

Beyond the Neutrino Fog: Dark Matter Discovery Limits
with the Full XENONnT Background Model
REU Program at Columbia University – Nevis
Laboratories

AlisonDawn Lewis¹

¹Department of Physics, University of Washington

August 1, 2026



Abstract

The imminent detection of solar and atmospheric neutrinos via coherent elastic neutrino–nucleus scattering (CE ν NS) in multi-tonne liquid-xenon detectors is expected to establish an irreducible background to weakly interacting massive particle (WIMP) searches, commonly termed the “neutrino fog.” We extend the background model used in the XENONnT dark-matter search to compute the full topology of the background fog for the sr1b science run, following the exposure-derivative formalism of O’Hare. Using a profile-likelihood framework with per-component nuisance parameters, we compute the expected 3σ discovery cross section as a function of exposure for a grid of WIMP masses and map the opacity index

$$n = - \left(\frac{d \ln \sigma}{d \ln N} \right)^{-1}$$

across the (m_χ, σ) plane. We find that including the radiogenic neutron background, arising from (α, n) and spontaneous-fission reactions in detector materials, raises the floor and introduces an additional high-mass fog feature that is absent when only cosmogenic and astrophysical neutrino sources are considered. We compare our derived floor to the published XENONnT SR0+SR1 upper limit and discuss the implications for future exposure scaling of tonne-scale xenon experiments.

Contents

1	Introduction	2
2	Background Model and Detector Response	2
3	Statistical Framework	3
4	Mapping the Neutrino Fog	4
5	The Role of the Radiogenic Neutron Background	6
6	Comparison to the Measured XENONnT Limit	8
7	Summary and Conclusions	9
8	Acknowledgements	10

1 Introduction

Weakly interacting massive particles (WIMPs) with masses between roughly $1\text{ GeV } c_0^{-2}$ and several $\text{TeV } c_0^{-2}$ remain among the most experimentally accessible candidates for particle dark matter. Dual-phase liquid-xenon time projection chambers (TPCs), such as XENONnT, LZ, and PandaX-4T, currently set the most stringent limits on the spin-independent (SI) WIMP–nucleon cross section by searching for nuclear recoils in ultra-low-background detectors [2, 1, 4]. As these experiments approach multi-tonne-year exposures, coherent elastic neutrino–nucleus scattering (CE ν NS) from solar, atmospheric, and diffuse-supernova neutrinos produces nuclear recoils that are effectively indistinguishable from a WIMP signal of a corresponding mass [7]. This background was historically referred to as the “neutrino floor,” implying a hard sensitivity limit, but subsequent work has shown that, with sufficiently large exposure and a well-constrained background model, discovery can in principle proceed past this point at a reduced statistical rate. The term “neutrino fog” is therefore more accurate [7].

O’Hare formalized this distinction by defining the fog in terms of the local scaling of a hypothetical discovery limit with the number of accumulated background events,

$$n = - \left(\frac{d \ln \sigma}{d \ln N} \right)^{-1}, \quad (1)$$

where $n = 2$ corresponds to ordinary Poissonian background subtraction, $\sigma \propto N^{-1/2}$, and $n > 2$ marks the onset of systematics-dominated saturation [7]. This definition has the advantage of being independent of an arbitrary choice of exposure or energy threshold, unlike earlier interpolation-based constructions of the neutrino floor [3].

The present work applies this exposure-derivative framework to the specific background model used in the XENONnT sr1b science run, rather than to a generic xenon target. In addition to the CE ν NS components considered by O’Hare, we incorporate the full set of backgrounds reported by the XENON Collaboration for its recent 3.1 tonne year WIMP search: flat and tritium-like electronic recoil components, ^{37}Ar , radiogenic neutrons, accidental coincidences, and surface events [2]. Because the radiogenic neutron background produces single-scatter nuclear recoils with a spectral shape that can closely track a WIMP signal at intermediate-to-high mass, we specifically investigate how its inclusion reshapes the fog relative to a neutrino-only background model. Background sources and detector strategies for suppressing them are reviewed extensively in Ref. [6].

2 Background Model and Detector Response

We use the sr1b background templates provided for the XENONnT analysis, binned in corrected S1 (cS1), corrected S2 (cS2), and radius. Each background component—flat ER, tritium-like ER, ^{37}Ar , radiogenic neutrons, solar and atmospheric/DSNB CE ν NS, accidental coincidences, and surface events—is stored as a `multihist.Histdd` object and combined into a total expected background rate per bin. Poisson and systematic uncertainties are propagated in quadrature, assuming that the uncertainties on the individual background components are uncorrelated with one another.

Figure 1 shows the combined sr1b background rate and its total per-bin uncertainty in the cS1–cS2 plane used for this analysis. The prominent diagonal band around $\text{cS2} \approx 2\text{--}5 \times 10^3$ photoelectrons is the electronic-recoil band, while the region beneath it contains the nuclear-recoil-like backgrounds that overlap the WIMP search region.




Figure 1: Combined background rate (left) and total background uncertainty (right) for the *sr1b* science run, shown in the *cS1*–*cS2* observable plane. The nuclear-recoil-like backgrounds populate the lower band beneath the dominant electronic-recoil band and directly overlap the WIMP-like signal region.

For a given WIMP mass m_χ , we generate the expected differential recoil spectrum using the standard spin-independent, isospin-conserving scattering formalism with a Helm form factor and the Standard Halo Model velocity distribution, following the conventions adopted in the XENONnT statistical framework [2, 5]. Signal and background expectations are compared using a binned Poisson likelihood in *cS1*–*cS2* with Gaussian constraint terms for each nuisance parameter.

3 Statistical Framework

Discovery significance is computed using the profile-likelihood-ratio test statistic of Cowan et al. [5], evaluated on Asimov datasets to avoid the computational cost of repeated toy Monte Carlo generation at every point in the $(m_\chi, \sigma, \text{exposure})$ grid. For a background-only model nested within a signal-plus-background model, we define

$$q_0 = -2 \ln \left[\frac{\mathcal{L}(0, \hat{\hat{\theta}})}{\mathcal{L}(\hat{\sigma}, \hat{\hat{\theta}})} \right], \quad \hat{\sigma} \geq 0, \quad (2)$$

and require $q_0 > 9$ for a nominal 3σ discovery.

We initially implemented a single global nuisance parameter scaling the total background rate, but found that this significantly underestimates the true discovery reach. With a single scale factor β multiplying the entire background template, every component is forced to rise and fall together: a fluctuation that should only be absorbed by, for example, the radiogenic-neutron normalization would instead also rescale the ER background by the same factor, even though the two arise from unrelated physical processes and their rates have no reason to be correlated. In other words, the fit lets each component’s normalization be constrained by the rates of physically unrelated components rather than by its own, independent systematic uncertainty. We therefore extended the fit to a multi-nuisance-parameter model with one scale factor β_i per background component,

$$i = \{\text{ER-flat}, {}^3\text{H}, {}^{37}\text{Ar}, \text{radiogenic}, \text{solar}, \text{atm} + \text{DSNB}, \text{AC}, \text{surface}\},$$

with each factor constrained by a Gaussian term whose width is given by the reported uncertainty for that component.

Figure 2 illustrates the impact of this change for a $200 \text{ GeV } c_0^{-2}$ WIMP. The single-nuisance-parameter model produces a smoothly decreasing discovery cross section that never flattens, because it cannot absorb an excess in one background alone. The full multi-nuisance-parameter model reproduces the expected saturation behavior, with a break in the power-law slope of cross section versus exposure once the background components become individually constraining.




Figure 2: Expected 3σ discovery cross section as a function of exposure for a $200 \text{ GeV } c_0^{-2}$ WIMP, using a single global background nuisance parameter (left) and independent per-component nuisance parameters β_i (right). Only the multi-nuisance-parameter model reproduces the flattening associated with the neutrino fog.

4 Mapping the Neutrino Fog

Following O’Hare [7], we quantify the opacity of the fog using Eq. (1), computed from the local logarithmic slope of the discovery cross-section curve as a function of exposure, or equivalently the number of background events N . The condition $n = 2$ corresponds to standard Poissonian background subtraction, while $n > 2$ indicates that systematic uncertainties on the background normalization are limiting further improvement in sensitivity.

Figure 3 shows this index as a function of exposure for a $200 \text{ GeV } c_0^{-2}$ WIMP using the full sr1b background model. A clear peak in n , reaching above 5, appears at exposures of order 10^4 – 10^5 years. Although these exposures are far beyond any realistic near-term experiment, the behavior illustrates that the discovery reach saturates well before neutrino and radiogenic backgrounds are individually resolved.




Figure 3: Fog opacity index n versus exposure for a $200 \text{ GeV } c_0^{-2}$ WIMP, computed from the derivative of the sr1b discovery curve. The dashed line marks $n = 2$, the boundary between the Poissonian and systematics-saturated regimes.

We repeat this calculation for a grid of WIMP masses, shown in Fig. 4. All masses show a similar qualitative structure: an initial rise toward $n \approx 2$ – 2.5 at intermediate exposures, driven primarily by the flat ER and radiogenic neutron backgrounds, followed by a secondary and substantially higher peak at very large exposures once the CE ν NS backgrounds begin to dominate and their flux-normalization systematics become the limiting factor. The height of the second peak grows with WIMP mass, from $n \approx 2.2$ at $10 \text{ GeV } c_0^{-2}$ to $n \gtrsim 5$ at $500 \text{ GeV } c_0^{-2}$ to $1000 \text{ GeV } c_0^{-2}$, reflecting the increasingly close spectral overlap between high-mass WIMP recoils and the harder recoil spectra of radiogenic neutrons and atmospheric neutrinos.




Figure 4: Fog opacity index n versus exposure for WIMP masses of 10, 50, 500, and 1000 GeV c_0^{-2} . The double-peak structure reflects sequential saturation by the local ER/radiogenic backgrounds and, at much larger exposures, the neutrino backgrounds.

5 The Role of the Radiogenic Neutron Background

A central question for this analysis is whether backgrounds intrinsic to the detector materials, rather than the astrophysical neutrino flux alone, meaningfully reshape the fog. We computed the full (m_χ, σ) fog map in two configurations: first, using the CE ν NS neutrino components (solar, atmospheric, and DSNB), surface, AC, and ER backgrounds, and second, restoring the radiogenic neutron background from (α, n) and spontaneous-fission reactions in detector materials using the rate reported for sr1b [2].

Figure 5 shows the fog map without the radiogenic background. The color scale traces a single, smoothly varying band of elevated n running from the sub-10 GeV c_0^{-2} solar-neutrino-dominated region toward higher masses, consistent with the CE ν NS-only picture described in the literature [7].




Figure 5: Fog opacity map computed using only the $CE\nu NS$ backgrounds (solar, atmospheric, and $DSNB$), surface, AC, and ER backgrounds, with no radiogenic neutron contribution. The color axis shows the opacity index n ; the single fog band tracks the neutrino-induced nuclear-recoil spectrum.

Figure 6 shows the same calculation with the radiogenic neutron background restored. A second, higher-opacity feature appears at masses above approximately $50 \text{ GeV } c_0^{-2}$ to $100 \text{ GeV } c_0^{-2}$, running roughly parallel to but distinct from the underlying $CE\nu NS$ band. The overall opacity in this region also increases, with the color scale extending to $n \approx 4$ rather than saturating below it. This reflects the fact that radiogenic neutrons, unlike $CE\nu NS$ events, are subject to comparatively large and only partially constrained normalization uncertainties tied to the radiopurity of detector construction materials. Their recoil spectrum tracks a WIMP-like shape closely enough at these masses that, at the normalization level assumed here, it produces a genuine statistical degeneracy with a WIMP signal. Unlike the neutrino background, however, the radiogenic-neutron rate is not a fundamental limit: it can in principle be reduced through active neutron vetoing, the use of low-radioactivity construction materials, and discrimination based on multiple-scatter topology.




Figure 6: Fog opacity map with the radiogenic neutron background included alongside other background sources. A second high-opacity feature emerges above approximately $50 \text{ GeV } c_0^{-2}$, reflecting the additional systematic-limited saturation introduced by detector-material neutron emission.

This result indicates that, for a real experiment such as XENONnT, the effective floor at intermediate-to-high WIMP mass is not set by astrophysical neutrinos alone. The radiopurity and neutron-emission budget of the detector materials contribute an additional and, at sr1b exposures, comparably sized source of background-normalization degeneracy. This component, unlike the astrophysical neutrino flux, is in principle reducible through improved detector radiopurity, active neutron vetoes, and scattering-topology discrimination, so it does not represent a hard, unavoidable limit in the same sense as the neutrino background.

6 Comparison to the Measured XENONnT Limit

Finally, we overlay our computed floor, defined as the $n = 2$ contour of the full radiogenic-plus-CE ν NS fog map, against the actual 90% confidence-level upper limit reported by the XENON Collaboration for the combined SR0+SR1 3.1 tonne year exposure [2]. Figure 7 reproduces the fog map of Fig. 6 with the published XENONnT exclusion limit drawn as a solid black curve and the reported -1σ sensitivity curve shown as a dashed line.




Figure 7: Fog opacity map, including the radiogenic neutron background, with the published XENONnT SR0+SR1 90% confidence-level upper limit overlaid as a solid black curve and the reported -1σ sensitivity curve shown as a dashed line. The experimental limit sits just above the onset of the fog, $n \approx 2-3$, across the probed mass range.

The XENONnT limit tracks closely along the upper edge of our computed $n \approx 2-3$ contour across the $10 \text{ GeV } c_0^{-2}$ to $1000 \text{ GeV } c_0^{-2}$ range probed by the experiment and reproduces the characteristic minimum near $30 \text{ GeV } c_0^{-2}$ to $50 \text{ GeV } c_0^{-2}$ [2]. This indicates that the current sr1b exposure is only beginning to approach the systematics-limited regime identified by our fog calculation, rather than being deep within it. This is consistent with XENONnT’s reported sensitivity still improving close to the Poissonian $\sigma \propto N^{-1/2}$ scaling for this dataset. The modest offset between the computed floor and the observed limit is consistent with the difference between an idealized Asimov-based projection and a real, unblinded dataset subject to the particular background realization observed in SR1.

7 Summary and Conclusions

We have applied the exposure-derivative definition of the neutrino floor introduced by O’Hare [7] to the full background model used in the XENONnT sr1b WIMP search, extending the calculation beyond CE ν NS backgrounds to include the radiogenic neutron background intrinsic to the detector. We find that this detector-material background introduces a distinct, higher-opacity feature in the fog above approximately $50 \text{ GeV } c_0^{-2}$ that is absent from a neutrino-only calculation. This indicates that the ultimate discovery-limited floor for a real xenon experiment is shaped not only by astrophysical neutrino flux uncertainties, but also by detector radiopurity.

Comparing our computed floor to the published XENONnT SR0+SR1 upper limit shows good qualitative agreement in both shape and the location of the sensitivity minimum. This supports the use of the exposure-derivative framework as a physically motivated benchmark for projecting the discovery reach of future tonne-to-multi-tonne xenon experiments such as XLZD/DARWIN.

8 Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. PHY-2349438.

References

- [1] J. Aalbers et al. Dark matter search results from the lux-zeplin experiment. *Physical Review Letters*, 135, 2025.
- [2] E. Aprile et al. Wimp dark matter search using a 3.1 tonne-year exposure of the xenonnT experiment. *arXiv preprint arXiv:2502.18005*, 2025.
- [3] J. Billard, L. Strigari, and E. Figueroa-Feliciano. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Physical Review D*, 89:023524, 2014.
- [4] Z. Bo et al. Dark matter search results from pandax. *Physical Review Letters*, 134:011805, 2025.
- [5] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71:1554, 2011.
- [6] Teresa Marrodán Undagoitia and Ludwig Rauch. Dark matter direct-detection experiments. *arXiv preprint arXiv:1509.08767*, 2015.
- [7] Ciaran A. J. O’Hare. Fog on the horizon: A new definition of the neutrino floor for direct dark matter searches. *arXiv preprint arXiv:2109.03116*, 2021.