

Approximate Differentiable Likelihoods for Astroparticle Physics Experiments

Juehang Qin¹ and Christopher Tunnell¹

¹Department of Physics and Astronomy, Rice University, Houston, TX, USA

E-mail: qinjuehang@rice.edu

Abstract. Traditionally, inference in liquid xenon direct detection dark matter experiments has used estimators of event energy or density estimation of simulated data. Such methods have drawbacks compared to the computation of explicit likelihoods, such as an inability to conduct statistical inference in high-dimensional parameter spaces, or a failure to make use of all available information. In this work, we implement a continuous approximation of an event simulator model within a probabilistic programming framework, allowing for the application of high performance gradient-based inference methods such as the No-U-Turn Sampler. We demonstrate an improvement in inference results, with percent-level decreases in measurement uncertainties. Finally, in the case where some observables can be measured using multiple independent channels, such a method also enables the incorporation of additional information seamlessly, allowing for full use of the available information to be made.

1 Introduction

Inference in particle physics is often based on black-box simulators that can faithfully reproduce the physics involved; however, such simulations do not provide tractable likelihoods or probability densities, leading to the use of either Approximate Bayesian Computation or approximate likelihoods obtained via density estimation [1]. Simulators constructed within probabilistic programming languages present an alternative that can allow us to conduct simulation-based inference by computing differentiable likelihoods directly using the models that our simulators are based on [1].

Prior work in particle physics which integrates existing event simulators into the probabilistic programming paradigm, but without making the simulators directly compatible with automatic differentiation [2, 3]. Making differentiable models is often challenging in particle physics as many physical processes, such as generation of quanta, are discrete. Despite these difficulties, compatibility with automatic differentiation frameworks such as JAX [4] confers significant benefits, enabling the application of high-performance gradient-based inference methods, such as variational inference [5, 6], and the No-U-Turn Sampler (NUTS) [7]. In this work, we demonstrate an approach using an approximate binomial distribution and the probabilistic programming framework `numpyro` [8, 9] to compute differentiable likelihoods. To that end, we introduce `bayesNEST`, a partial implementation of the NEST model for liquid noble element particle physics detectors, including dark matter experiments XENON and LZ [10, 11].

Dark matter is one of the greatest mysteries in physics—while most of the matter in the universe is dark matter, its particle nature is not understood [12]. Consequently, liquid xenon detectors are built to search for WIMP dark matter [13]. These experiments have other physics channels, including measurements of double electron capture in ^{124}Xe [14, 15]. These measurements are typically conducted in a low-statistics regime, necessitating sophisticated approaches to inference. Typically, an estimator for event energy is used for inference when searching for electronic recoil signals in these liquid xenon experiments [16, 17]. In analyses using nuclear recoil events, likelihoods are instead estimated using density estimation from

simulated events [18]. Density estimation techniques suffer from the curse of dimensionality [1], whereas estimators computed from observed data might not make use of all available information.

Another approach to this problem is to marginalise over all discrete parameters the likelihood [19]. While this approach is more exact, it can be very computationally and memory intensive to construct the large arrays of likelihoods needed to enumerate over the plausible range of discrete parameters. As such, this work is complementary and can be used to reduce the computational and memory requirements of inference using explicit likelihoods where the approximations we introduce are acceptable.

2 Background

2.1 Dual-phase Xenon Time Projection Chambers

Dual-phase xenon time projection chambers (TPCs) use an active target of liquid xenon, with a layer of gaseous xenon in the upper portion of the detector. The detector is instrumented with arrays of photomultiplier tubes (PMTs) on the top and bottom. When a particle interaction occurs in the liquid xenon target, scintillation photons are produced. This signal, termed S1, is detected by the PMTs. In addition, electrons are liberated and drifted towards the gaseous region at the top of the detector using an electric field, the drift field. These electrons produce a secondary S2 signal in the gaseous phase.

In the current-generation dual-phase xenon detectors experiments LZ and XENONnT, there are 494 PMTs distributed between the top and bottom sensor arrays [20, 10]. Signals are classified as S1 or S2 based on the shape of the signal waveform and the distribution of the signal across the sensors [17, 21], and the amplitude of signals are determined by integrating the signal waveforms. These TPCs are capable of 3D position reconstruction; (x, y) positions can be estimated using the pattern of S2 signal amplitudes across the top PMT array, whereas the z positions can be estimated from the time delay between the S1 and S2 signals. Events can be separated into several categories, including electronic recoils, typically produced by gamma rays and beta decays, and nuclear recoils, produced by neutrons, neutrinos, and WIMP dark matter. More details about liquid xenon TPCs can be found in the following review: [13].

2.2 The Signal Generation Model

In this work, we focus on inference involving electronic recoil events, and thus have implemented a probabilistic model for the S1 and S2 signals expected from beta decays; our implementation will be described in section 3.2. In this section, we will describe how these signals are modelled for electronic recoil events according to the Nobel Element Simulation Technique (NEST). This description is based on [22] and codebase of NEST version 2.3.11 [22].

Given an interaction of energy E , the total number of quanta is computed as the following truncated gaussian discretised by rounding,

$$\mu_q = \frac{E}{W}, \quad N_q \sim [\text{TN}(\mu_q, \sqrt{\mu_q}, a = 0, b = \infty) + 0.5], \quad (1)$$

where $W[\text{eV}] = 21.94 - 2.93\rho$ is the workfunction, the average energy required to produce a single quanta, and (a, b) is the support of the truncated distribution. ρ refers to the density of liquid xenon. Typical values of the workfunction lie within the range of $13 \text{ eV} \sim 14 \text{ eV}$.

The quanta (N_q) are sorted into the categories of excitons and ions using a binomial distribution, $N_i \sim \text{B}(N_q, p_i)$, $N_{ex} = N_q - N_i$, where N_i is the number of ions, N_{ex} is the number of excitons, and p_i is the probability of each quanta being an ion. Finally, some ions represented by N_i might undergo recombination to be produce S1 photons, instead of being part of the S2 ionisation signal. This is modelled using a discretised and truncated skew-normal distribution, as given by

$$\mu_e = (1 - p_r)N_i, \quad N_e \sim [\text{TSkewN}(\mu_e, \sigma_e, \alpha, a = 0, b = \infty) + 0.5], \quad N_\gamma = N_q - N_e, \quad (2)$$

where p_r is the recombination probability, μ_e and σ_e are the mean and standard deviation of the skew-normal distribution, and α is the shape parameter controlling skewness.

From N_γ , the number of photons, and N_e , the number of electrons, the S1 signal simulated using the following detector model:

$$N'_\gamma \sim \text{B}\left(N_\gamma, \frac{g1}{p_{DPE}}\right), \quad N''_\gamma \sim \text{B}(N'_\gamma, p_{DPE}), \quad S1 \sim \mathcal{N}\left(N'_\gamma + N''_\gamma, \sigma_{PS}\sqrt{N'_\gamma + N''_\gamma}\right) \quad (3)$$

where $\frac{g1}{p_{DPE}}$ is the probability of detecting a photon, p_{DPE} is the probability of a sensor detecting a doubled signal from a single scintillation photon due to double photoelectron emission [23], and σ_{PS} is

the the sensor energy resolution. The S2 detector model is instead given by

$$N'_e \sim \text{B}(N_e, p_e), \quad S2 \sim \mathcal{N}\left(N'_e, \sigma_e \sqrt{N'_e}\right), \quad (4)$$

where p_e is the probability that an electron is successfully drifted to the top of the liquid xenon and then extracted into the gas phase, and σ_e is the spread of the S2 signal. It can be seen that many of the latent parameters are discrete, posing a problem for direct implementation into a differentiable model.

3 Methods

3.1 Continuous Approximation for the Binomial Distribution

Approximations can be used to work around the aforementioned discrete random variables. The normal distribution is a commonly used approximation, but it does not perform well when p is far away from 0.5 as there is no skewness. This can be remedied using either transformations of approximating distributions such as a normal distribution [24], or a skew-normal distribution [25]. Finding the skew-normal shape parameters that approximate a binomial distribution, however, requires one to solve an equation where the solution does not have a simple functional form. This makes it tricky to use the skew-normal distribution as an automatic-differentiation-compatible approximation of the binomial distribution. To address this issue, we use a mixture of a normal distribution and a beta distribution scaled by N , the number of trials. We believe that this is the first time a beta-normal mixture distribution has been used as a continuous approximation of a binomial distribution to obtain a differentiable likelihood.

We match the first two moments of the distributions to determine the correct parameters of the normal and beta distributions that approximate a given binomial distribution. If a mixture distribution comprises distributions with equal mean, then the higher-order central moment of the mixture distribution is simply given by a weighted sum of the moments of the constituent distributions. This can be derived as such:

$$\mathbb{E}[(X - \mu)^n] = \int_{-\infty}^{\infty} (x - \mu)^n \sum_{i=0}^k w_i p_i(x) = \sum_{i=0}^k w_i \int_{-\infty}^{\infty} (x - \mu)^n p_i(x) = \sum_{i=0}^k w_i \mathbb{E}[(X_i - \mu)^n], \quad (5)$$

where X is a random variable with the probability density $f(x) = \sum_{i=0}^k w_i p_i(x)$. Thus, if the normal and beta distributions match the mean and variance of the binomial distribution, any mixture of the two will match the first two moments of the binomial distribution. Similarly, the non-central moments of a mixture distribution can also be given by a weighted-sum of the constituents.

The mean and variance of binomial distribution $\text{B}(n, p)$ are np and $np(1-p)$, respectively [26]. This corresponds to the normal distribution $\mathcal{N}[\mu = np, \sigma = \sqrt{np(1-p)}]$. As the beta distribution is the conjugate distribution of the binomial distribution [27], one might intuitively expect the corresponding Beta distribution to be $\text{Beta}(\alpha, \beta)$, where $\alpha = np$, and $\beta = n(1-p)$. However, it turns out that this distribution does not match the variance of the binomial distribution. Instead, the matching beta distributions has the following parameters:

$$\alpha' = \alpha - \frac{\alpha}{\alpha + \beta}, \quad \beta' = \beta - \frac{\beta}{\alpha + \beta}. \quad (6)$$

The next step is to match the third standardised moment of the distributions, given by

$$\gamma = \mathbb{E}\left[\left(\frac{X - \mu}{\sigma}\right)^3\right] = \frac{\mathbb{E}[X^3] - 3\mu\sigma^2 - \mu^3}{\sigma^3}. \quad (7)$$

It can be seen that the skewness of a distribution does not depend on the scale of a distribution, and thus the following derivation can be carried out without needing to additionally account for the scale of the beta distribution. As $\mathbb{E}[X^3]$ can also be obtained via a weighted sum for a mixture distribution, and the means and variances are already matched, we can simply focus on this raw moment. The skewness of a normal distribution is zero; as such, $\mathbb{E}[X_{\mathcal{N}}^3] = 3\mu\sigma^2 + \mu^3$. The skewness of $\text{Beta}(\alpha', \beta')$ is $\gamma_\beta = [2(\beta' - \alpha')\sqrt{\alpha' + \beta' + 1}]/[(\alpha' + \beta' + 2)\sqrt{\alpha'\beta'}]$ [26]. The skewness of a binomial distribution is given by $\gamma_{\text{Binom}} = \frac{1-2p}{\sqrt{Np(1-p)}}$ [26].

We can thus find the relative weights by solving the following system of equations for the weight w , where w is the weight for the beta distribution in the mixture, and correspondingly the weight for the normal distribution is $1 - w$:

$$\frac{x - 3\mu\sigma^2 - \mu^3}{\sigma^3} = \frac{2(\beta' - \alpha')\sqrt{\alpha' + \beta' + 1}}{(\alpha' + \beta' + 2)\sqrt{\alpha'\beta'}} \cdot \frac{1 - 2p}{\sqrt{Np(1-p)}} = \frac{wx + (1-w)(3\mu\sigma^2 + \mu^3) - 3\mu\sigma^2 - \mu^3}{\sigma^3}. \quad (8)$$

The solution is given by $w = \frac{\alpha' + \beta' + 2}{2(\alpha' + \beta' + 1)}$.




Figure 1: Comparison of the performance of the beta-normal mixture distribution introduced in section 3.1 versus that of a normal distribution (top) and a skew-normal distribution (bottom), when approximating a binomial distribution. A comparison using the earth mover’s distance is shown on the left, and one using the maximum absolute error is shown on the right. Blue and red indicate regions of the binomial distribution parameter space where the beta-normal mixture distribution outperforms and under-performs, respectively. The leftmost corners of the bottom plots are empty because for those regions of parameter space, we were unable to obtain numerical solutions for the skew-normal distribution.

We now compare the performance of the beta-normal mixture distribution introduced in this work, a skew-normal distribution, and a normal distribution when approximating a binomial approximation using both the maximum absolute error [25, 28] and the earth mover’s distance [29]. This is shown in fig. 1. We can see that the normal distribution is a poorer approximation for the binomial distribution than the beta-normal mixture distribution, however, there are regions of relative over- and under-performance in the parameter space when comparing the mixture and the skew-normal distributions. The chief advantage of the mixture distribution is thus the simple closed form of the distribution parameters, enabling the implementation of a model that is fully compatible with automatic differentiation in JAX [4].

3.2 Model Implementations

A graphical representation of the model introduced in this work, termed bayesNEST, is shown in fig. 2. This model is implemented in `numpyro` [8, 9]. The NEST emissions model (see section 2.2) is represented within the plate, as variables within the NEST emissions model are repeated for each event. μ_I represents the ^{125}I rate (see section 4.2) and is not used for the line search described in section 4.1. μ_{sig} is the signal, and μ_{bg} is the rate of background events. N_{obs} is the total number of events observed, and is given a Poisson likelihood where $\lambda = \sum_i \mu_i$. E has a prior parameterised by the signal rate parameters. In the case of a line search, the prior corresponds to a mixture distribution with two components: a uniform distribution and a normal distribution, corresponding to the uniform background and the signal respectively. The weights of the mixture distribution are given by the ratio of μ_{bg} to μ_{sig} , normalised such that the weights sum to 1. For the double-electron capture measurement, another normal distribution is added, corresponding to the ^{125}I background (see section 4.2). Mono-energetic components are represented using normal distributions with standard deviations much smaller than the energy resolution, instead of a delta function, so that the model has a smooth and continuous derivative.

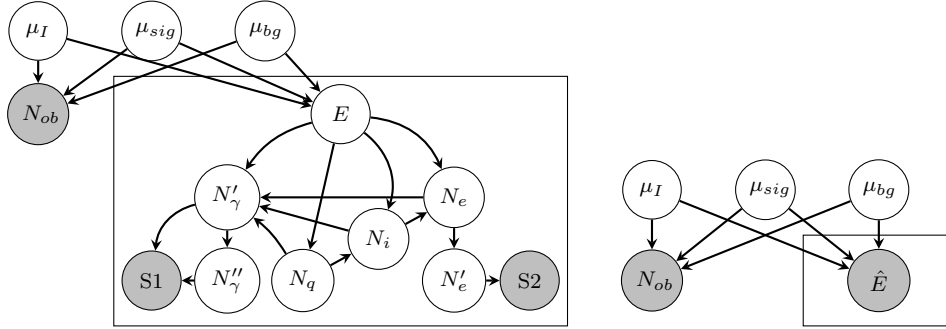


Figure 2: Left: Graphical representation of bayesNEST model. A plate, indicated by the rectangle, is used to represent the parts of the model that represent the emission model described in section 2.2. Right: Graphical representation of the model for energy spectrum fit. The only variable within the plate is the estimated energy. Observed variables are shown in grey.

We benchmark inference with the bayesNEST model (see fig. 2) against an energy spectrum fit. For this model, the observable is an estimator of energy, $\hat{E} = W(S1/g_1 + S2/g_2)$, where W is the workfunction, g_1 is the average signal amplitude observed from a single emitted photon, and g_2 is the average signal amplitude observed from a single S2 electron. The prior for \hat{E} is defined as a mixture distribution where the normally-distributed components have a larger standard deviation, corresponding to the standard deviation of the estimator \hat{E} given a mono-energetic component. This energy resolution is estimated for each mono-energetic component by generating 5×10^4 events using the NEST simulator, computing \hat{E} for each event, and then computing the corrected sample standard deviation.

4 Results

4.1 Mono-energetic Signal Search with Known Background Shape




Figure 3: Left: 90% credible intervals for the signal rate parameter for the mono-energetic signal search. 20 intervals are computed for each of $\mu_{true} \in \{0, 20, 40\}$. Right: 90% credible intervals for the signal rate parameter for the double-electron capture search. 20 intervals are computed for each of $\mu_{bg} \in \{75, 100, 150\}$. Intervals computed using bayesNEST are shown in dark blue, and those computed from a spectral fit are shown in light blue. The mean from all samples for each set of 20 runs is shown in red and pink for bayesNEST and for the spectral fit, respectively.

We first compare the difference between inference using a simple spectral fitting model and the bayesNEST model (see section 3.2) using a search for a mono-energetic signal with a flat background. This simulates searches for dark matter or physics beyond the standard model using the electronic recoil channel in liquid xenon detectors, such as those conducted in [16] and [17].

We simulated background data with a flat spectrum between 50 keV and 350 keV and a rate of 50 events over the interval using NEST version 2.3.11 [22] and nestpy [30]. A search for a mono-energetic peak at an energy of 200 keV is conducted on events selected to have estimated energies of between 100 keV and 300 keV. The NUTS sampler implemented in numpyro [9, 8] is used. This is repeated 20 times each for true signal rates (μ_{true}) of 0, 20, and 40 events. These experiments were run using 8 CPU

Table 1: Left: The mean 90% credible intervals from bayesNEST ($\mathbb{E}[\sigma_{90\%}]$), spectral fitting ($\mathbb{E}[\sigma'_{90\%}]$), and the mean relative difference in interval size ($\mathbb{E}[1 - \sigma'_{90\%}/\sigma_{90\%}]$) for the mono-energetic signal search. Right: The mean 90% credible intervals from bayesNEST ($\mathbb{E}[\sigma_{90\%}]$), spectral fitting ($\mathbb{E}[\sigma'_{90\%}]$), and the mean relative difference in interval size ($\mathbb{E}[1 - \sigma'_{90\%}/\sigma_{90\%}]$) for the double-electron capture search.

μ_{true}	$\mathbb{E}[\sigma_{90\%}]$	$\mathbb{E}[\sigma'_{90\%}]$	$\mathbb{E}\left[1 - \frac{\sigma'_{90\%}}{\sigma_{90\%}}\right]$	μ_{bg}	$\mathbb{E}[\sigma_{90\%}]$	$\mathbb{E}[\sigma'_{90\%}]$	$\mathbb{E}\left[1 - \frac{\sigma'_{90\%}}{\sigma_{90\%}}\right]$
0	4.8 ± 0.2	4.9 ± 0.2	$3.6 \pm 0.9\%$	75	34.6 ± 0.6	35.0 ± 0.6	$0.9 \pm 0.7\%$
20	16.1 ± 0.1	16.4 ± 0.1	$1.6 \pm 0.4\%$	100	35.2 ± 0.6	35.8 ± 0.6	$2.0 \pm 0.7\%$
40	22.2 ± 0.1	22.3 ± 0.1	$0.5 \pm 0.3\%$	150	37.8 ± 0.7	38.5 ± 0.7	$1.9 \pm 0.8\%$

cores, and took less than 1 day. The results are shown on the left in fig. 3 and table 1. We can see that inference using bayesNEST has reduced variance, and that the mean is closer to the true value μ_{true} .

4.2 Searching for Rare Nuclear Decays

We then compare the difference between the two methods of inference by simulating a search for the radioactive decay of ^{124}Xe via double-electron capture. This is a challenging problem, as in addition to a smooth background component, there exists an ^{125}I background which is also mono-energetic at 67.3 keV, very close to the signal we are interested in, at 64.3 keV [14].

For this test, we fix the signal rate to be 50 events and the ^{125}I background rate to be 5 events when simulating data. We consider 3 scenarios with differing rates of the flat background (μ_{bg}): 75 events, 100 events, and 150 events. This is defined on an energy range of 40 keV to 110 keV, and for inference, data with estimated energies of between 50 keV and 100 keV are selected. These experiments were run using 8 CPU cores and took less than 1 day to complete as well. The results are shown on the right in fig. 3 and table 1. As with the results in section 4.1, we can see that inference with bayesNEST produces results with lower uncertainty, and the mean over 20 runs is closer to the true signal rate of 50 events.

5 Discussion

Inference in liquid xenon dark matter experiments has largely been based on the use of estimators of event energy [14, 16, 17], or comparisons of data with empirical distributions from simulated data [31, 18]. Prior work on computing explicit event-by-event likelihoods [19] has relied on marginalising a discrete likelihood over a large array, requiring computationally and memory intensive enumeration of discrete latent variables. In this work, we demonstrate the use of continuous approximations to construct such a likelihood, alleviating the need for the marginalisation procedure, and enabling inference using gradient-based inference methods, such as variational inference [5, 6] and the No-U-Turn Sampler (NUTS) [7].

The approximate procedure introduced in this work produces improvements over inference using an energy estimator, offering percent-level improvements in the uncertainty of inferred parameters. In the background-dominated regime, sensitivity is asymptotically proportional to signal/ $\sqrt{\text{background}}$ [32]; as such, if both signals and background event rates are proportional to the detector target mass, the sensitivity scales as the square-root of exposure. A 4% improvement such as that demonstrated in section 4.1 would thus correspond to a 7% increase in exposure. The high cost of xenon [20] means that even small improvements from analysis or inference technique can represent very cost-effect investments.

In addition to improvements in inference, the use of explicit likelihoods allow for inference to easily incorporate additional information if available. For example, N'_γ can be directly estimated by counting pulses in photosensors independently of the S1 amplitude [33, 34]. This can be easily incorporated into the model introduced in this work, either by directly making N'_γ in the model an observable, or by introducing a new observable depending on N'_γ while modelling for detection efficiency. Discrete measurements from pulse counting are often better at lower energies and continuous measurements based on integration of waveforms are often better at higher energies. The ability to integrate discrete and continuous measurements is also invaluable, as it would allow for a single unified scheme that can make use of both sources of information, without having a sharp transition point between the two methods that might complicate inference or data analysis.

Acknowledgements

This work was supported by the National Science Foundation under Grant No. 2046549.

References

- [1] Cranmer K, Brehmer J, Louppe G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*. 2020 May;117(48):30055–30062. Available from: <http://dx.doi.org/10.1073/pnas.1912789117>.
- [2] Baydin AG, Heinrich L, Bhimji W, Munk A, Gram-Hansen B, Louppe G, et al. Efficient Probabilistic Inference in the Quest for Physics Beyond the Standard Model. In: *Neural Information Processing Systems*; 2018. p. 5459–5472. Available from: <https://api.semanticscholar.org/CorpusID:49903021>.
- [3] Baydin AG, Shao L, Bhimji W, Heinrich L, Meadows L, Liu J, et al. Etalumis: bringing probabilistic programming to scientific simulators at scale. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. SC '19*. ACM; 2019. p. 1–24. Available from: <http://dx.doi.org/10.1145/3295500.3356180>.
- [4] Bradbury J, Frostig R, Hawkins P, Johnson MJ, Leary C, Maclaurin D, et al. JAX: composable transformations of Python+NumPy programs; 2018. Available from: <http://github.com/google/jax>.
- [5] Hoffman MD, Blei DM, Wang C, Paisley JW. Stochastic variational inference. *J Mach Learn Res*. 2012;14:1303–47. Available from: <https://api.semanticscholar.org/CorpusID:5652538>.
- [6] Wingate D, Weber T. *Automated Variational Inference in Probabilistic Programming*; 2013.
- [7] Hoffman MD, Gelman A. The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J Mach Learn Res*. 2011;15:1593–623. Available from: <https://api.semanticscholar.org/CorpusID:12948548>.
- [8] Bingham E, Chen JP, Jankowiak M, Obermeyer F, Pradhan N, Karaletsos T, et al. Pyro: Deep Universal Probabilistic Programming. *J Mach Learn Res*. 2019;20:28:1–28:6. Available from: <http://jmlr.org/papers/v20/18-403.html>.
- [9] Phan D, Pradhan N, Jankowiak M. Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro. *arXiv preprint arXiv:1912.11554*. 2019.
- [10] Aprile E, et al. Projected WIMP sensitivity of the XENONnT dark matter experiment. *JCAP*. 2020;11:031.
- [11] Akerib DS, et al. The LUX-ZEPLIN (LZ) Experiment. *Nucl Instrum Meth A*. 2020;953:163047.
- [12] Workman RL, et al. Review of Particle Physics. *PTEP*. 2022;2022:083C01.
- [13] Aalbers J, et al. A next-generation liquid xenon observatory for dark matter and neutrino physics. *J Phys G*. 2023;50(1):013001.
- [14] Aprile E, et al. Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T. *Nature*. 2019;568(7753):532–5.
- [15] Aprile E, et al. Double-Weak Decays of ^{124}Xe and ^{136}Xe in the XENON1T and XENONnT Experiments. *Phys Rev C*. 2022;106(2):024328.
- [16] Aprile E, et al. Excess electronic recoil events in XENON1T. *Phys Rev D*. 2020;102(7):072004.
- [17] Aprile E, et al. Search for New Physics in Electronic Recoil Data from XENONnT. *Phys Rev Lett*. 2022;129(16):161805.
- [18] Aprile E, et al. XENON1T dark matter data analysis: Signal and background models and statistical inference. *Phys Rev D*. 2019;99(11):112009.
- [19] James RS, Palmer J, Kaboth A, Ghag C, Aalbers J. FlameNEST: explicit profile likelihoods with the Noble Element Simulation Technique. *JINST*. 2022;17(08):P08012.
- [20] Akerib DS, et al. LUX-ZEPLIN (LZ) Conceptual Design Report. 2015 9.

- [21] Aalbers J, et al. First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment. *Phys Rev Lett.* 2023;131(4):041002.
- [22] Szydagis M, Brown E, Carrara N, Kamaha AC, Kozlova ES, McKinsey DN, et al.. Noble Element Simulation Technique. Zenodo; 2023. Available from: <https://doi.org/10.5281/zenodo.7577399>.
- [23] Faham CH, Gehman VM, Currie A, Dobi A, Sorensen P, Gaitskell RJ. Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes. *JINST.* 2015;10(09):P09010.
- [24] Gebhardt F. Some Numerical Comparisons of Several Approximations to the Binomial Distribution. *Journal of the American Statistical Association.* 1969;64(328):1638-46. Available from: <http://www.jstor.org/stable/2286095>.
- [25] Chang CH, Lin JJ, Pal N, Chiang MC. A Note on Improved Approximation of the Binomial Distribution by the Skew-Normal Distribution. *The American Statistician.* 2008;62(2):167-70. Available from: <http://www.jstor.org/stable/27643999>.
- [26] Walck C. Hand-book on statistical distributions for experimentalists; 1996.
- [27] Hoff PD. A first course in Bayesian statistical methods Peter D. Hoff. 1st ed. Springer texts in statistics. New York, N.Y: Springer; 2009.
- [28] Schader M, Schmid F. Two Rules of Thumb for the Approximation of the Binomial Distribution by the Normal Distribution. *The American Statistician.* 1989;43(1):23-4. Available from: <https://www.tandfonline.com/doi/abs/10.1080/00031305.1989.10475601>.
- [29] Ramdas A, Trillos NG, Cuturi M. On Wasserstein Two-Sample Testing and Related Families of Nonparametric Tests. *Entropy.* 2015;19:47. Available from: <https://api.semanticscholar.org/CorpusID:7725237>.
- [30] Farrell S, Tunnell C, Angevaare JR, Rischbieter GRC, Carrara N, mszydagis, et al.. NESTCollaboration/nestpy: Sync with NESTv2.3.12beta. Zenodo; 2023. Available from: <https://doi.org/10.5281/zenodo.7552304>.
- [31] Aprile E, et al. Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys Rev Lett.* 2018;121(11):111302.
- [32] Cowan G, Cranmer K, Gross E, Vitells O. Asymptotic formulae for likelihood-based tests of new physics. *Eur Phys J C.* 2011;71:1554. [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)].
- [33] Akerib DS, et al. Tritium calibration of the LUX dark matter experiment. *Phys Rev D.* 2016;93(7):072009.
- [34] Akerib DS, et al. Results from a search for dark matter in the complete LUX exposure. *Phys Rev Lett.* 2017;118(2):021303.