Haxton, R. G. Hamish Robertson, and A. M. Serenelli, *Ann. Rev. Astron. Astrophys.* **51**, 21 (2013), [arXiv:1208.5723 \[astro-ph.SR\]](#).
- [164] G. Bellini *et al.* (BOREXINO), *Nature* **512**, 383 (2014).

- [165] L. E. Strigari, *Phys. Rev. D* **93**, 103534 (2016), [arXiv:1604.00729 \[astro-ph.CO\]](#).
- [166] N. Grevesse and A. J. Sauval, *Space Sci. Rev.* **85**, 161 (1998).
- [167] F. L. Villante, A. M. Serenelli, F. Delahaye, and M. H. Pinsonneault, *Astrophys. J.* **787**, 13 (2014), [arXiv:1312.3885 \[astro-ph.SR\]](#).
- [168] N. Vinyoles, A. M. Serenelli, F. L. Villante, S. Basu, J. Bergström, M. C. Gonzalez-Garcia, M. Maltoni, C. Peña-Garay, and N. Song, *Astrophys. J.* **835**, 202 (2017), [arXiv:1611.09867 \[astro-ph.SR\]](#).
- [169] J. Billard, L. E. Strigari, and E. Figueroa-Feliciano, *Phys. Rev. D* **91**, 095023 (2015), [arXiv:1409.0050 \[astro-ph.CO\]](#).
- [170] D. G. Cerdeño, M. Fairbairn, T. Jubb, P. A. N. Machado, A. C. Vincent, and C. Boehm, *JHEP* **05**, 118 (2016), [Erratum: *JHEP*09,048(2016)], [arXiv:1604.01025 \[hep-ph\]](#).
- [171] D. G. Cerdeno, J. H. Davis, M. Fairbairn, and A. C. Vincent, *JCAP* **1804**, 037 (2018), [arXiv:1712.06522 \[hep-ph\]](#).
- [172] M. Agostini *et al.* (BOREXINO), (2020), [arXiv:2006.15115 \[hep-ex\]](#).
- [173] J. Seguinot, T. Ypsilantis, and A. Zichichi, *Conf. Proc. C* **920310**, 289 (1992).
- [174] C. Arpesella, C. Brogini, and C. Cattadori, *Astropart. Phys.* **4**, 333 (1996).
- [175] H.-T. Janka, K. Langanke, A. Marek, G. Martinez-Pinedo, and B. Mueller, *Phys. Rept.* **442**, 38 (2007), [arXiv:astro-ph/0612072](#).
- [176] A. Mirizzi, I. Tamborra, H.-T. Janka, N. Saviano, K. Scholberg, R. Bollig, L. Hudepohl, and S. Chakraborty, *Riv. Nuovo Cim.* **39**, 1 (2016), [arXiv:1508.00785 \[astro-ph.HE\]](#).
- [177] C. J. Horowitz, K. J. Coakley, and D. N. McKinsey, *Phys. Rev. D* **68**, 023005 (2003), [arXiv:astro-ph/0302071 \[astro-ph\]](#).
- [178] A. Mirizzi, G. G. Raffelt, and P. D. Serpico, *JCAP* **0605**, 012 (2006), [arXiv:astro-ph/0604300 \[astro-ph\]](#).
- [179] Supernova data accessed via <https://www.mpa.mpg.de/ccsnarchive/index.html> Accessed: 2019-04-26.
- [180] I. Tamborra, B. Muller, L. Hudepohl, H.-T. Janka, and G. Raffelt, *Phys. Rev. D* **86**, 125031 (2012), [arXiv:1211.3920 \[astro-ph.SR\]](#).
- [181] N. Raj, V. Takhistov, and S. J. Witte, *Phys. Rev. D* **101**, 043008 (2020), [arXiv:1905.09283 \[hep-ph\]](#).
- [182] K. Scholberg, *Ann. Rev. Nucl. Part. Sci.* **62**, 81 (2012), [arXiv:1205.6003 \[astro-ph.IM\]](#).
- [183] R. F. Lang, C. McCabe, S. Reichard, M. Selvi, and I. Tamborra, *Phys. Rev. D* **94**, 103009 (2016), [arXiv:1606.09243 \[astro-ph.HE\]](#).
- [184] J. F. Beacom and P. Vogel, *Phys. Rev. D* **60**, 033007 (1999), [arXiv:astro-ph/9811350 \[astro-ph\]](#).
- [185] R. Tomas, D. Semikoz, G. G. Raffelt, M. Kachelriess, and A. S. Dighe, *Phys. Rev. D* **68**, 093013 (2003), [arXiv:hep-ph/0307050 \[hep-ph\]](#).
- [186] K. Abe *et al.* (Super-Kamiokande), *Astropart. Phys.* **81**, 39 (2016), [arXiv:1601.04778 \[astro-ph.HE\]](#).
- [187] C. Lunardini, *Astropart. Phys.* **79**, 49 (2016), [arXiv:1007.3252 \[astro-ph.CO\]](#).
- [188] J. Barranco, A. Bernal, and D. Delepine, *J. Phys. G* **45**, 055201 (2018), [arXiv:1706.03834 \[astro-ph.CO\]](#).
- [189] A. de Gouvea, I. J. Martinez Soler, Y. F. Perez-Gonzalez, and M. Sen, (2020), [arXiv:2007.13748 \[hep-ph\]](#).
- [190] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, *Phys. Rev. D* **70**, 023006 (2004), [arXiv:astro-ph/0403630 \[astro-ph\]](#).
- [191] G. Battistoni, A. Ferrari, T. Montaruli, and P. R. Sala, *Astropart. Phys.* **19**, 269 (2003), [Erratum: *Astropart. Phys.*19,291(2003)], [arXiv:hep-ph/0207035 \[hep-ph\]](#).
- [192] E. Richard *et al.* (Super-Kamiokande), *Phys. Rev. D* **94**, 052001 (2016), [arXiv:1510.08127 \[hep-ex\]](#).
- [193] M. G. Aartsen *et al.* (IceCube), *Phys. Rev. D* **91**, 122004 (2015), [arXiv:1504.03753 \[astro-ph.HE\]](#).
- [194] K. Abe *et al.* (Super-Kamiokande), *Phys. Rev. D* **91**, 052019 (2015), [arXiv:1410.2008 \[hep-ex\]](#).
- [195] M. G. Aartsen *et al.* (IceCube), *Phys. Rev. Lett.* **117**, 071801 (2016), [arXiv:1605.01990 \[hep-ex\]](#).
- [196] G. Mitsuka *et al.* (Super-Kamiokande), *Phys. Rev. D* **84**, 113008 (2011), [arXiv:1109.1889 \[hep-ex\]](#).
- [197] A. Esmaili and A. Yu. Smirnov, *JHEP* **06**, 026 (2013), [arXiv:1304.1042 \[hep-ph\]](#).
- [198] J. Salvado, O. Mena, S. Palomares-Ruiz, and N. Rius, *JHEP* **01**, 141 (2017), [arXiv:1609.03450 \[hep-ph\]](#).
- [199] B. Abi *et al.* (DUNE), (2018), [arXiv:1807.10334 \[physics.ins-det\]](#).
- [200] J. H. Davies and D. R. Davies, *Solid Earth* **1**, 5 (2010).
- [201] A. Gando *et al.* (KamLAND), *Nature Geo.* **4**, 647 (2011).
- [202] Y. Huang, V. Chubakov, F. Mantovani, R. L. Rudnick, and W. F. McDonough, *Geochemistry, Geophysics, Geosystems* **14**, 2003 (2013), [arXiv:1301.0365 \[physics.geo-ph\]](#).
- [203] T. Araki *et al.*, *Nature* **436**, 499 (2005).
- [204] G. Bellini *et al.* (Borexino), *Phys. Lett. B* **687**, 299 (2010), [arXiv:1003.0284 \[hep-ex\]](#).
- [205] S. M. Usman, G. R. Jocher, S. T. Dye, W. F. McDonough, and J. G. Learned, *Sci. Rep.* **5**, 13945 (2015), [arXiv:1509.03898 \[physics.ins-det\]](#).
- [206] A. C. Hayes and P. Vogel, *Ann. Rev. Nucl. Part. Sci.* **66**, 219 (2016), [arXiv:1605.02047 \[hep-ph\]](#).
- [207] A. Bungau *et al.*, *Phys. Rev. Lett.* **109**, 141802 (2012), [arXiv:1205.4419 \[hep-ex\]](#).
- [208] D. Akimov *et al.* (COHERENT), (2015), [arXiv:1509.08702 \[physics.ins-det\]](#).
- [209] C. Volpe, in *European strategy for future neutrino physics. Proceedings, Workshop, Geneva, Switzerland, October 1-3, 2009* (2009) [arXiv:0911.2630 \[hep-ph\]](#).
- [210] J. Barranco, O. Miranda, and T. Rashba, *J. High Energy Phys.* **2005**, 021 (2005), [hep-ph/0508299](#).
- [211] K. Scholberg, *Phys. Rev. D* **73**, 033005 (2006), [arXiv:hep-ex/0511042 \[hep-ex\]](#).
- [212] J. Barranco, O. G. Miranda, and T. I. Rashba, *Phys. Rev. D* **76**, 073008 (2007), [arXiv:hep-ph/0702175 \[hep-ph\]](#).
- [213] P. Coloma, M. C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz, *Phys. Rev. D* **96**, 115007 (2017), [arXiv:1708.02899 \[hep-ph\]](#).
- [214] J. Liao and D. Marfatia, *Phys. Lett. B* **775**, 54 (2017), [arXiv:1708.04255 \[hep-ph\]](#).
- [215] J. Papavassiliou, J. Bernabeu, and M. Passera, *PoS HEP2005*, 192 (2006), [arXiv:hep-ph/0512029 \[hep-ph\]](#).
- [216] T. S. Kosmas, O. G. Miranda, D. K. Papoulias, M. Tortola, and J. W. F. Valle, *Phys. Rev. D* **92**, 013011 (2015), [arXiv:1505.03202 \[hep-ph\]](#).

- [217] J. A. Formaggio, E. Figueroa-Feliciano, and A. J. Anderson, *Phys. Rev. D* **85**, 013009 (2012), [arXiv:1107.3512 \[hep-ph\]](#).
- [218] A. J. Anderson, J. M. Conrad, E. Figueroa-Feliciano, C. Ignarra, G. Karagiorgi, K. Scholberg, M. H. Shaevitz, and J. Spitz, *Phys. Rev. D* **86**, 013004 (2012), [arXiv:1201.3805 \[hep-ph\]](#).
- [219] K. Patton, J. Engel, G. C. McLaughlin, and N. Schunck, *Phys. Rev. C* **86**, 024612 (2012), [arXiv:1207.0693 \[nucl-th\]](#).
- [220] K. M. Patton, G. C. McLaughlin, and K. Scholberg, *Int. J. Mod. Phys. E* **22**, 1330013 (2013).
- [221] M. Cadeddu, C. Giunti, Y. F. Li, and Y. Y. Zhang, *Phys. Rev. Lett.* **120**, 072501 (2018), [arXiv:1710.02730 \[hep-ph\]](#).
- [222] P. deNiverville, M. Pospelov, and A. Ritz, *Phys. Rev. D* **84**, 075020 (2011), [arXiv:1107.4580 \[hep-ph\]](#).
- [223] P. deNiverville, M. Pospelov, and A. Ritz, *Phys. Rev. D* **92**, 095005 (2015), [arXiv:1505.07805 \[hep-ph\]](#).
- [224] S.-F. Ge and I. M. Shoemaker, *JHEP* **11**, 066 (2018), [arXiv:1710.10889 \[hep-ph\]](#).
- [225] D. Akimov *et al.* (COHERENT), (2019), [arXiv:1911.06422 \[hep-ex\]](#).
- [226] M. Abdullah, D. Aristizabal Sierra, B. Dutta, and L. E. Strigari, *Phys. Rev. D* **102**, 015009 (2020), [arXiv:2003.11510 \[hep-ph\]](#).
- [227] B. Armbruster *et al.* (KARMEN), *Phys. Lett. B* **423**, 15 (1998).
- [228] L. Auerbach *et al.* (LSND), *Phys. Rev. C* **64**, 065501 (2001), [arXiv:hep-ex/0105068](#).
- [229] A. Bolozdynya *et al.* (2012) [arXiv:1211.5199 \[hep-ex\]](#).
- [230] E. Bertuzzo, F. F. Deppisch, S. Kulkarni, Y. F. Perez Gonzalez, and R. Zukanovich Funchal, *JHEP* **04**, 073 (2017), [arXiv:1701.07443 \[hep-ph\]](#).
- [231] M. C. Gonzalez-Garcia, M. Maltoni, Y. F. Perez-Gonzalez, and R. Zukanovich Funchal, *JHEP* **07**, 019 (2018), [arXiv:1803.03650 \[hep-ph\]](#).
- [232] C. Boehm, D. G. Cerdeño, P. A. N. Machado, A. Olivares-Del Campo, and E. Reid, *JCAP* **1901**, 043 (2019), [arXiv:1809.06385 \[hep-ph\]](#).
- [233] A. Aleksandrov *et al.* (NEWS), (2016), [arXiv:1604.04199 \[astro-ph.IM\]](#).
- [234] S. Rajendran, N. Zobrist, A. O. Sushkov, R. Walsworth, and M. Lukin, *Phys. Rev. D* **96**, 035009 (2017), [arXiv:1705.09760 \[hep-ph\]](#).
- [235] A. Drukier, K. Freese, A. Lopez, D. Spergel, C. Cantor, G. Church, and T. Sano, (2012), [arXiv:1206.6809 \[astro-ph.IM\]](#).
- [236] D. Nygren, *J. Phys. Conf. Ser.* **460**, 012006 (2013).
- [237] D. P. Snowden-Ifft and J. L. Gauvreau, *Rev. Sci. Instrum.* **84**, 053304 (2013), [arXiv:1301.7145 \[physics.ins-det\]](#).
- [238] E. Baracchini, G. Cavoto, G. Mazzitelli, F. Murtas, F. Renga, and S. Tomassini, *JINST* **13**, P04022 (2018), [arXiv:1710.01994 \[physics.ins-det\]](#).
- [239] D. P. Snowden-Ifft, C. J. Martoff, and J. M. Burwell, *Phys. Rev. D* **61**, 101301 (2000), [arXiv:astro-ph/9904064 \[astro-ph\]](#).
- [240] C. J. Martoff, D. P. Snowden-Ifft, T. Ohnuki, N. Spooner, and M. Lehner, *Nucl. Instrum. Methods A* **440**, 355 (2000).
- [241] J. B. R. Battat *et al.* (DRIFT), *Phys. Dark Univ.* **9-10**, 1 (2014), [arXiv:1410.7821 \[hep-ex\]](#).
- [242] J. B. R. Battat *et al.* (DRIFT), *Astropart. Phys.* **91**, 65 (2017), [arXiv:1701.00171 \[astro-ph.IM\]](#).
- [243] S. Burgos *et al.*, *Astropart. Phys.* **31**, 261 (2009), [arXiv:0809.1831 \[astro-ph\]](#).
- [244] J. B. R. Battat *et al.*, *JINST* **11**, P10019 (2016), [arXiv:1606.05364 \[physics.ins-det\]](#).
- [245] G. De Geronimo *et al.*, *IEEE Trans. Nucl. Sci.* **58**, 1376 (2011).
- [246] J. L. Gauvreau, Private communication.
- [247] A. Oed, *Nucl. Instrum. Methods A* **263**, 351 (1988).
- [248] F. Sauli, *Nucl. Instrum. Meth. A* **477**, 1 (2002).
- [249] A. Ochi, T. Nagayoshi, S. Koishi, T. Tanimori, T. Nagae, and M. Nakamura, *Imaging 2000. Proceedings, Stockholm, Sweden, June 27-July 1, 2000*, *Nucl. Instrum. Meth. A* **471**, 264 (2000).
- [250] Y. Giomataris, P. Rebourgeard, J. P. Robert, and G. Charpak, *Nucl. Instrum. Methods A* **376**, 29 (1996).
- [251] R. De Oliveira, Private communication.
- [252] F. J. Iguaz *et al.*, *Eur. Phys. J. C* **76**, 529 (2016), [arXiv:1512.01455 \[physics.ins-det\]](#).
- [253] Q. Riffard *et al.*, *JINST* **12**, P06021 (2017), [arXiv:1504.05865 \[astro-ph.IM\]](#).
- [254] Q. Riffard, *Non-baryonic dark matter directional detection with MIMAC*, Ph.D. thesis, Université Grenoble Alpes (2015).
- [255] K. Nakamura *et al.*, *PTEP* **2015**, 043F01 (2015).
- [256] F. Sauli, *Nucl. Instrum. Methods A* **386**, 531 (1997).
- [257] K. Nakamura *et al.*, *J. Phys. Conf. Ser.* **469**, 012003 (2013).
- [258] D. A. Dwyer, “Demonstration of a true 3D micro-power sensor for liquid argon time projection chambers,” In preparation (2019).
- [259] D. A. Dwyer, Private communication (2019).
- [260] M. Garcia-Sciveres, D. Arutinov, M. Barbero, R. Beccherle, S. Dube, *et al.*, *Nucl. Instrum. Methods A* **636**, S155 (2011).
- [261] V. M. Khoa, Master Thesis **Tokyo Institute of Technology**, 1 (February 2010).
- [262] R. Plackett, R. Ballabriga, M. Campbell, X. Llopart, T. Tick, L. Tlustos, D. Turecek, S. Vahanen, and W. Wong, *PoS VERTEX2010*, 030 (2010).
- [263] R. Bates, *PoS VERTEX2015*, 006 (2015).
- [264] P. Lewis, S. Vahsen, M. Hedges, I. Jaegle, I. Seong, *et al.*, *Nucl. Instrum. Methods A* **789**, 81 (2015), [arXiv:1410.1131 \[physics.ins-det\]](#).
- [265] T. Kim, M. Freytsis, J. Button-Shafer, J. Kadyk, S. E. Vahsen, and W. A. Wenzel, *Nucl. Instrum. Meth. A* **589**, 173 (2008).
- [266] S. Vahsen, K. Oliver-Mallory, M. Lopez-Thibodeaux, J. Kadyk, and M. Garcia-Sciveres, *Nucl. Instrum. Methods A* **738**, 111 (2014).
- [267] S. Vahsen, M. Hedges, I. Jaegle, S. Ross, I. Seong, *et al.*, *Nucl. Instrum. Methods A* **788**, 95 (2015).
- [268] K. N. Buckland, M. J. Lehner, G. E. Masek, and M. Mojaver, *Phys. Rev. Lett.* **73**, 1067 (1994).
- [269] F. A. F. Fraga *et al.*, *Nucl. Instrum. Meth. A* **478**, 357 (2002).
- [270] S. Ahlen, J. Battat, T. Caldwell, C. Deaconu, D. Dujmic, *et al.*, *Phys. Lett. B* **695**, 124 (2011), [arXiv:1006.2928 \[hep-ex\]](#).
- [271] J. B. R. Battat *et al.*, *Nucl. Instrum. Meth. A* **755**, 6 (2014).
- [272] C. Deaconu *et al.*, *Phys. Rev. D* **95**, 122002 (2017), [arXiv:1705.05965 \[astro-ph.IM\]](#).

- [273] L. M. S. Margato, F. A. F. Fraga, S. T. G. Fetal, M. M. F. R. Fraga, E. F. S. Balau, A. Blanco, R. Ferreira-Marques, and A. J. P. L. Policarpo, *Nucl. Instrum. Meth. A* **535**, 231 (2004).
- [274] S. T. G. Fetal *et al.*, *Nucl. Instrum. Methods A* **581**, 202 (2007).
- [275] V. C. Antochi, E. Baracchini, G. Cavoto, E. D. Marco, M. Marafini, G. Mazzitelli, D. Pinci, F. Renga, S. Tomassini, and C. Voena, *JINST* **13**, P05001 (2018), [arXiv:1803.06860](https://arxiv.org/abs/1803.06860) [physics.ins-det].
- [276] A. Kaboth *et al.*, *Nucl. Instrum. Methods A* **592**, 63 (2008), [arXiv:physics.ins-det/0803.2195](https://arxiv.org/abs/physics.ins-det/0803.2195) [physics.ins-det].
- [277] D. Pinci, E. Di Marco, F. Renga, C. Voena, E. Baracchini, G. Mazzitelli, A. Tomassini, G. Cavoto, V. C. Antochi, and M. Marafini, *PoS EPS-HEP2017*, 077 (2017).
- [278] G. Mazzitelli *et al.* (CYGNUS), (2019), [arXiv:1901.04190](https://arxiv.org/abs/1901.04190) [physics.ins-det].
- [279] R. Campagnola, *Study and optimization of the light-yield of a triple-GEM detector*, Ph.D. thesis, University of Rome, La Sapienza (2018).
- [280] I. A. Costa *et al.*, (2019), [arXiv:1905.04066](https://arxiv.org/abs/1905.04066) [physics.ins-det].
- [281] E. Baracchini, Private communication.
- [282] T. Naka *et al.*, *Proceedings, 12th Pisa Meeting on Advanced Detectors: La Biodola, Isola d'Elba, Italy, May 20-26, 2012*, *Nucl. Instrum. Meth. A* **718**, 519 (2013).
- [283] J. J. Mock, M. Barbic, D. R. Smith, D. A. Schultz, and S. Schultz, *J. of Chem. Phys.* **116** (2002), 10.1063/1.1462610.
- [284] Y. Hochberg, Y. Kahn, M. Lisanti, C. G. Tully, and K. M. Zurek, *Phys. Lett. B* **772**, 239 (2017), [arXiv:1606.08849](https://arxiv.org/abs/1606.08849) [hep-ph].
- [285] S. Betts *et al.* (2013) [arXiv:1307.4738](https://arxiv.org/abs/1307.4738) [astro-ph.IM].
- [286] P. H. Heckmann, H. Hansen, and A. Flammersfeld, *Zeitschrift für Physik* **162**, 84 (1961).
- [287] P. Belli, R. Bernabei, C. Bacci, A. Incicchitti, and D. Prosperi, *Nuovo Cim. C* **15**, 475 (1992).
- [288] N. J. C. Spooner, J. W. Roberts, and D. R. Tovey, in *Proceedings, 1st International Workshop on The identification of dark matter (IDM 1996)* (1996) pp. 481–486.
- [289] Y. Shimizu, M. Minowa, H. Sekiya, and Y. Inoue, *Nucl. Instrum. Methods A* **496**, 347 (2003), [arXiv:astro-ph/0207529](https://arxiv.org/abs/astro-ph/0207529) [astro-ph].
- [290] F. Cappella *et al.*, *Eur. Phys. J. C* **73**, 2276 (2013).
- [291] H. Sekiya, M. Minowa, Y. Shimizu, Y. Inoue, and W. Suganuma, *Phys. Lett. B* **571**, 132 (2003), [arXiv:astro-ph/0307384](https://arxiv.org/abs/astro-ph/0307384) [astro-ph].
- [292] M. Cadeddu *et al.*, *JCAP* **1901**, 014 (2019), [arXiv:1704.03741](https://arxiv.org/abs/1704.03741) [astro-ph.CO].
- [293] G. Jaffé, *Annalen der Physik* **347**, 303 (1913).
- [294] Y. Nakajima, A. Goldschmidt, M. Long, D. Nygren, C. Oliveira, and J. Renner, *J. Phys. Conf. Ser.* **650**, 012003 (2015), [arXiv:1505.03586](https://arxiv.org/abs/1505.03586) [physics.ins-det].
- [295] Y. Nakajima, A. Goldshmidt, H. S. Matis, T. Miller, D. R. Nygren, C. A. B. Oliveira, and J. Renner, *JINST* **11**, C03041 (2016), [arXiv:1511.02257](https://arxiv.org/abs/1511.02257) [physics.ins-det].
- [296] H. Cao *et al.* (SCENE), *Phys. Rev. D* **91**, 092007 (2015), [arXiv:1406.4825](https://arxiv.org/abs/1406.4825) [physics.ins-det].
- [297] L. M. Capparelli, G. Cavoto, D. Mazzilli, and A. D. Polosa, *Phys. Dark Univ.* **9-10**, 24 (2015), [Erratum: *Phys. Dark Univ.* **11**, 79 (2016)], [arXiv:1412.8213](https://arxiv.org/abs/1412.8213) [physics.ins-det].
- [298] D. P. Snowden-Ifft, *Rev. Sci. Instrum.* **85**, 013303 (2014).
- [299] N. S. Phan, R. Lafler, R. J. Lauer, E. R. Lee, D. Loomba, J. A. J. Matthews, and E. H. Miller, *JINST* **12**, P02012 (2017), [arXiv:1609.05249](https://arxiv.org/abs/1609.05249) [physics.ins-det].
- [300] T. Ikeda *et al.*, *EPJ Web Conf.* **174**, 02006 (2018), [arXiv:1709.06219](https://arxiv.org/abs/1709.06219) [astro-ph.IM].
- [301] T. N. Thorpe, *Gain resolution studies and first dark matter search with novel 3D nuclear recoil detectors*, Ph.D. thesis, University of Hawaii at Manoa (2018).
- [302] T. Ikeda, T. Shimada, H. Ishiura, K. Nakamura, T. Nakamura, and K. Miuchi, *JINST* **15**, P07015 (2020), [arXiv:2004.09706](https://arxiv.org/abs/2004.09706) [physics.ins-det].
- [303] I. Jaegle, P. Lewis, M. Garcia-Sciveres, M. Hedges, T. Hemperek, J. Janssen, Q. Ji, D.-L. Pohl, S. Ross, J. Schueler, I. Seong, T. Thorpe, and S. Vahsen, *Nucl. Instrum. Meth. A* **945**, 162296 (2019).
- [304] M. Cannoni, *Phys. Rev. D* **87**, 075014 (2013), [arXiv:1211.6050](https://arxiv.org/abs/1211.6050) [astro-ph.CO].
- [305] P. C. Divari, T. S. Kosmas, J. D. Vergados, and L. D. Skouras, *Phys. Rev. C* **61**, 054612 (2000).
- [306] J. Bergstrom, M. C. Gonzalez-Garcia, M. Maltoni, C. Pena-Garay, A. M. Serenelli, and N. Song, *JHEP* **03**, 132 (2016), [arXiv:1601.00972](https://arxiv.org/abs/1601.00972) [hep-ph].
- [307] J. F. Ziegler and P. Biersack, *The Stopping and Range of Ions in Matter* (Pergamon Press, New York, 1985).
- [308] S. Biagi, <http://magboltz.web.cern.ch/magboltz/>.
- [309] C. Deaconu, *A model of the directional sensitivity of low-pressure CF₄ dark matter detectors*, Ph.D. thesis, MIT (2015).
- [310] Y. H. Hilal and L. G. Christophorou, *Journal of Physics D: Applied Physics* **20**, 975 (1987).
- [311] A. Scarff, *Developments Towards a Scaled-Up One-Dimensional Directional Dark Matter Detector*, Ph.D. thesis, Sheffield U. (2017-11).
- [312] H. Ishiura, “Basic characteristics of MPGD in anion gas for 3D track detector used for direct search of dark matter with direction sensitivity,” (2019).
- [313] R. Lafler, D. Loomba, Private communication.
- [314] E. Baracchini *et al.*, (2020), [arXiv:2004.10493](https://arxiv.org/abs/2004.10493) [physics.ins-det].
- [315] E. Baracchini *et al.*, (2020), [arXiv:2007.01763](https://arxiv.org/abs/2007.01763) [physics.ins-det].
- [316] N. Smirnov, *Ann. Math. Statist.* **19**, 279 (1948).
- [317] D. Loomba, Private communication.
- [318] S. E. Vahsen, Private communication.
- [319] P. Sadowski, Private communication.
- [320] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Eur. Phys. J. C* **71**, 1554 (2011), [Erratum: *Eur. Phys. J. C* **73**, 2501 (2013)], [arXiv:1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].
- [321] S. J. Ross *et al.* (2012) pp. 1760–1765, [arXiv:1304.0507](https://arxiv.org/abs/1304.0507) [physics.ins-det].
- [322] E. Aprile *et al.* (XENON100), *J. Phys. G* **40**, 115201 (2013), [arXiv:1306.2303](https://arxiv.org/abs/1306.2303) [astro-ph.IM].
- [323] S. Agostinelli *et al.* (GEANT4), *Nucl. Instrum. Meth. A* **506**, 250 (2003).
- [324] W. B. Wilson *et al.*, SOURCES 4C : a code for calculating ([alpha],n), spontaneous fission, and delayed neutron sources and spectra. (2002).
- [325] V. A. Kudryavtsev, *Computer Physics Communications* **180**, 339 (2009), [arXiv:0810.4635](https://arxiv.org/abs/0810.4635) [physics.comp-ph].
- [326] “UK DM project: radioactivity test results,” .
- [327] B. J. Mount *et al.*, (2017), [arXiv:1703.09144](https://arxiv.org/abs/1703.09144) [physics.ins-det].

- [328] V. Alvarez *et al.*, *JINST* **8**, T01002 (2013), [arXiv:1211.3961 \[physics.ins-det\]](#).
- [329] A. Hallin, Private communication (2017).
- [330] E. Sala, I. S. Hahn, W. G. Kang, G. W. Kim, Y. D. Kim, M. H. Lee, D. S. Leonard, and S. Y. Park, *Journal of Physics: Conference Series* **718** (2016).
- [331] T. Hashimoto, K. Miuchi, K. Nakamura, R. Yakabe, T. Ikeda, R. Taishaku, M. Nakazawa, H. Ishiura, A. Ochi, and Y. Takeuchi, *AIP Conf. Proc.* **1921**, 070001 (2018), [arXiv:1707.09744 \[physics.ins-det\]](#).
- [332] T. Hashimoto *et al.*, *NIM A* **977**, 164285 (2020).
- [333] S. Cebrián *et al.*, *Astropart. Phys.* **34**, 354 (2011), [arXiv:1005.2022 \[physics.ins-det\]](#).
- [334] N. Phan, R. Lauer, E. Lee, D. Loomba, J. Matthews, and E. Miller, *Astroparticle Physics* **84**, 82 (2016).
- [335] R. Oliveira, Private communication (2017).
- [336] K. Nakamura *et al.*, *Physics Procedia* **61**, 737 (2015).
- [337] Surface Mount Precision Plate Resistors https://www.jfine.co.jp/eng/e_c/resistors/surface_mount/pdf/e_sm.pdf.
- [338] J. Castel, S. Cebrian, I. Coarasa, *et al.*, *The European Physical Journal C* **79**, 782 (2019).
- [339] K. Miuchi *et al.*, *PTEP* **2019**, 063H01 (2019), [arXiv:1903.01663 \[physics.ins-det\]](#).
- [340] M. Robinson, V. A. Kudryavtsev, R. Luscher, J. E. McMillan, P. K. Lightfoot, N. J. C. Spooner, N. J. T. Smith, and I. Liubarsky, *Nucl. Instrum. Meth. A* **511**, 347 (2003), [arXiv:hep-ex/0306014 \[hep-ex\]](#).
- [341] J. Brack *et al.*, *Phys. Procedia* **61**, 130 (2015), [arXiv:1404.2253 \[physics.ins-det\]](#).
- [342] C. Couturier, Q. Riffard, N. Sauzet, O. Guillaudin, F. Naraghi, and D. Santos, *Journal of Instrumentation* **12**, P11020 (2017).
- [343] S. Cebrián, *Int. J. Mod. Phys. A* **32**, 1743006 (2017), [arXiv:1708.07449 \[nucl-ex\]](#).
- [344] L. Baudis, A. Kish, F. Piastra, and M. Schumann, *Eur. Phys. J. C* **75**, 485 (2015), [arXiv:1507.03792 \[astro-ph.IM\]](#).
- [345] J. J. Back and Y. A. Ramachers, *Nucl. Instrum. Meth. A* **586**, 286 (2008), [arXiv:0709.3472 \[nucl-ex\]](#).
- [346] T. Gamble, Private communication (2018).
- [347] V. Álvarez *et al.* (NEXT Collaboration), *J. Instrum.* **9**, P03010 (2014), [arXiv:1311.3242 \[physics.ins-det\]](#).
- [348] J. L. Bessede, I. Huet, G. Montillet, E. Barbier, and J. Micozzi, *Implementation Of Treatment & Recovery Of the SF6 Gas Containing A High Amount Of Decomposition Products Due To High Voltage Electrical Interruptions* (2018).
- [349] A. C. Ezeribe, W. Lynch, R. R. M. Gregorio, J. Mckeand, A. Scarff, and N. J. C. Spooner, *JINST* **12**, P09025 (2017), [arXiv:1707.07772 \[physics.ins-det\]](#).
- [350] Q. Riffard *et al.*, *JINST* **11**, P08011 (2016), [arXiv:1602.01738 \[astro-ph.IM\]](#).