

Dark Matter Search with XENON1T

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Agnietenkapel
op dinsdag 19 juni 2018, te 10:00 uur

door

Jelle Aalbers

geboren te Apeldoorn.

| | | |
|----------------|---|--|
| Promotores: | Prof. dr. M.P. Decowski Prof. dr. S.C.M. Bentvelsen | Universiteit van Amsterdam Universiteit van Amsterdam |
| Overige Leden: | dr. G. Bertone Prof. dr. J. Conrad Prof. dr. A.P. Colijn Prof. dr. ir. P.J. de Jong dr. M.E.J. Postma Prof. dr. W. Verkerke dr. I.B. van Vulpen | Universiteit van Amsterdam Stockholm University Universiteit Utrecht Universiteit van Amsterdam Nikhef Universiteit van Amsterdam Universiteit van Amsterdam |

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Copyright © 2018, Jelle Aalbers
ISBN 9789462339842

Cover created using neural style transfer. Original pictures: XENON1T cryostat and support structure viewed from the top of the water tank (courtesy of Sander Breur) and M51 ‘whirlpool’ galaxy (courtesy of the Hubble Space Telescope). Style network trained on *Rain Princess* from Leonid Afremov. Created with `fast-style-transfer` from Logan Engstrom.

Printed by Gildeprint

Contents

| | |
|--|------------|
| S Dark matter search with XENON1T (summary) | 6 |
| S.1 Outline | 11 |
| 1 Dark matter | 12 |
| 1.1 Evidence | 12 |
| 1.2 Detection methods | 15 |
| 1.3 Particle models | 20 |
| 2 Detecting dark matter with xenon | 26 |
| 2.1 Nuclear recoil scattering rates | 27 |
| 2.2 Alternative dark matter signatures | 34 |
| 2.3 Particle detection with liquid xenon | 37 |
| 2.4 XENON1T and LXe TPCs | 47 |
| 3 Signals of the XENON1T detector | 54 |
| 3.1 Introduction | 54 |
| 3.2 Electronic noise | 57 |
| 3.3 Single photons and photoelectrons | 60 |
| 3.4 Single electrons | 68 |
| 3.5 S2s | 76 |
| 3.6 S1s | 83 |
| 4 The XENON1T event builder | 89 |
| 4.1 The XENON1T DAQ | 89 |
| 4.2 Trigger design | 92 |
| 4.3 Event builder | 101 |
| 5 First dark matter search results | 107 |
| 5.1 Introduction | 107 |
| 5.2 Data analysis | 109 |
| 5.3 Background model | 112 |
| 5.4 Results | 115 |
| 5.5 Outlook | 116 |

| | |
|--|------------|
| 6 Statistical data analysis | 117 |
| 6.1 Methodology | 117 |
| 6.2 ER response fitting | 121 |
| 6.3 Dark matter limit setting | 128 |
| 6.4 Recommendations for future runs | 141 |
| 7 Expanding the dark matter search | 143 |
| 7.1 Introduction | 143 |
| 7.2 Spin-dependent search | 147 |
| 7.3 Bremsstrahlung search | 150 |
| 7.4 S2-only search | 153 |
| 7.5 S2-only bremsstrahlung search | 169 |
| 7.6 Outlook | 171 |
| Bibliography | 172 |
| Z Donkere materie zoeken met XENON1T (samenvatting) | 186 |
| Acknowledgements | 192 |

Chapter S

Dark matter search with XENON1T (summary)

Physics aims to understand the universe at its most fundamental level, as precisely as possible. Recent decades saw great progress towards this goal, including the spectacular discoveries of the Higgs boson and gravitational waves. Surprisingly, however, we still do not understand what the majority of the universe is made of. For every kg of ordinary matter, our universe contains 5.4 kg of *dark matter* [1] – unknown particles that interact very weakly with atoms and light. Discovering the nature of dark matter is among the primary unsolved problems of particle physics.

Figure 1 summarizes our current understanding of dark matter and ordinary matter in the universe. The early universe was so hot that (radiation) pressure resisted gravitational collapse of overdensities of matter. In today’s cooler universe, matter can collapse to stars by radiating away the energy gained during collapse. Dark matter does not interact with radiation, and therefore collapsed already in the early universe, dragging matter along with it to form the seeds of galaxies, clusters, and the cosmic web of superclusters. However, since dark matter cannot lose energy by radiation, it does not collapse all the way to stars. Instead, it settles in dynamical (virial) equilibrium structures such as spherical haloes around galaxies.

Evidence for dark matter comes from several independent observations. Dark matter haloes around galaxies and clusters (bottom right in figure 1) are visible through their gravitational effects on matter (the motion of stars and galaxies) and light (gravitational lensing). The tug-of-war between gravity and pressure in the early universe (top left in figure 1) left its signatures in the cosmic microwave background (CMB), a faint leftover radiation liberated when the early universe cooled enough to become transparent, and carries unmistakable imprints of dark matter’s gravitational effects. More strikingly, the CMB shows that the early universe was so homogeneous that galaxies could not have formed without the seeds of structure laid down by dark matter

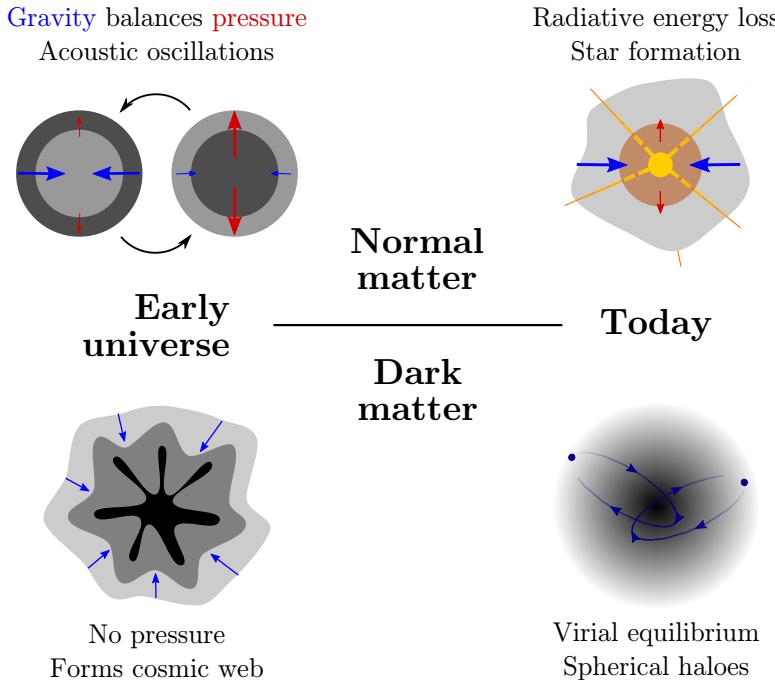


Figure 1: Simplified history of matter and dark matter in the universe.

(bottom left in figure 1). Without dark matter, the night sky would be dark, and there would be no one to see it.

This astrophysical and cosmological evidence shows dark matter interacts gravitationally with matter. However, gravity is too weak to measure individual particles (other than black holes), so experiments try to detect dark matter particles in one of three other ways:

- produce dark matter by smashing together ordinary matter at *colliders*,
- observe visible decay products of *annihilation* of dark matter and anti-dark matter in the universe, and
- search for tiny motions of normal atoms that result from elastic *scattering* ('bouncing') of passing dark matter particles.

Each method has limitations: colliders cannot make dark matter if that requires making a heavy intermediate particle first; annihilation searches must assume dark antiparticles exist and rule out ordinary astrophysical origins of any signal they see; while scattering searches cannot see dark matter light enough to bounce off atoms without transferring detectable amounts of energy.

This thesis describes a dark matter search with XENON1T, the world's most sensitive dark matter scattering detector. Figure 2 shows how XENON1T works. First, dark matter recoils off a xenon atom in the detector (specifically,

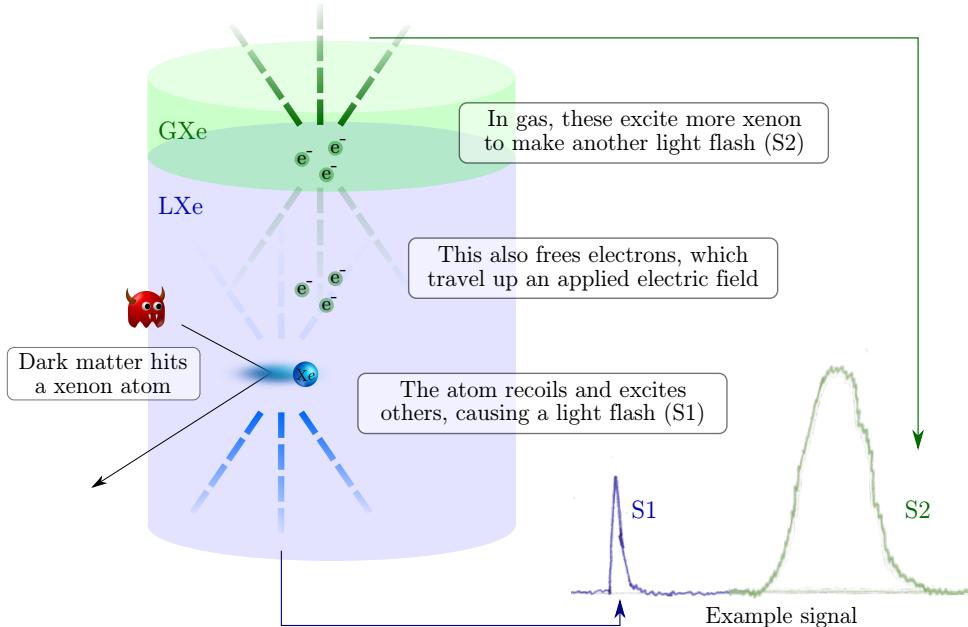


Figure 2: Operating principle of a dual-phase liquid-xenon (LXe) gaseous-xenon (GXe) time projection chamber (TPC) such as XENON1T.

its nucleus). The xenon atom, set in motion, excites some of its neighbours and ionizes others. Excited xenon atoms (after forming short-lived Xe_2 molecules) emit photons which are detected by sensitive light sensors called photomultipliers (PMTs) at the top and bottom of the detector. Electrons freed during ionization are pulled by an electric field towards a region with gaseous xenon (GXe). Here, a stronger electric field accelerates them enough to excite more xenon, causing a second, larger light flash (S2). The light distribution over the PMTs and the time between the S1 and S2 signals together indicate the 3D-position of the original dark matter interaction.

Unfortunately, many other processes cause light flashes in xenon, which XENON1T must protect itself against. To shield from cosmic rays, XENON1T operates deep underground in the *Laboratori Nazionali del Gran Sasso* under the Italian Appenines. To shield from natural background radioactivity, XENON1T is made from specially screened materials and submerged in a 10 m high instrumented water tank. Moreover, the dark matter search uses only events from the inner tonne of the three tonnes of liquid xenon used for XENON1T, so the other two tonnes act as additional shielding. The dominant remaining background is radioactive contaminants (in particular radon) released into the xenon from detector materials. Background from neutrino scattering will be dominant in future dark matter detectors. Most of these backgrounds cause recoils of *electrons* rather than whole xenon atoms, and can be separated (imperfectly) from dark-matter induced nuclear recoils

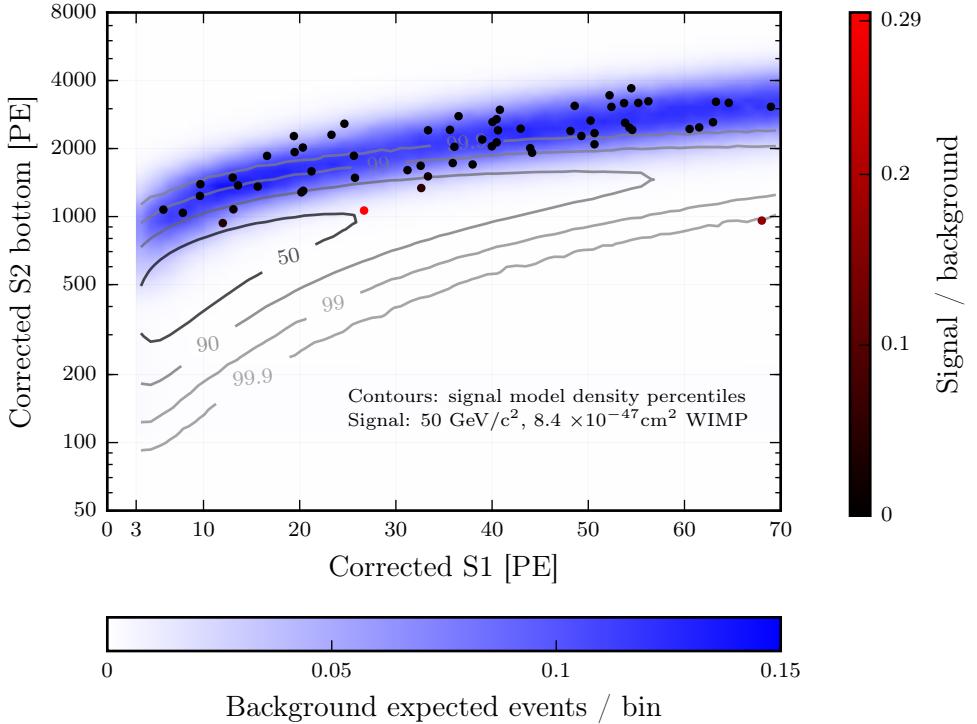


Figure 3: Dots: events found in XENON1T’s first dark matter search, as a function of S1 and S2 signal size measured in photoelectrons (PE). Blue shading shows where we would expect background events, contours where we would expect dark matter events (for a particular dark matter model indicated on the figure). The contours are density percentiles, so half the dark matter events would fall in the contour labeled ‘50’. The red shading of the dots indicates the ratio of signal/background likelihoods for each event: a value of 1 would mean an event came equally likely from background or dark matter (if the indicated dark matter model is true).

by their different ratio of S1 and S2 light.

Figure 3 shows the events XENON1T observed during its first dark matter search run of 34.2 days (described in detail in chapter 5). Most events are consistent with background models derived from calibration data, in which radioactive sources were placed near the detector on purpose. Even the most dark-matter like event (red dot in the center of figure 3) is consistent with background from (primarily) electronic recoils. One event at high S1 and low S2 (just above the center right in figure 3) defies simple explanation, but is also not in a region where dark matter signals are expected. Most likely this comes from a very rare background, such as accidental coincidence of unrelated S1 and S2 signals.

These first results show XENON1T’s background, in the region relevant for

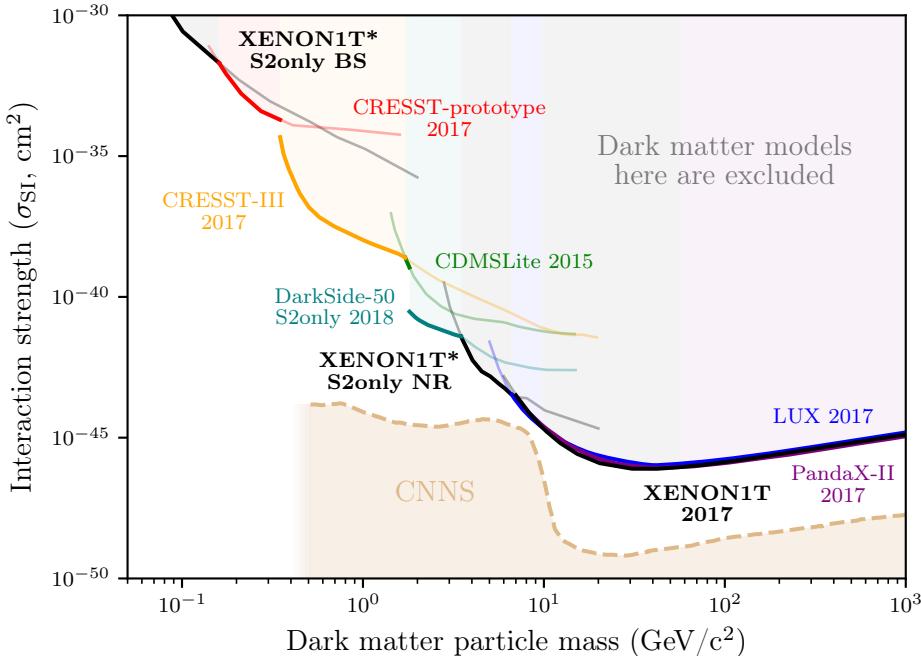


Figure 4: Constraints on the strength with which dark matter and normal matter interact for different possible dark matter masses. World-leading constraints are drawn as thick lines (with coloured shading above), others as thin lines. The black constraints (from XENON1T) are derived in this work, including the unpublished proof-of-concept results indicated by asterisks. The region labeled CNNS indicates models which are difficult to detect due to backgrounds from nuclear recoils of (primarily solar) neutrinos [4] (for the lowest WIMP masses, this is far beyond reach of current technology and was not computed).

dark matter searches, is lower than any experiment ever achieved. As we see no evidence for dark matter, we also set new limits on the strength with which dark matter and ordinary matter interact, shown in figure 4. Constraints from the LUX and PandaX-II experiments are similar, but these smaller experiments are nearing or past their end-of-life, as their larger backgrounds make longer searches futile. In contrast, XENON1T will soon release results of a year-long search, expected to set up to an order of magnitude stronger constraints – or discover dark matter.

In the coming years, XENON1T will be upgraded to XENONnT, containing ~ 3 times as much xenon [2]. The next generation of experiments will include detectors with more than 10 t of xenon such as DARWIN [3]. The coming decade will show a dark matter discovery, or falsify most currently established dark matter models. Let dark matter beware – XENON1T is coming for you!

S.1 Outline

This thesis is structured as follows. Chapter 1 discusses evidence for dark matter and some of the popular dark matter models that XENON1T is sensitive to. In chapter 2, we calculate the expected rate of dark matter interactions in XENON1T and review the detector physics that converts these to measurable S1 and S2 signals. Next, chapter 3 presents measurements and models of the S1 and S2 signal shapes that underpin the open-source signal processor (**pax**) and simulator (**fax**) developed during the research for this thesis [5]. One piece of this software, tasked with low-level monitoring and deciding which of XENON1T’s signals to record, is discussed in chapter 4.

Chapter 5 presents the first dark matter search results of XENON1T, as published in physical review letters [6]. We discuss technical details of the search in chapter 6, including recommendations for future science runs of XENON1T. Finally, chapter 7 re-examines the first dark matter search data to look for more general signatures of dark matter, including spin-dependent interactions, inelastic interactions (specifically, bremsstrahlung emission), and low-mass dark matter recoils for which only the S2 can be detected. While this last analysis is a proof-of-concept, it shows XENON1T is sensitive to light dark matter models usually considered out of reach for liquid xenon detectors.

Chapter 1

Dark matter

Dark matter is the missing link between our understanding of the universe at small and large scales. The evolution of galaxies, the motions of galaxy groups, particle physics in the early universe, the very laws of dynamics at large scales – all remain in doubt until the dark matter conundrum is resolved. The merest glimpse of a dark matter particle would send shockwaves through modern physics; as the first particle beyond the standard model, it may augur the end of this theory which has stood for generations. Perhaps it would be the first of a new family of particles, described by a new sector of particle physics, as is the hope of supersymmetry [7].

This chapter introduces the observational evidence for dark matter’s existence, the methods by which we try to detect it, and the results of these efforts so far for the most popular model, primarily targeted by XENON1T: weakly interacting massive particles (WIMPs).

1.1 Evidence

There is ample evidence for an additional source of gravity beyond standard model matter [8], from dynamics of galaxies and clusters (section 1.1.1) and especially from cosmology (section 1.1.2). Dark matter (DM), broadly understood as whatever generates this extra gravity, cannot be removed from our cosmic inventory without major alterations to physical theories. Some scientists consider modifications of general relativity as a possible explanation for some observations generally attributed to dark matter, so less dark matter would need to be postulated [9, 10]. While modified gravity theories offer a simple alternative account of galactic dynamics (especially rotation curves), they are less successful for galaxy clusters, and show dramatic discrepancies when confronted with contemporary high-precision cosmology data [11]. Additionally, modifications of general relativity are increasingly constrained by high-precision tests of general relativity, such as those arising from the recently detected gravitational waves [12, 13, 14].

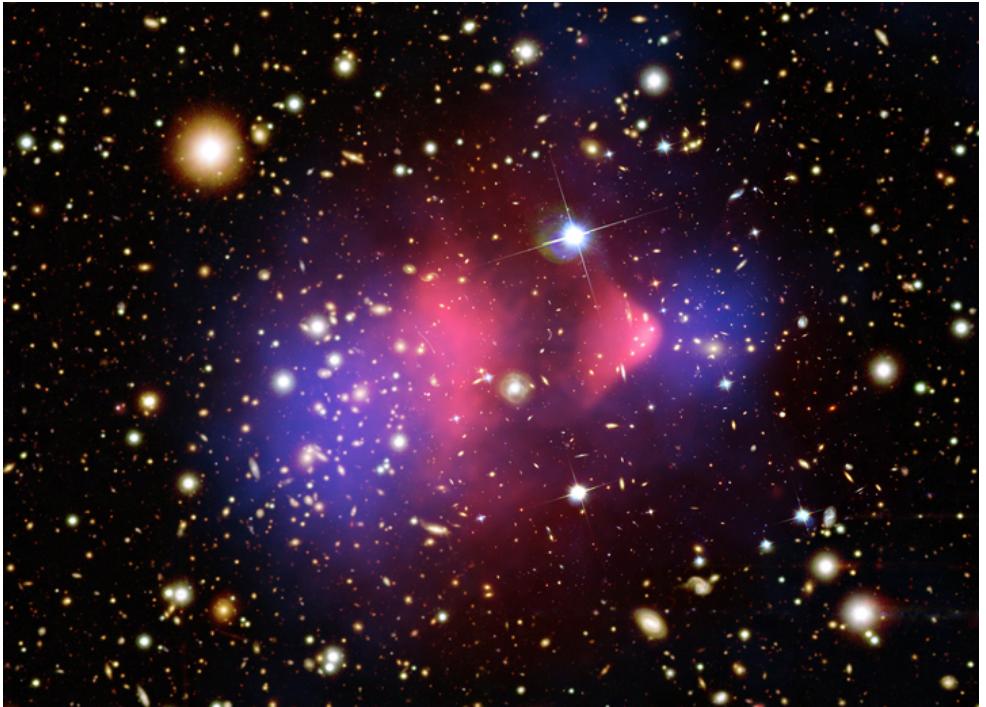


Figure 1.1: From [20] (based on [21]): The Bullet cluster of galaxies as seen by the Hubble and Magellan optical telescopes, with the density of hot gas as seen by the Chandra X-ray observatory overlaid in pink, and the cluster's mass density inferred from gravitational lensing in blue. This shows the primary source of gravity in the cluster is something other than the intergalactic gas (or the visible galaxies).

1.1.1 Galaxies and clusters

Discrepancies in the dynamics of galaxies in clusters and the motions of stars inside galaxies have been known since the early twentieth century [15, 16, 17]. Contemporary astrophysical data clearly rules out a universe made only of luminous matter interacting through standard gravity – specifically, dynamical mass estimates (from rotation curves for spiral galaxies, velocity dispersion for elliptic galaxies or galaxy clusters, gravitational lensing for all of these) converge on estimates about five times higher than methods based on luminous matter (star counting for galaxies, X-ray spectroscopy for clusters) [18, 19]. For galaxies, gravitational methods locate most of the mass elsewhere than the luminous mass: in an approximately spherical halo extending far beyond the galaxy's stars. For galaxy clusters, whose baryonic mass is mostly in sparse intergalactic gas, evidence of such a discrepancy is more rare, yet dramatically exhibited in colliding clusters such as the Bullet Cluster (see figure 1.1).

Compared to the cosmological evidence for dark matter, the evidence from galaxies and clusters is much less specific. Postulating unobserved masses

to solve dynamical problems in astronomy has a long history, perhaps most famously exemplified by the prediction of Neptune from perturbations in the orbit of Uranus [22]. On the other hand, so do proposals to change the laws of dynamics, as with Einstein’s explanation of the precession of Mercury’s perihelion [23]. Even with contemporary data, galactic data is consistent with both dark matter and modified gravity models [24]. Moreover, the observed dynamical anomalies can also be explained (in principle) by postulating appropriate amounts of ordinary matter that is difficult to detect, such as cold dust. In fact, dark matter was but one of several hypotheses used to explain a diverse collection of anomalies until the rise of physical cosmology in the 1970’s [25].

1.1.2 Cosmology

Dark matter is an integral part of the standard Λ CDM model of cosmology. Cosmological evidence for dark matter comes from anisotropies in the cosmic microwave background and the clustering of matter in the universe today. The effect of dark matter on the latter is more dramatic, but the former especially has been measured with exquisite precision in recent years. As a result, we know the universe is flat, with a dark matter density of $\Omega_{\text{dm}}h^2 = 0.1197 \pm 0.0022$, about five times higher than that of baryonic matter [1] (in agreement with the astrophysical evidence).

Temperature anisotropies in the cosmic microwave background inform us about density perturbations at the time of recombination – when electrons and nuclei form neutral atoms and thus cease to couple strongly to the photons. The CMB is homogeneous to one part in $\mathcal{O}(10^{-5})$, so these density perturbations were very small, and evolved linearly [28]. A perturbation on a length scale λ causes an oscillation in the photon-baryon fluid with period λ/v_{sound} , transforming the initially scale-invariant angular power spectrum of perturbations into one with a distinctive series of peaks (see figure 1.2, black). The detailed interplay of radiation, matter, and dark matter determines the exact location and amplitude of the peaks [28]; at high multipole numbers (small scales) the oscillations are damped by photon movement during recombination itself.

The Λ CDM model can describe the CMB spectrum excellently with just six parameters (or slightly more, depending on which external constraints are included). The high-precision CMB measurements from Planck [1] pose a formidable challenge for alternative cosmologies – the relatively high third peak at $\ell \approx 800$, in particular, is not seen in models without dark matter, or in the few modified gravity models for which the CMB spectrum calculation has been performed (see figure 1.2).

Dark matter made an even larger impact on the universe as the seed of structure formation. While matter in the early universe resisted clumping due to strong self-interactions, dark matter was free to aggregate, forming

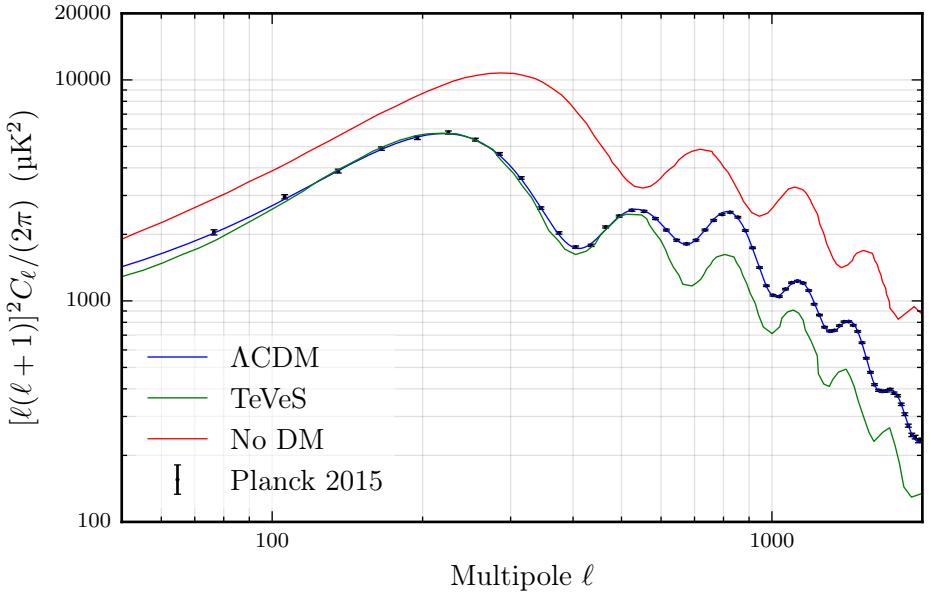


Figure 1.2: CMB temperature anisotropy measurements (black) and best-fit Λ CDM model (blue) from Planck [1]. The red line is a similar fit from which dark matter has subsequently been removed completely [26]; the green line is a fit of TeVeS, a leading modified gravity model, to WMAP data [27]. Distortions for the red and green lines at high ℓ are due to limited resolution on the figures of [26] and [27].

the seeds of the galaxies and galaxy clusters we see today. At increasingly smaller scales, the universe appears less and less homogeneous, as shown in figure 1.3. An array of experiments has validated the Λ CDM predictions over many orders of magnitude [29]. In contrast, replacing dark matter by baryonic matter yields the nowadays absurd prediction of a universe without any large-scale structures; modified gravity theories can form some structure, but still show very large disagreement with the data at large scales [30, 11].

1.2 Detection methods

There is currently no convincing evidence for any non-gravitational interaction of dark matter, let alone identification of individual dark matter particles. Several experimental efforts aim to change this.

Usually dark matter experiments are divided in ‘direct’ and ‘indirect’ detection experiments. Since ‘direct’ is a value-laden term, this is inevitably contested. The Bullet Cluster lensing results were heralded as “direct empirical proof” of dark matter [21]; others would only consider evidence of

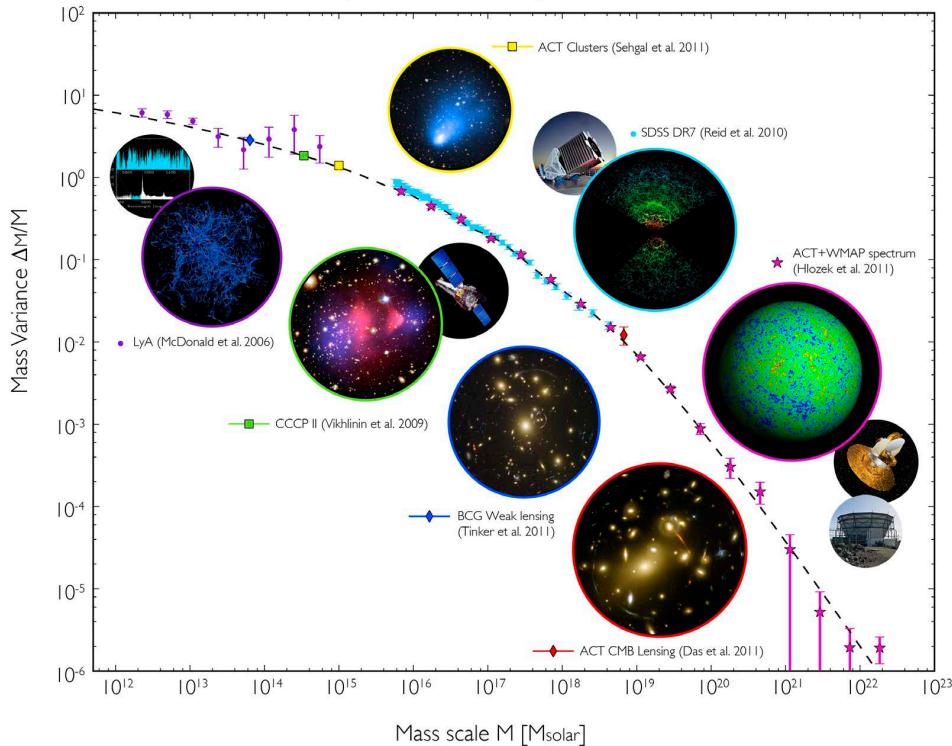


Figure 1.3: From [29]: Variations in the matter distribution ('clumpiness') at different scales in the universe. Here $\Delta_M/M = \sqrt{P(k)k^3/(2\pi^2)}$, with $P(k)$ the power spectrum of the matter distribution in the universe, and k the wavenumber corresponding to the indicated mass scale (via $M = (4/3)\pi\rho_m(\pi/k)^3$, with ρ the universe's mean matter density). For comparison, the mass of the milky way is $\sim 10^{12} M_\odot$, and the mass of the observable universe $1.7 \times 10^{24} M_\odot$. The Λ CDM prediction is shown by the dashed line. Observations are indicated by coloured markers, from low to high mass scale: hydrogen absorption lines in the intergalactic medium (LyA-forest, purple), weak lensing by the brightest galaxies in clusters (dark blue), X-ray observations of clusters (green), small-scale CMB fluctuations due to foreground galaxy clusters (kSZ effect, yellow), redshift surveys of galaxies (light blue), lensing of the CMB by foreground overdensities (red), and CMB anisotropy data (magenta, see also figure 1.2). Several datasets shown in this plot have already been superseded by measurements with even higher precision.

a non-gravitational interaction as direct detection; still others only count an individual dark matter particle interacting in an earth-based detector as direct evidence. A more physically meaningful distinction is the interaction by which an experiment searches for dark matter: gravitational, or otherwise. Non-gravitational experiments can be further categorized by the process they search for: production (colliders), annihilation (from the galactic center, halo, as well as other galaxies), scattering (in earth-based detectors), and scattering followed by annihilation (neutrinos from the Sun and Earth).

1.2.1 Gravitational searches

If dark matter consists of extended objects, we might detect these individually through their gravitational interaction. Ordinary dark astrophysical objects, such as stellar remnant black holes, are not dark matter candidates, as they would have to be formed from matter that was baryonic at the time of the CMB (and big-bang nucleosynthesis), running afoul of the cosmological observations discussed above.

The most popular extended dark matter candidates are primordial black holes. These are constrained by a wide variety of astrophysical studies, using e.g. gravitational lensing or the stability of galactic morphology. The most well-established bounds currently leave a few mass windows in which primordial black holes could make up all the dark matter; perhaps most interestingly $(25 - 100)M_\odot$ [31], which includes the recently observed black hole mergers [12, 32] at its lower end. However, these windows disappear when constraints with more astrophysical uncertainties are included [31].

Even if primordial black holes make up a small fraction of the dark matter, they would be important as a accretion seed for particle dark matter candidates [33]. Depending on the particle dark matter model, this can result in ultracompact minihaloes that would be prime targets for annihilation-based searches [34].

1.2.2 Production: colliders

Dark matter particles could be produced at colliders, where they would be detected as events with ‘missing’ energy-momentum, i.e. events in which the momentum of the detected decay products does not add up to zero for electron-positron colliders, or zero in the directions transverse to the beam for hadron colliders. The exact list of possible signatures (and their backgrounds from e.g. standard model neutrinos) is long, and depends greatly on the specific dark matter model assumed [8]; for a review, see [35]. Current colliders do not have the center of mass energy necessary to produce dark matter candidates at the upper end of the most popular mass range (see section 1.3.2), but can probe lower dark matter masses than scattering experiments, since scatters of low-mass dark matter deposit very little energy (see section 2.1.2).

A new ‘invisible’ particle discovered at a collider is not automatically dark matter; if it decays to regular matter in a few seconds (or a few billion years), detectors in collider experiments would not notice. On the other hand, collider searches are not affected by astrophysical backgrounds and associated uncertainties.

1.2.3 Annihilation: neutrinos, gamma and cosmic rays

Annihilation experiments look for standard-matter endproducts of DM-antiDM annihilation, and set limits on the thermally averaged annihilation cross-section $\langle\sigma_A v\rangle$ under present-day dark matter velocities.

The central challenge for these experiments is dealing with astrophysical backgrounds that produce competing signals. Experiments with directional sensitivity (gamma ray telescopes such as Fermi-LAT and HESS, neutrino telescopes such as IceCube and ANTARES) focus on regions with very little normal matter (such as galactic halos and dwarf spheroidal galaxies) or where particularly high dark matter excesses are predicted (such as the galactic center). Cosmic-ray experiments instead focus on particles (such as positrons) which have comparatively fewer astrophysical production mechanisms. There are many interesting observations in this field (such as excess positrons in cosmic rays, or excess gamma rays from the galactic center) that could be interpreted as hints of dark matter, but astrophysical explanations remain quite viable. For reviews, see [36, 37].

1.2.4 Scattering: earth-based detectors

Scattering experiments try to detect dark matter particles in a controlled detector on earth. They pack a high density of target matter in a tightly monitored volume, then shield it from as many background signals as possible. Dark matter particles in the most popular mass range (see section 1.3.2) produce signals in the keV-MeV range when recoiling off heavy nuclei. Recoils off electrons are highly suppressed except in special models (see section 2.2.1). Thus, scattering experiments primarily constrain the DM-nucleon cross-section σ_n (see section 2.1.4).

Elastic nuclear recoils deposit their energy through up to three channels, depending on the target material: heat (elastic scattering), charge (ionization) and light (scintillation). The most significant background in these experiments is often electronic recoils produced by background radioactivity (see section 1.2.5). Since recoiling ions have a much lower charge / mass ratio than recoiling electrons of a given energy, they generally produce less ionization and lose more energy through heat. Thus, many experiments are sensitive to multiple signal channels, to enable signal/background discrimination. For example, dual-phase noble element TPCs (see section 2.4), such as XENON1T, detect ionization and scintillation, while solid-state detectors generally detect heat and either ionization (e.g. CDMS) or scintillation (e.g. CRESST).

Due primarily to the Sun’s motion around the galaxy, the rate of DM-induced recoils will be anisotropic; and due to the earth’s motion around the sun, it will have an annual modulation. The extent of both effects depends highly on the dark matter mass and the energy threshold of a detector (see section 2.1.1). Directional sensitivity requires special detectors; while these traditionally use low-density materials, which are not optimal for stopping much dark matter, this field is undertaking significant R&D efforts [38].

In the early 2010’s, several scattering experiments – CoGenT[39], CRESST-II [40], CDMS-Si [41], DAMA and DAMA/LIBRA [42] – claimed evidence for low-mass ($\mathcal{O}(10\text{ GeV})$) WIMPs. These claims were in tension with each other and with results from competing experiments such as XENON100 [43] and LUX [44]. The CoGenT and CRESST-II excesses are nowadays understood as underestimated radioactive backgrounds, while the CDMS-Si excess has been undercut by the more sensitive SuperCDMS experiment by the same collaboration [45]. The claim from the NaI scintillation detectors DAMA and DAMA/LIBRA is based on a significant annual modulation signal, but in tension with null results from essentially all modern scattering experiments (even for highly specific dark matter models, see e.g. [46]).

Dark matter with very strong couplings to ordinary matter cannot be probed by ground-based and underground detectors, as this dark matter would lose most of its kinetic energy via scatters in the earth’s atmosphere (and for underground detectors, the rock overburden) [47]. However, such models are mostly excluded by other observations, including particle detectors on balloons and space stations, and the earth heating constraint discussed in section 1.2.5 below.

For reviews of DM-scattering experiments and detection mechanisms, see [48, 49, 50, 51]. Chapter 2 discusses expected DM-scattering rates, the technology of noble gas detectors, and XENON1T in particular.

1.2.5 Scattering and annihilation: neutrinos and heat

Celestial bodies such as the Sun and Earth may trap dark matter in their core, where they might annihilate in significant numbers. If the decay products include neutrinos, these could be detected in neutrino detectors such as IceCube or Super-Kamiokande.

To produce competitive limits, trapping dark matter particles gravitationally in Keplerian orbits would be insufficient: scattering off ordinary matter is required to slow the dark matter down enough to produce a dense core. If the thermally averaged annihilation cross-section $\langle\sigma_A v\rangle$ and DM-nucleon scattering cross-section σ_n are both large enough, capture and annihilation equilibrate within the lifetime of the system, and limits on neutrino fluxes can be converted to limits on both cross-sections.

Decay products other than neutrinos are converted to heat before they reach detectors. For the earth, the lack of a substantial new source of core

heating can be used to constrain strongly interacting dark matter models [52] out of reach of conventional scattering experiments.

1.3 Particle models

1.3.1 General models

Particle dark matter models span a vast space. We can divide them roughly in three classes, based on their balance between simplicity, range of validity, and independent motivation [53]:

- On the one extreme, there are *full theories* (such as models with supersymmetry or extra dimensions), usually additionally motivated by other puzzles in physics (such as the hierarchy problem or even quantum gravity). An individual theory usually has many degrees of freedom, and postulates many more phenomena than dark matter candidates. The space of all possible theories is limited only by our imagination.
- On the other extreme are *effective field theories* (EFTs), which capture all possible low-energy behaviours of more elaborate (but unknown) models, using a very restricted set of parameters. At energies higher or comparable to the scales of the unknown full theory, these models break down.
- In between are so-called *simplified models*, which suppose the simplest models for dark matter that could in principle be valid up to scales well beyond the reach of modern phenomenology. Generally this means postulating a dark matter particle and a single ‘mediator’ which interacts with it and (some) standard model particles. In the most basic models, the mediator is a standard model particle. Unlike full theories, these models have no independent motivation beyond explaining dark matter; like EFTs, however, they may arise as low-energy limits of full theories (but with a less restricted range of validity).

For presenting results of low-energy scattering experiments such as XENON1T, EFTs seems ideal, as long as the dark matter mass is not extremely small itself. In practice only two types of effective interactions (spin-dependent and spin-independent, see section 2.1.4) are commonly used for presenting results. In contrast, modern LHC dark matter searches are most sensitive to models for which EFTs do not apply at the relevant LHC energy scales [35], and instead present their results in terms of simplified models. Section 1.3.3 considers an example of a simplified model.

Dark matter models also differ in the mechanism explaining their production and survival. A great variety exists, with some exceedingly specific and complex mechanisms (e.g. for axions). Some of the best-studied models are the following:

- *Freeze-out*, or decoupling, is arguably the most popular mechanism. In the hot early universe, all particles are relativistic and presumed to be in local thermal equilibrium, freely converting between the various species. When the universe cools enough to make a particle species non-relativistic, its number density is reduced much more rapidly than the relativistic species, provided thermal equilibrium is maintained. This net annihilation continues until the annihilation rate, which is falling due to increasing rarefaction of the species, falls below the expansion rate of the universe. After this freeze-out time, the number density of the species is further diluted only by the expansion of space. The sooner a particle decouples, especially after the particle becomes nonrelativistic, the larger its number density today.
- *Matter-antimatter asymmetry* also starts from an initial thermal equilibrium between dark matter and normal matter, but assumes the annihilation halts because of an excess of particles over antiparticles. This is how standard model electrons and hadrons survived until today. Standard model neutrinos are expected to have frozen out early instead, so the number of relic neutrinos and antineutrinos could be similar (though a substantial neutrino-antineutrino asymmetry is also possible [54]).
- *Freeze-in* is an alternative mechanism, where dark matter never enters thermal equilibrium with ordinary particles. If dark matter couples very weakly to standard matter, but was also very rare in the early universe, most dark matter is produced when it is already out of equilibrium with standard matter [55, 56].

1.3.2 WIMP dark matter

Frozen-out, or thermal relic dark matter has traditionally received most attention in the literature, and is the most relevant for XENON1T.

The mass of thermal relics must be below $\mathcal{O}(100 \text{ TeV})$ by unitarity of the total S-matrix for dark matter particle-antiparticle scattering [57]. There are also several constraints on low-mass thermal relics, of varying degrees of strength:

- $m_\chi \lesssim 5 \text{ keV}$ would mean the dark matter stays relativistic too long to seed the large-scale structures observed in Lyman- α forest data [58];
- $m_\chi \lesssim \text{MeV}$ implies dark matter would be relativistic at the time of neutrino decoupling, which causes disagreements with CMB and big-bang nucleosynthesis observations [59], except in certain special models (e.g. [60]);
- $m_\chi \lesssim \text{GeV}$ would cause perturbations of the CMB by dark matter annihilation products, unless the dark matter decays mostly to neutrinos, or has a strongly velocity-suppressed (p-wave) annihilation [1, 61, 62];

- $m_\chi \lesssim 20$ GeV is in tension with absence of annihilation signals in Fermi-LAT’s observation of dwarf spheroidal galaxies [63]. As this limit is based on annihilation in the universe today, it requires an even stronger lack of velocity-dependence of annihilation.

This leaves a mass window for straightforward thermal relic models between $\mathcal{O}(\text{GeV})$ – or $\mathcal{O}(\text{MeV})$ for models with velocity-dependent annihilation – and $\mathcal{O}(100 \text{ TeV})$. Such particles are generally called weakly interacting massive particles (WIMPs).

A thermally-averaged annihilation cross section $\langle \sigma_A v \rangle$ of $3 \times 10^{-26} \text{ cm}^3/\text{s}$ is often stated as necessary to obtain the observed dark matter density through freeze-out. This is borne out in precise calculations for simple models up to a factor 2 [64], but in other models it is not accurate at all [65]. For precise calculations of the relation between $\langle \sigma_A v \rangle$ and $\Omega_\chi h^2$, numerical routines such as **DarkSUSY** or **micrOMEGAs** are generally used.

The ‘weak’ in WIMP does not have to refer to the standard model weak interaction – though the characteristic scale of weak interactions ($1/\sqrt{G_F} \approx 293 \text{ GeV}$) does lie right in the middle of the allowed mass window. Moreover, a new fermion with this mass scale and interacting via the weak force has, by dimensional arguments, roughly the correct $\langle \sigma_A v \rangle$ mentioned above to yield the observed dark matter density via thermal freeze-out [66], a fact sometimes called the *WIMP miracle*. This apparent coincidence, together with the appearance of suitable WIMP candidates in more complete theories of physics beyond the standard model [7], has caused substantial experimental and theoretical interest in WIMPs.

1.3.3 Discovery potential of XENON1T

To set the stage for XENON1T, it is useful to compare its physics reach to complementary constraints from other species of experiments. We have to assume a concrete dark matter model to even start such a comparison: at the least, a simplified model (see section 1.3.1). As mentioned, in the simplest of these models, the dark matter interacts with ordinary matter via a standard-model mediator: most commonly, the Z boson or the Higgs boson. This still leaves open a range of possible dark matter particle types, with different couplings to the mediator particle, for which the experimental constraints will be different.

As an example, consider a thermal-relic fermionic WIMP χ that couples to the standard model Z boson with an axial coupling:

$$\mathcal{L} = \mathcal{L}_{\text{SM}} + \mathcal{L}_\chi + g_{\chi a} a \bar{\chi} \gamma^\mu \gamma^5 \chi Z_\mu. \quad (1.1)$$

Here \mathcal{L}_{SM} is the standard model Lagrangian, \mathcal{L}_χ the Lagrangian for the free dark matter fermion χ , $g_{\chi a}$ the coupling constant, and $a = 1$ for a Dirac fermion and $1/2$ for a Majorana fermion.

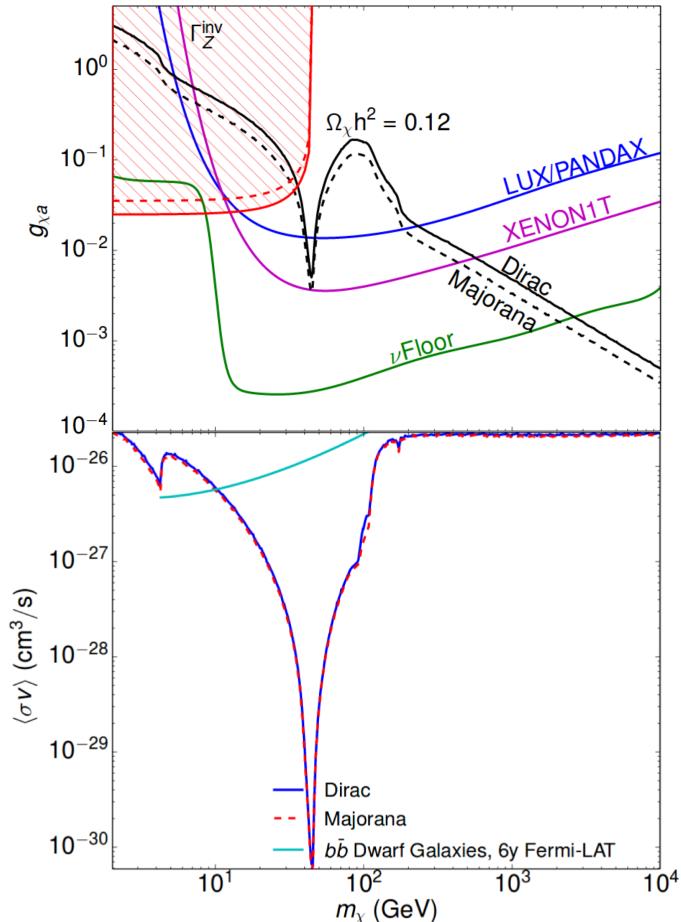


Figure 1.4: Adapted from [67]: experimental constraints on fermionic dark matter interacting with the standard model Z via a purely axial coupling. Top panel: constraints on the coupling constant $g_{\chi,a}$ for different dark matter masses from lack of invisible Z decays at LEP (red shaded region), direct detection experiments LUX and PandaX-2016 (blue line), and projected limits for both XENON1T [68] and a hypothetical DM-scattering detector with neutrino-floor sensitivity (green, see section 2.3.4.2). The black line indicates couplings that produce the observed relic density through freeze-out for a Dirac (solid) and Majorana (dashed) fermion. Bottom panel: constraints on the thermally averaged cross-section $\langle \sigma_A v \rangle$ at velocities relevant for annihilation-based searches. Note $g_{\chi,a}$ determines $\langle \sigma_A v \rangle$, since the model is otherwise fully specified. Models yielding the observed relic density lie on the solid dark blue and dashed red lines for Dirac and Majorana fermions, respectively. The light blue line shows the constraint on $b\bar{b}$ annihilation from Fermi-LAT observations of dwarf spheroidal galaxies [63] (the $\tau^+\tau^-$ annihilation limits in [63] are weaker for this model).

As shown in figure 1.4, the lowest mass WIMPs under this model are ruled out, most strongly by collider data, but also by Fermi-LAT’s annihilation constraints. Scattering experiments rule out most of the remaining mass ranges, except at high masses ($m_\chi \gtrsim 200 \text{ GeV}$) and a small region near $m_\chi = m_Z/2$. In the latter region, $\chi\bar{\chi}$ annihilation (with intermediary Z production) is resonantly enhanced, hence much lower couplings are needed to explain the observed relic abundance. Similar but smaller effects are seen when production of particular standard model particles becomes kinematically accessible, e.g. $\chi\bar{\chi} \rightarrow b\bar{b}$ at $m_\chi = m_b \approx 4 \text{ GeV}/c^2$. Overall, smaller masses result in a smaller phase-space factor in the cross-section, and thus require stronger couplings to produce the same relic abundance.

For this model, XENON1T will be able to discover WIMPs under this model if their mass is between 200 and 500 GeV, or near $m_Z/2$. Similar future experiments can probe masses $\lesssim 20 \text{ TeV}/c^2$; even higher masses will remain beyond the range of scattering experiments not using new technologies for mitigating the coherent neutrino nucleus scattering background (see section 2.3.4.2).

For other standard-model mediator models, the potential of XENON1T is shown in figure 1.5. Much of the parameter space is already ruled out, some of it will be out of reach for some time, and there are still significant regions where XENON1T and experiments in the near future can make a discovery [67, 69]. As we can see, XENON1T has most potential for high-mass, Higgs-mediated scalar or vector WIMPs.

Other simplified models postulate an additional new mediator coupling the dark and visible sectors. Such models have considerably more degrees of freedom, most importantly the mass and particle type of the new mediator, but also the couplings between the mediator and standard model particles. Complementary constraints on a generic class of such models are considered in [69]. Global constraints on supersymmetric models can be found in [70, 71].

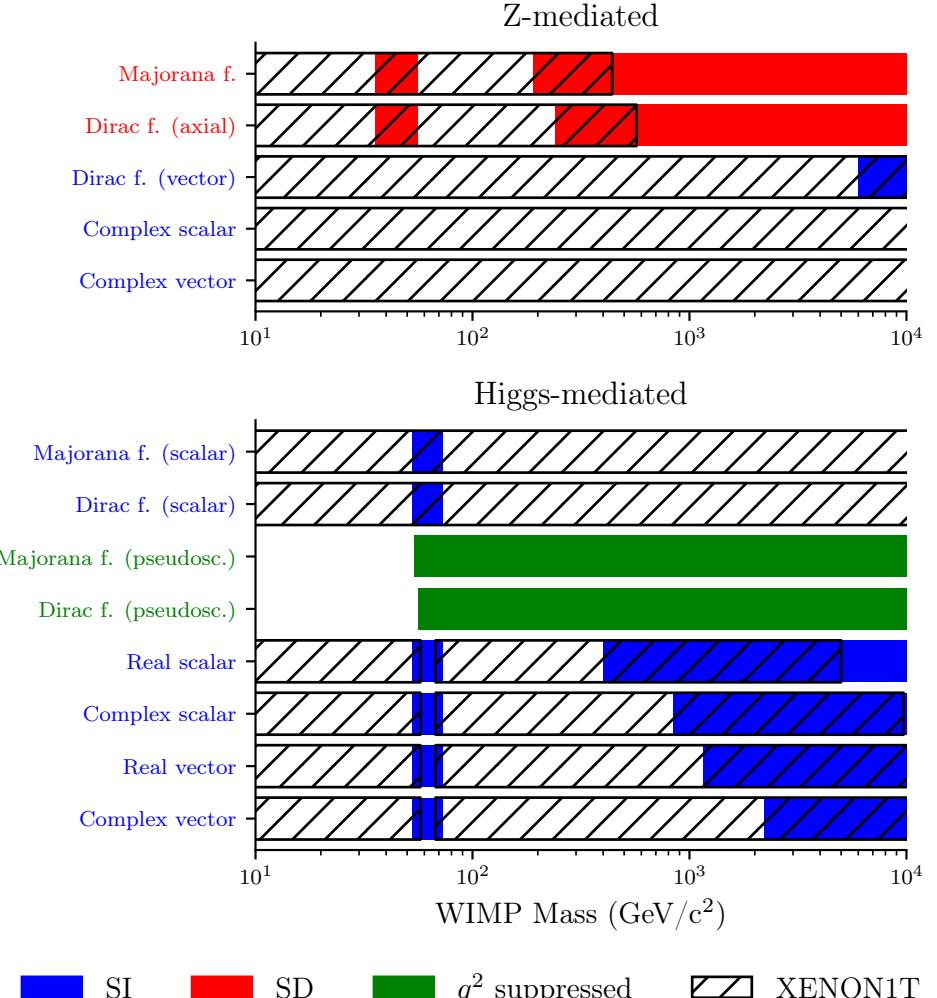


Figure 1.5: Projected XENON1T sensitivity to WIMP models with a standard-model mediator and a cross-section that yields the observed dark matter abundance via freeze-out. All data is from [67]; region boundaries should be considered approximate (especially around the $m_Z/2$ and $m_H/2$ poles). The dark matter particle type is indicated on the vertical axis, with the coupling type in parentheses when several elementary couplings are possible. Coloured mass ranges are allowed by experimental data from 2016; the colour code indicates the dominant scattering mechanism for each model, discussed in section 2.1.4 (SI: spin-independent, SD: spin-dependent, q^2 suppressed: suppressed by momentum transfer factors). Hatched regions will be probed by XENON1T. None of the models are allowed for masses below 10 GeV/c²; masses above 10 TeV/c² are not considered in [67].

Chapter 2

Detecting dark matter with xenon

This chapter introduces the technology of XENON1T: detecting dark matter (DM) scatters on earth using the noble element xenon. The structure of this chapter and subsequent chapters mirrors XENON1T's overall detection mechanism, shown in figure 2.1: WIMP collision rates (1 in fig. 2.1) are discussed in section 2.1, energy deposition in liquid xenon (2) in section 2.3; XENON1T's signal response (3) in chapter 3, and the XENON1T DAQ (4) in chapter 4. The remaining chapters of this thesis discuss the analysis of XENON1T's data.

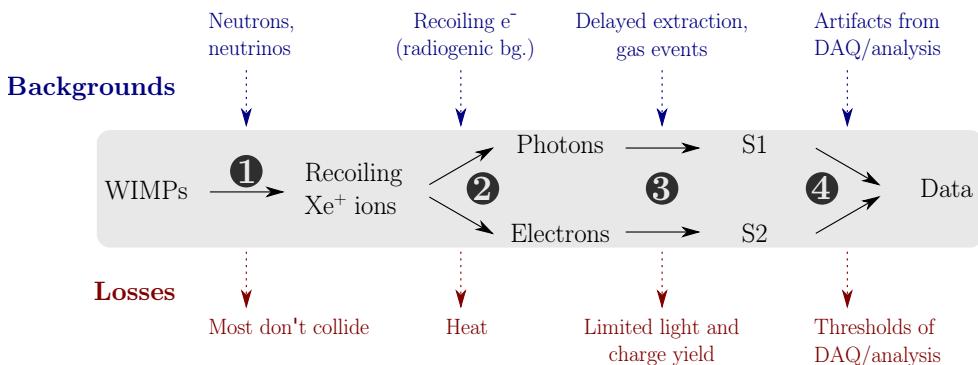


Figure 2.1: Dark matter detection in XENON1T (middle row). A WIMP collides with a xenon nucleus (1); the recoiling ion transfers its energy to carriers we can detect - scintillation light and ionization electrons (2); these produce signals in XENON1T (3); and finally these signals are digitized and interpreted (4). Each step introduces backgrounds (top row) and losses (bottom row), which are discussed later in the text.

2.1 Nuclear recoil scattering rates

In particle physics, the total interaction cross-section σ with a single (target) particle relates the rate r of interactions to the number flux of (dark matter) particles:

$$r = nv\sigma, \quad (2.1)$$

with n the (DM) particle number density and v the (DM) particle speed (here assumed equal for all incoming particles).

A DM-scattering experiment measures the differential rate of DM interactions per unit recoil energy E_R and detector mass. This is simply eq. 2.1, divided by the mass of a target nucleus m_N , with differential rates and cross sections taking the place of r and σ , and with an integral over v since dark matter has a distribution of velocities:

$$\frac{dR}{dE_R} = \frac{\rho}{m_\chi m_N} \int d^3\mathbf{v} F(\mathbf{v}) |\mathbf{v}| \frac{d\sigma}{dE_R}(\mathbf{v}),$$

with ρ the local dark matter mass density, m_χ the dark matter mass, $F(\mathbf{v})$ the dark matter velocity distribution in the detector's frame, and $d\sigma/dE_R(\mathbf{v})$ the differential DM-nucleus cross section per unit nuclear recoil energy E_R for a given \mathbf{v} . Not every \mathbf{v} will be kinematically able to yield a recoil energy E_R . This is sometimes indicated by integration bounds, but in the discussion below we will set $d\sigma/dE_R = 0$ when this is the case. Moreover, we will restrict to the case of a non-directional detector such as XENON1T, for which only the observed DM speed distribution $f(v)$ is relevant:

$$\frac{dR}{dE_R} = \frac{\rho}{m_\chi m_N} \int dv f(v) v \frac{d\sigma}{dE_R}(v). \quad (2.2)$$

To use 2.2, we must express it in terms of the fundamental parameters of the interaction, usually the dark matter mass m_χ and the DM-nucleon cross section σ_n , as well as astrophysical and detector parameters. This is detailed in the next subsections. While some of these are standard results, we will need them in chapters 5 and 7.

2.1.1 Velocity distribution

The velocity distribution of dark matter particles kinematically determines the energy they can deposit in elastic recoils with ordinary matter. The *standard halo model* assumes the distribution $F_g(\mathbf{v}_g)$ of dark matter particle velocities \mathbf{v}_g in the galactic rest frame is a Maxwell-Boltzmann distribution with a cutoff at an escape velocity v_{esc} . That is,

$$F_g(\mathbf{v}_g) = \begin{cases} k \left(\frac{1}{\sqrt{\pi} v_0} \right)^3 \exp \left(- (v_g/v_0)^2 \right) & v_g < v_{\text{esc}} \\ 0 & v_g \geq v_{\text{esc}} \end{cases}, \quad (2.3)$$

where

$$v_0 = \sqrt{2k_B T/m} \quad (2.4)$$

is the most probable (modal) speed of the dark matter particles, and the normalization constant k (1 if $v_{\text{esc}} = \infty$) is

$$k = \text{erf}(v_{\text{esc}}/v_0) - \frac{2}{\sqrt{\pi}} \frac{v_{\text{esc}}}{v_0} \exp\left(-\left(v_{\text{esc}}/v_0\right)^2\right), \quad (2.5)$$

as derived in appendix A of [72]. The distribution is isotropic: integrating over the unit sphere produces the Maxwell-Boltzmann speed distribution if $v_{\text{esc}} = \infty$. Note this distribution with v_{esc} finite is *not* a thermal equilibrium solution. If dark matter had a significant self-interaction, the halo would dilute and cool due to continuous escape of high-velocity particles.

The Earth moves with respect to the galactic rest frame with a velocity \mathbf{v}_E , so the velocity distribution of dark matter particles detected in an earth-based detector is not isotropic. However, for non-directional detectors such as XENON1T, we only need to compute the observed speed distribution $f(v)$ for equation 2.2. In principle, this has an annual modulation due to the Earth's 30 km/s orbital velocity around the sun; we will ignore this modulation here, as it small compared to the average $v_E \approx 232$ km/sec.

Taking the polar angle origin in the direction of earth's motion, the ϕ coordinates of the local and galactic frames are identical, and the observed speed v relates to the speed in the galactic frame v_g as:

$$v_g^2 = v^2 + v_E^2 - 2v_E v_g \cos(\pi - \theta). \quad (2.6)$$

where θ is the angle between \mathbf{v} and \mathbf{v}_E . The speed distribution is therefore

$$f(v) = \int_0^{2\pi} d\phi \int_0^\pi d\theta v^2 \sin(\theta) F_g(v_g).$$

The ϕ integral yields 2π . The θ integral is simplified by changing variables to $x = \cos \theta$. For $v \geq v_{\text{esc}} + v_E$, the integrand is zero at all angles. For lower velocities, eq. 2.6 implies $v_g < v_{\text{esc}}$ as long as

$$x < x_{\text{max}}(v) = \min\left(1, \frac{v_{\text{esc}}^2 - v_E^2 - v^2}{2v_E v}\right). \quad (2.7)$$

Thus, for $v < v_{\text{esc}} + v_E$:

$$\begin{aligned} f(v) &= 2\pi v^2 k \left(\frac{1}{\sqrt{\pi} v_0}\right)^3 \int_{-1}^{x_{\text{max}}(v)} dx \exp(v_g(v, x)/v_0)^2 \\ &= \frac{kv}{\sqrt{\pi} v_0 v_E} \left(e^{-(v-v_E)^2/v_0^2} - e^{-(v^2+v_E^2+2vv_E x_{\text{max}}(v))/v_0^2} \right), \end{aligned} \quad (2.8)$$

which reduces to eq. 8.15 in [7] for $v_{\text{esc}} = \infty$. Figure 2.2 compares the speed distribution in the galactic and local frame.

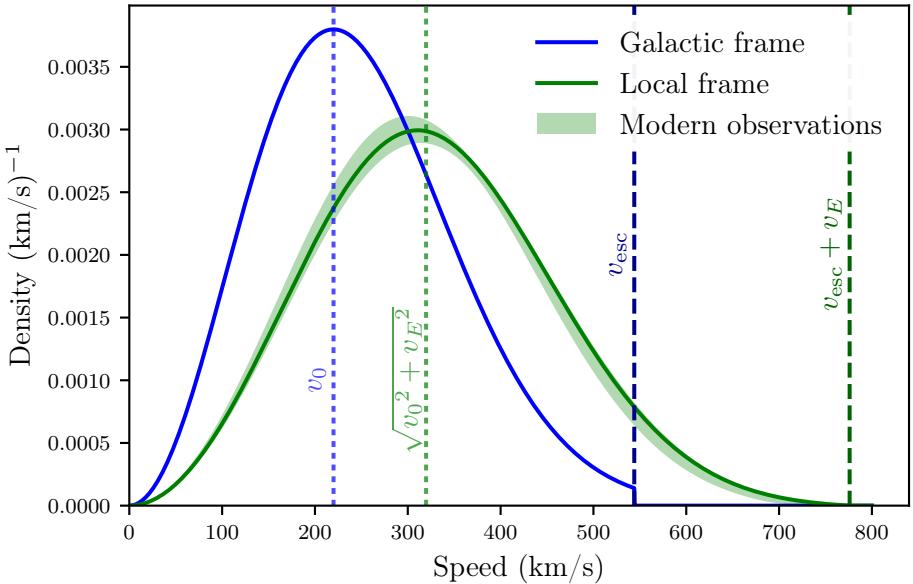


Figure 2.2: Dark matter speed distribution according to the standard halo model in the galactic rest frame (blue) and the local frame (green). Also shown are the galactic escape velocity v_{esc} , the maximum observed dark matter speed $v_E + v_{\text{esc}}$, the modal dark matter velocity v_0 in the galactic frame and the approximately modal velocity $\sqrt{v_0^2 + v_E^2}$ in the local frame. The light green band shows the local-frame velocity distribution (with $\pm 1\sigma$ uncertainty) under the standard halo model updated with modern measurements (see table 2.1 and further description in the text).

| Parameter | SHM | Modern | Units |
|------------------|-----|-------------------|--|
| ρ | 0.3 | 0.22-0.33 | $\text{GeV}/(\text{c}^2 \text{ cm}^3)$ |
| v_0 | 220 | 218 ± 6 | km/s |
| v_E | 232 | * | km/s |
| v_{esc} | 544 | 533^{+54}_{-41} | km/s |

Table 2.1: Parameters of the standard halo model (SHM) compared to those favoured by modern observations summarized in [73]. In modern literature, v_E is not reported separately, but calculated as the norm of the vector sum of v_0 and three independently measured components of the solar peculiar velocity (and the earth's orbital velocity, which we neglected here); its mean value is 230.6 km/s.

DM-scattering experiments almost universally assume the standard halo model with parameters $\rho = 0.3 \text{ GeV}/(\text{c}^2 \text{ cm}^3)$, $v_0 = 220 \text{ km/s}$, $v_E = 232 \text{ km/s}$ and $v_{\text{esc}} = 544 \text{ km/s}$, to ensure their results can be fairly compared with each other. Modern observations of galactic motions support a slightly different set of parameters, summarized in table 2.1. Moreover, the standard halo model is nowadays seen as a “convenient, but not realistic, benchmark using a distribution that dates to before simulations were capable of creating realistic galaxies” [73].

The impact of astrophysical uncertainties is not quantified by experimental collaborations themselves, and would take us too far afield here as well. Nonetheless, it can be significant for light WIMPs (see e.g. [74]). These only produce detectable signals if their speed is at the far right tail of the distribution, so small changes in the speed distribution greatly affect DM-scattering rates. For the same reason, annual modulation effects (and anisotropy) are greatest for light WIMPs, but are often ignored in DM-scattering results (as they were for XENON1T). Fortunately, for many light WIMP models, collider and annihilation experiments give stronger constraints than scattering experiments (see e.g. figure 1.4).

2.1.2 Scattering kinematics

Since WIMP velocities are nonrelativistic, the kinematics of the scattering is well-described by Newtonian mechanics. Conservation of energy and momentum can be used to show the nuclear recoil energy E_R obeys [72]:

$$E_R(v) = \frac{m_\chi v^2}{2} \frac{1 - \cos \theta}{2} \frac{4m_\chi m_N}{(m_\chi + m_N)^2} \quad (2.9)$$

$$= \frac{\mu_N^2 v^2}{m_N} (1 - \cos \theta), \quad (2.10)$$

where θ is the scattering angle in the center-of-mass frame and μ_N the DM-nucleus reduced mass. Note the maximum recoil energy is achieved for a head-on collision ($\theta = \pi$):

$$E_{\text{max}}(v) = \frac{2\mu_N^2 v^2}{m_N}. \quad (2.11)$$

Figure 2.3 shows E_{max} as a function of m_χ . When $m_\chi \gg m_N$, E_R becomes independent of m_χ ; in head-on collisions with very massive dark matter particles, the nucleus obtains four times the velocity of the dark matter particle; $E_{\text{max}} \approx 1.64 \text{ MeV}$ for xenon.

For isotropic hard-sphere scattering, $\cos \theta$ is uniformly distributed between -1 and 1 , so the recoil energy would likewise be homogeneously distributed between 0 and E_{max} . WIMP-nucleus scattering, however, is more complicated due to quantum-mechanical effects (in particular, interference caused by the non-pointlike nature of the nucleus).

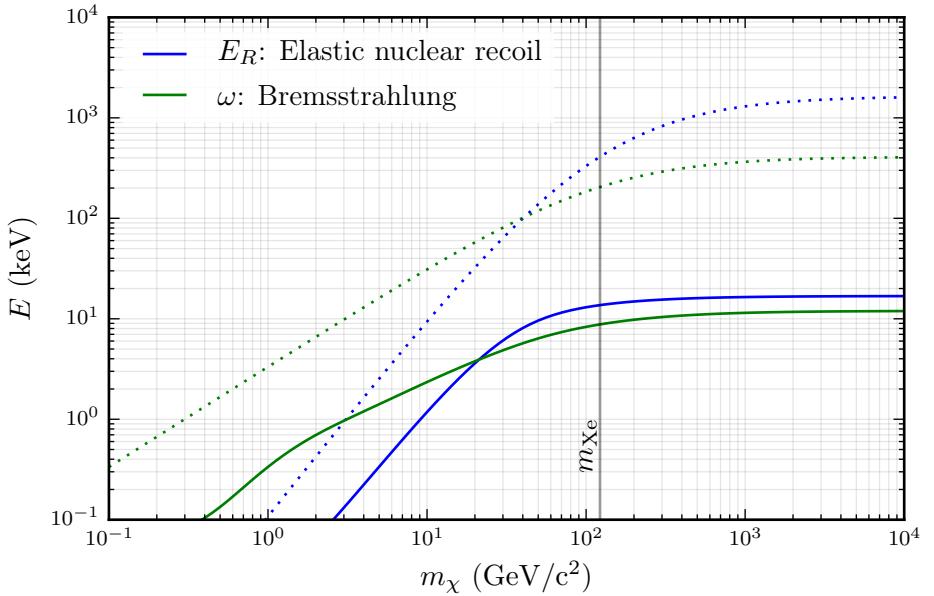


Figure 2.3: Mean (solid) and maximum (dotted) energy deposited in xenon by a dark matter particle scatter, as a function of the dark matter particle mass m_χ . Blue lines show the energy E_R of elastic nuclear recoils, and green the photon energy ω of produced Bremsstrahlung photons. The maximum energy deposition is for a head-on collision with the highest possible local-frame dark matter particle speed $v_{\text{esc}} + v_E$ according to the standard halo model (see section 2.1.1). Clearly the bremsstrahlung channel allows xenon-based detectors (with $\mathcal{O}(\text{keV})$ detection thresholds) to probe much lower WIMP masses.

2.1.3 Form factor

For elastic scattering, the *form factor* is a function $F(E_R)$ with $F(0) = 1$ defined so the differential cross section is:

$$\frac{d\sigma}{dE_R}(v) = \begin{cases} \sigma_{N,0} F(E_R)^2 / E_{\max}(v) & E_R \leq E_{\max} \\ 0 & \text{otherwise} \end{cases}, \quad (2.12)$$

where $\sigma_{N,0}$ is not dependent on E_R . The form factor is sometimes defined more specifically as accounting only for the extended nature of the nucleus, with other energy-dependent terms absorbed by other functions; we use it in its generic sense of including all energy-dependent terms here.

The form of F depends on the WIMP-nucleus interaction; interactions where F depends directly on v (aside from its dependence on E_R) can also be constructed. For $F(E_R) = 1$ we recover the hard-sphere result of a homogeneously distributed energy spectrum, with total cross-section $\sigma_{N,0}$.

The constant $\sigma_{N,0}$ is sometimes called the *point-like* WIMP-nucleus cross-section; the subscript 0 is a reminder that it is *not* the total cross-section (unless $F = 1$). No physical DM-nucleus interaction has $F = 1$, though this is often a reasonable approximation for light WIMPs (which cause only low-energy recoils).

WIMP scattering rates thus depend on both $\sigma_{N,0}$ (describing the strength of the interaction) and F (its energy dependence). Fortunately, dark matter models to which scattering experiments are sensitive generally have one of two form factors (or a mixture thereof), called *spin-independent* (SI) and *spin-dependent* (SD) scattering:

$$\frac{d\sigma}{dE_R}(v) = \frac{1}{E_{\max}} (\sigma_0^{\text{SI}} F_{\text{SI}}^2 + \sigma_0^{\text{SD}} F_{\text{SD}}^2), \quad (2.13)$$

where as in 2.12, the σ_0 terms are energy-independent, while the form factors are 1 at $E_R = 0$. (for $E_R \leq E_{\max}$). Only nuclei with non-zero nuclear spin can participate in spin-dependent scattering.

The restriction to these two scattering mechanisms follows from nonrelativistic effective field theory, which is applicable if the momentum transfer q in the interaction (times c) is much smaller than the mass of the particle(s) mediating the scattering [75, 76]. For xenon recoils,

$$q = \sqrt{2E_R m_N} \approx 15.6 \sqrt{\frac{E}{\text{keV}}} \text{ MeV/c} \leq 633 \text{ MeV/c}. \quad (2.14)$$

In the last step we used the maximum value of E_R from section 2.1.2; as we will see below, typical recoil energies are at least 100x lower (depending on the dark matter mass), so the typical q is below 64 MeV/c. Light mediators can give rise to long-range interactions and significant self-interactions for dark matter [77, 78], so are generally disfavoured. Moreover, a mediator lighter than the dark matter opens up a channel through which the dark matter could decay (so the ‘mediator’ would actually be the dark matter candidate that survives today).

Low-energy / heavy-mediator interactions can be parametrized by about ten effective field theory operators – the exact number depends on the dark matter particle type [79]. Most of these are suppressed by the momentum transfer q , or the dark matter velocity $v \approx 10^{-3}c$, both very small compared to the rest energies of the particles involved. The remaining operators give rise to the aforementioned spin-dependent and spin-independent interactions.

With some exceptions [80], DM-scattering experiments present their results as constraints on one or both of these scattering models. Spin-independent constraints are easiest to compute, and strongest, since heavy nuclei such as xenon have a large number of nucleons, but not especially high nuclear spin. For spin-independent elastic nuclear scattering, the form factor is well-

approximated by the Helm function [72]:

$$\begin{aligned}
 F_{\text{SI}} &= \frac{3j_1(\hat{q}r)}{\hat{q}r} e^{-(\hat{q}s)^2/2}, \\
 \hat{q} &= q/(\hbar c), \\
 r^2 &= (1.23A^{1/3} - 0.6)^2 \text{fm}^2 + \frac{7}{3}\pi^2(0.52 \text{ fm})^2 - 5s^2, \\
 s &= 0.9 \text{ fm}.
 \end{aligned} \tag{2.15}$$

Here A is the number of nucleons in the nucleus, and $q = \sqrt{2E_R m_N}$ the momentum transfer. For xenon, this suppresses the cross-section by 1.74×10^{-8} at the maximum possible recoil energy; for a typical recoil energy of 17 keV (see section 2.1.4), the suppression is about 0.5.

For spin-dependent elastic nuclear scattering, the form factor is usually expressed through a *structure function* S_A :

$$F_{\text{SD}}^2(E_R) = S_A(E_R)/S_A(0). \tag{2.16}$$

This structure function depends on the detailed interaction of the dark matter with the different nucleons. Computing S_A requires advanced nuclear physics calculations, and unlike for the spin-independent case, a relevant theoretical uncertainty remains [81]. Recent spin-dependent constraints from xenon-based detectors used the structure functions from [82] (used by XENON100 [83]) or its updated result [84] (used by LUX [85, 86]). We will discuss spin-dependent scattering further in section 7.2

2.1.4 Nuclear cross-section

We can now compute WIMP-nucleus scattering rates as a function of a single parameter, $\sigma_{0,N}$. While it is possible to present experimental results as constraints on $\sigma_{0,N}$, this $\sigma_{0,N}$ still depends on the specific target nucleus used. To facilitate easy comparison between the results of experiments with different target media, this dependence can (mostly) be factored out. The factorization is different for spin-independent and spin-dependent scattering.

For a scalar WIMP that couples via a scalar mediator to quarks, it can be shown the interaction is completely spin-independent, and [8, p.116]:

$$\sigma_{N,0,\text{SI}} = \frac{4\hbar^2 c^2}{\pi} \mu_N^2 (Z f_p + (A - Z) f_n)^2, \tag{2.17}$$

with Z and A the proton and nucleon count and f_p and f_n the dark matter coupling to protons and neutrons, respectively. These couplings can be derived from couplings to quarks if those are supplied by an underlying theory. If $f_p \approx f_n$, this simplifies to $(4\hbar^2 c^2/\pi) \mu_N^2 (A f_n)^2$: each of the A nucleons contributes a scattering amplitude proportional to f_n , which all add coherently.

DM-scattering experiments quote their results in terms of the parameter σ_n , defined through:

$$\sigma_{N,0,SI} = \sigma_{n,SI} \frac{\mu_N^2}{\mu_n^2} A^2, \quad (2.18)$$

with μ_n the DM-nucleon reduced mass. In the scalar mediator model discussed above, $\sigma_n = (4\hbar^2 c^2/\pi) \mu_n^2 f_n^2$, i.e. the pointlike cross-section for a nucleus with $A = 1$. This is manifestly independent of the nucleus used in the experiment. For general models, this independence may not hold, e.g. because $f_p \neq f_n$. In such cases, a correction is needed to compare results from experiments with different target media.

For spin-dependent scattering of WIMPs on the protons in the nucleus, the equivalent convention is to use:

$$\sigma_{N,0,SD} = \sigma_{p,SD} \frac{4\pi \mu_N^2}{3\mu_p^2 (2J+1)} S_{A,p}(E_R = 0), \quad (2.19)$$

with J the total nuclear spin and μ_p the DM-proton reduced mass [83]. For neutrons, the same formula with $p \rightarrow n$ applies.

Putting the results above (eqs. 2.8, 2.12, 2.15 and 2.18) into the master rate equation 2.2, we find the rate and spectrum of elastic nuclear recoils for a spin-independent WIMP interaction. Rates for different energy thresholds are shown in figure 2.4; an example spectrum for a 50 GeV/c² WIMP is shown in figure 5.1. The average nuclear recoil energy for different WIMP masses is shown in solid blue in figure 2.3; for high-mass WIMPs, this is about 16.8 keV. The large discrepancy with the kinematic maximum (dashed blue in the same figure) for high WIMP masses is caused by the nuclear form factor.

2.2 Alternative dark matter signatures

2.2.1 Electronic recoil

Dark matter could scatter off atomic electrons rather than the nuclei. Since electrons are much lighter than nuclei, and cross-sections are proportional to the particle masses (determining the phase-space density), the DM-electron cross-section is suppressed by $\mathcal{O}(10^{-5})$ compared to the DM-nucleus cross-section. Moreover, replacing $m_N \rightarrow m_e$ in equation 2.11, we see dark matter that recoils off electrons at rest would deposit at most 6.8 eV, and generally much less – well below the ionization threshold of xenon (see section 2.3.1), let alone the detection threshold of contemporary detectors. To increase the energy transfer to $\mathcal{O}(\text{keV})$, the electron must have a momentum much higher than typical of bound atomic electrons; this adds another similarly-sized suppression factor [87].

Nonetheless, direct detection experiments can set limits on the DM-electron cross section σ_e through searches for electronic recoils [88, 89]. Even for most

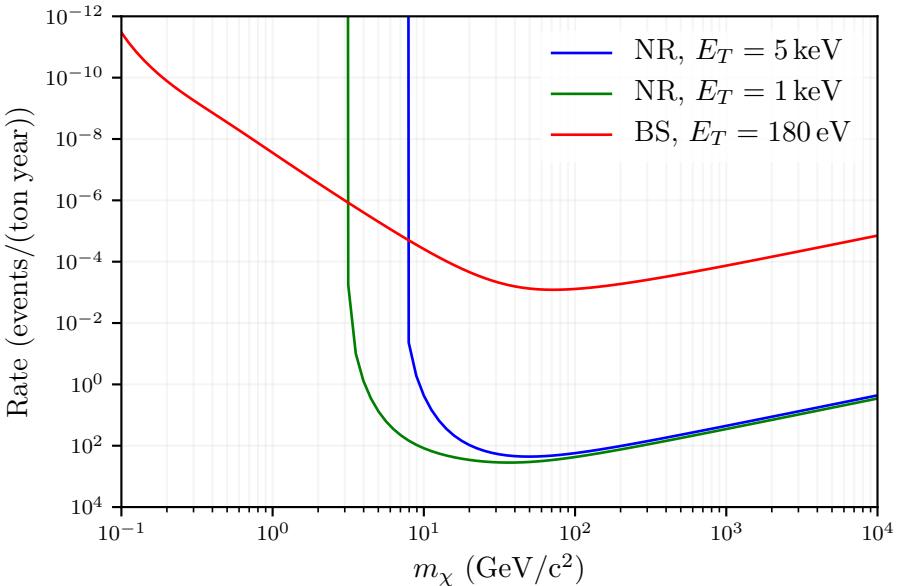


Figure 2.4: Expected event rates for spin-independent dark-matter nucleus scattering for $\sigma_n = 10^{-45}\text{cm}^2$. Vertical scale is inverted to show similarity to exclusion limit plots (e.g. figure 5.4). Blue: nuclear recoil search with an energy threshold of 5 keV, corresponding roughly to XENON1T’s S1-S2 search (see chapter 5). Green: same, but with a 1 keV threshold, corresponding to the S2-only search of chapter 7. Red: bremsstrahlung photon search with a threshold of 180 eV, corresponding to the S2-only bremsstrahlung search of chapter 7.

so-called *leptophilic* dark matter models, where dark matter couples only to electrons (or additionally to the other leptons), constraints on σ_e from ordinary nuclear recoil searches derived through loop-induced DM-nucleon scattering are orders of magnitude stronger [90]. The exceptions are WIMPs much lighter than nuclei, for which nuclear recoil limits are weak due to increasingly inefficient energy transfer (see figure 2.3) and detector thresholds; and leptophilic models with axial-vector couplings, in which the virtual lepton loop contributions cancel [87]. The latter model had been proposed as a way of reconciling the DAMA/LIBRA signal claims (see section 1.2.4) with null results from other experiments, but XENON100 has excluded this interpretation at 4.4σ [46].

2.2.2 Inelastic scattering

Above, we considered only elastic WIMP-nucleus scattering, since the typical energies deposited in dark matter scattering are insufficient to excite $\mathcal{O}(\text{MeV})$

nuclear energy levels. However, the odd-mass isotopes ^{129}Xe and ^{131}Xe have low-lying excited states at 39.6 keV and 80.2 keV respectively, which could practically be excited by WIMP recoils [91, 92]. This yields a nuclear recoil followed (after an $\mathcal{O}(\text{ns})$ decay time) by a de-excitation photon. The observed signal would predominantly consist of quanta created due to photoelectric absorption of the photon, pushing it into the background-rich electronic recoil band (see section 2.3). Sensitivity calculations show searches for this inelastic scattering do not extend the discovery reach of modern detectors – though they would be valuable in constraining the nature of a dark matter signal once observed [93].

In certain dark matter models, the dark matter particle can itself be excited by the scatter, while elastic scattering is highly suppressed, changing the energy spectrum of the expected signals. This has been suggested as another way to reconcile the DAMA/LIBRA signals with null results from other experiments [94], but again, data from XENON100 excludes the models that could support such an interpretation [95].

2.2.3 Bremsstrahlung

Bremsstrahlung photons are occasionally (but rarely) emitted in ordinary nuclear recoil scattering. For low-energy scatters, the photon acquires a negligible amount of momentum, but its energy ω can carry away up to all of the energy available in the center of mass frame:

$$\omega_{\max} = \frac{1}{2} \mu_N v^2 = E_{\max, \text{elastic}} \frac{m_N}{4\mu_N}, \quad (2.20)$$

with $E_{\max, \text{elastic}}$ as given by 2.11. This is shown in green dashed in figure 2.3. The nuclear recoil energy for head-on elastic recoils with high-mass WIMPs (observed in the frame in which the nucleus is initially at rest) can be larger; however, for light WIMPs ($m_\chi < m_N/3$), the maximum energy of Bremsstrahlung photons exceeds that of elastic recoils. Moreover, as discussed in 2.3, the resulting X-ray photon will deposit almost all of its energy in visible channels (mostly ionization), while an elastic nuclear recoil loses $\sim 90\%$ of its energy as heat. As such, bremsstrahlung has been suggested as a viable detection channel for low-mass dark matter [96, 97].

The minimum and maximum E_R that can produce a photon with energy ω for a given v are kinematically determined [96]:

$$E_{\max/\min} = \frac{\mu_N v^2}{m_N} \left(\left(1 - \frac{\omega}{\mu_N v^2} \right) \pm \sqrt{1 - \frac{2\omega}{\mu_N v^2}} \right). \quad (2.21)$$

Assuming we only observe the photon, we are interested in the differential rate with respect to ω (a slight modification of equation 2.2):

$$\frac{dR}{d\omega} = \frac{\rho}{m_\chi m_N} \frac{1}{\int dv f(v)} \int dv f(v) v \int dE_R \frac{d\sigma}{d\omega dE_R}(v). \quad (2.22)$$

The differential cross-section for bremsstrahlung is derived in [96]:

$$\frac{d\sigma}{d\omega dE_R}(v) = \begin{cases} \frac{4\alpha}{3\pi} \frac{1}{\omega} \frac{E_R}{m_N c^2} |f_X(\omega)|^2 \frac{d\sigma}{dE_R} & E_{\min} \leq E_R \leq E_{\max} \\ 0 & \text{otherwise} \end{cases}. \quad (2.23)$$

Here α is the fine-structure constant, f_X the atomic X-ray form factor, and $d\sigma/dE_R$ the differential elastic DM-nucleus cross-section, equation 2.12.

The mean energy of bremsstrahlung photons as a function of WIMP mass is shown in figure 2.3, and is indeed substantially higher than mean elastic nuclear recoil energies for sufficiently low WIMP masses. Thus, bremsstrahlung searches can expect substantial event rates for low-mass WIMP models inaccessible by traditional searches, as shown in figure 2.4.

2.3 Particle detection with liquid xenon

The detection medium of a dark matter detector converts energy depositions of DM scatters, usually in the form of nuclear recoils (NRs), into signals that can be registered by conventional sensors – in the case of XENON1T, ultimately photons detected by photomultipliers. This section discusses the response of liquid xenon (LXe), XENON1T’s detection medium. We also discuss signals caused by recoils of electrons (electronic recoils or ERs); although these are unlikely to be produced by WIMPs (see section 2.2.1), they are the dominant signal caused by radioactive background.

We describe the physical processes underlying signal formation in liquid xenon (section 2.3.1), the quantitative charge and light yields (section 2.3.2) and finally the possible sources of ER (section 2.3.3) and NRs (section 2.3.4) signals. While we occasionally mention the relevance of a background source for XENON1T, the aim is to provide context for the description of XENON1T and its calibration that follows in section 2.4, not to discuss the full XENON1T background model (see sections 5.3 and 6.3.1 for that).

2.3.1 Interaction microphysics

The name ‘nuclear recoil’ could be seen as misleading. While the first ionization energy of Xe is only 12.13 eV, it takes 42 keV to fully ionize a xenon atom [98]. Thus, nuclei set in motion by ~ 17 keV WIMP recoils will retain at least some electrons. These initial xenon ions gradually lose their energy through the three mechanisms mentioned in section 1.2.4: heat (elastic scattering), charge (ionization) and light (scintillation); the latter two yield detectable signals in XENON1T.

Noble gases such as xenon are excellent scintillators. When ionized or excited, noble gas atoms produce light to which they are themselves transparent, unlike light from ordinary atomic energy level transitions. An excited xenon atom (exciton) rapidly combines with a ground-state atom to form an excited

Xe_2^* dimer (excimer), which relaxes to one of two electronic ground states (mostly vibrationally, with occasional infrared emission) with an essentially equal energy level: a spin singlet and spin triplet [99]. From either of these states, the dimer decays back to separate Xe -atoms and a characteristic 7 eV ultraviolet photon, which passes freely through the xenon. The infrared photon yield is estimated to be less than a percent of the UV-photon light yield [99].

The singlet state has a shorter lifetime (~ 3 ns) than the triplet (~ 24 ns), and singlet/triplet ratios are found to depend on the type (ER vs NR) and energy of the interaction [100, 101]. However, unlike in liquid argon, this difference in lifetimes is too small to be of much use in dark matter detectors, considering the timing resolution of large photomultipliers and especially the small S1 signals produced by WIMPs.

Noble elements are also excellent media for charge detection due to their ultralow electron affinity. Free electrons can be drifted through liquid xenon at speeds of $\mathcal{O}(\text{km/s})$ (see section 3.5.1) by application of electric fields of $\mathcal{O}(100 \text{ V/cm})$. The Xe^+ ions created during ionization rapidly combine with neutral Xe to make Xe_2^+ ions. Like electrons, their drift can be guided by electric fields, but in liquid xenon, ion drift is so slow ($\mathcal{O}(\text{mm/s})$ in XENON1T [102]) that the positive charges behave as ‘holes’, i.e. they mainly drift along the field by means of electron exchange. In this process, Xe_2^+ steals an electron from a neutral Xe atom, which then takes the place of one of the Xe atoms in the excimer [103].

The spatial extent of nuclear recoil tracks in LXe can be calculated by numerical codes such as SRIM/TRIM [104]. For the energies of WIMP recoils, the track length is $\ll 100 \text{ nm}$, as shown in figure 2.5. This is much smaller than the $\sim 4.6 \mu\text{m}$ mean thermalization length in LXe for low-energy ($\sim 10 \text{ eV}$) electrons created by ionizations along the track [105], indicated by the purple dashed line in figure 2.5. Thus, the charge signal will not contain information about the track’s structure, and LXe detectors cannot be used as directional dark matter detectors.

The approximate extent of electronic recoil tracks can be computed from electronic stopping power tables [106], see figure 2.5. Electronic recoil tracks are substantially longer and have a lower energy density than nuclear recoil tracks: since electrons have 4.2×10^{-6} the mass of a xenon atom, elastic energy losses in electronic recoils are insignificant.

2.3.2 Charge and light yields

Describing the energy distribution over heat, charge, and light is a difficult task, due to the variety of processes involved:

- *Elastic energy loss.* Nuclear recoils lose a substantial part of their energy in elastic nuclear recoils. Since xenon atoms have the same mass, a single collision could transfer all kinetic energy. While such secondary recoils can create further signals (if they are energetic enough, and especially

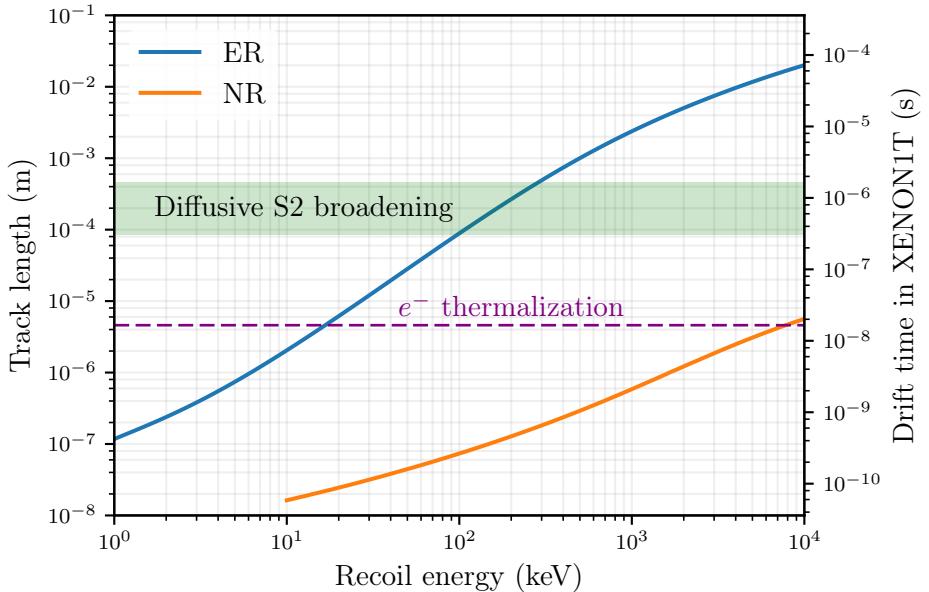


Figure 2.5: Track length of electronic (blue) and nuclear (orange) recoils in LXe of 3 g/cm^3 . The right y-axis indicates the time it takes in XENON1T to drift an electron the same distance. The purple line indicates the mean thermalization distance of ionization electrons, and the green band indicates electron cloud smearing from electron diffusion during drift (discussed in section 3.5).

if they are ionized), much energy is lost to heat, which can only be measured by specialized detectors (e.g. ultralow temperature crystals or bubble chambers). As mentioned above, the elastic energy loss for electrons is insignificant.

- *Exciton/ion balance.* The recoiling particle loses energy through some combination of direct excitation and ionization (and, for nuclear recoils, elastic collisions), depending on the ratio between the (generally poorly known) cross sections of these processes.
- *Recombination* of electron/ion pairs diminishes the charge signal in favour of light emission (since the recombined atom will be in an excited state). For tracks with high ionization density (such as high-energy NRs and as), recombination can happen very rapidly; for tracks with ionization density (such as high-energy ERs), the ejected electron has enough time to thermalize. Afterwards, when an electron approaches an ion within the Onsager radius ($\mathcal{O}(50 \text{ nm})$), the electrostatic potential of the electron/ion pair exceeds the thermal energy scale $k_B T$, and recombination is likely. Without a drift field, many recombination events

would eventually occur as electron/ion pairs accidentally meet during their random motion; with a field, the electron cloud is pulled away from the ion cloud, and recombination quickly ceases.

- *Biexcitonic quenching* [107], sometimes called Penning ionization, results when two Xe^* excitons collide to form an electron/ion pair and one unexcited Xe atom, rather than combine with two other Xe atoms to produce two Xe_2^* excimers. This reduces the scintillation yield; even if the electron/ion pair recombines, this causes the emission of one photon instead of two. Such a meeting of two Xe^* excitons before excimer formation is likely only if the track is very dense.

The effects of these processes are reflected in the observed light and charge yields, i.e. the number of photons and electrons, respectively, emitted per keV deposited energy. Unfortunately, the effects are partially degenerate, and no currently available model is both simple, physically motivated, and has good agreement with most observations. The models that do exist (see e.g. [99] for a review) have low epistemic status; experiments usually measure LXe yields in-situ with calibration sources, or rely (partially) on similar measurements from other detectors. Attempts to use the fledgling LXe yield models beyond the ranges in which they have been measured are met with skepticism [108]. In WIMP searches, it is conventional to assume no light and charge yield for nuclear recoils below the lowest calibration energy achieved (currently 0.7 keV [109]).

A global fit of experimental data for nuclear recoils has been published in [110]. The predictions of this semi-empirical model are shown in figure 2.6 (for xenon subjected to an electric field of 120 V/cm). Under this model, nuclear recoils primarily produce ions at low energies ($\sim 80\%$ at 1 keV), but an equal amount of excitons and ions beyond ~ 10 keV. Recombination of ion/electron pairs is essentially absent at low energies (0.04% of ions recombine at 1 keV) but nearly complete at high energies (99.7% at 1 MeV). Biexcitonic quenching is absent at low energy and a few-percent effect at dark matter recoil energies, but increases with energy until 75% of excimers undergo this quenching at 1 MeV. The elastic energy loss decreases from 87% at 1 keV to 60% at 1 MeV.

For the distribution of ER energy over charge and light, no explicit analytical model or a global fit to experimental data has been published. While the simulation package NEST [114] includes predictions for ER yields, none of the NEST papers published as of this writing compare its predictions to measurements of light and charge yield data for ER energies relevant to dark-matter searches (< 10 keV). In fact, recent measurements from LUX and neriX show substantial tension with NEST's prediction, especially around 20 keV, as shown in figure 2.7 [111, 113].

However, many observations have found that for electronic recoils, charge yield and light yield sum to a constant of 73 eV^{-1} quanta (electrons or photons) over the entire keV - MeV energy range where LXe detectors have been

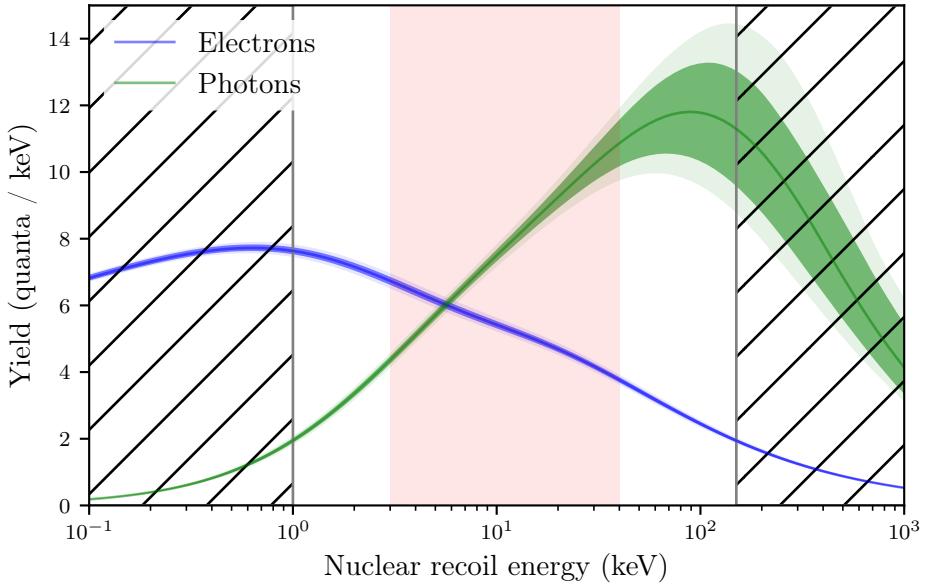


Figure 2.6: Charge and light yields for nuclear recoils in liquid xenon at 120 V/cm according to the NEST global fit [110]. Hatched ranges lie beyond the last datapoints included in the fit. The red region indicates the (5-40) keV range of most interest to conventional dark matter searches with LXe TPCs (see chapter 5). Colored bands represent estimated 1 and 2 σ confidence intervals from varying the global fit’s parameters based on the intervals in table I of [110].

calibrated, and at any drift field at which measurements have been done. Recently, neriX found some deviations from this picture below 3 keV, but it is generally still assumed to be approximately accurate [113]. Since for ERs we can neglect elastic energy loss and (due to the low density of the track) biexcitonic quenching, this suggests that creating a single detectable quantum in LXe takes a constant energy $W = 13.7$ eV. For ERs, the original energy E deposited in the interaction can then be reconstructed from estimates of the produced number of photons n_{ph} and electrons n_e :

$$E = W(n_{\text{ph}} + n_e). \quad (2.24)$$

Energy estimates of features in data using equation 2.24 are sometimes called ER-equivalent energies, indicated by an “ee” subscript on the energy units, e.g. keV_{ee}.

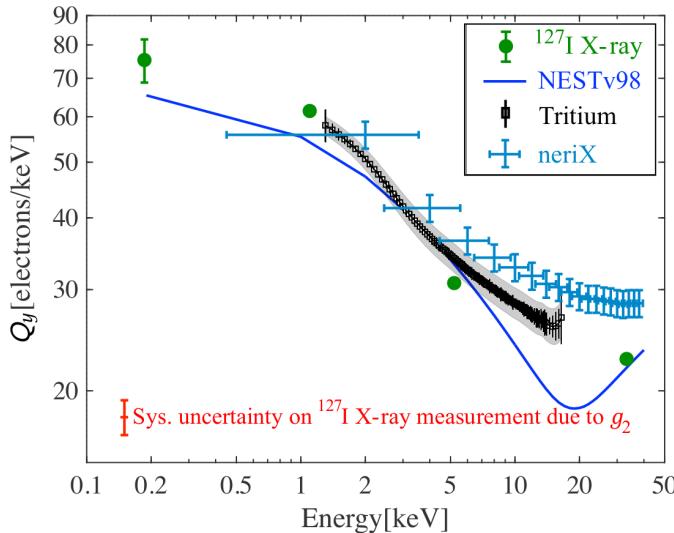


Figure 2.7: From [111]: ER charge yields according to NEST v0.98, LUX calibrations with tritium [112] and ^{127}I [111], all at 180 V/cm. The light blue points show neriX’s Compton coincidence measurements at 170 V/cm [113].

2.3.3 Sources of electronic recoils

Electronic recoils can be distinguished from nuclear recoils due to their different charge and light yields. However, statistical fluctuations of the amount of detected charge and light quanta (and to a lesser extent, fluctuations in the amount of recombination in the electron/ion cloud) make the distinction imperfect, so electronic recoils ≤ 10 keV (see figure 5.2) are a relevant (indeed, dominant) background for nuclear recoil WIMP searches.

2.3.3.1 X-rays

Electron recoils < 10 keV can be produced photoelectrically by X-ray photons.¹ These are not a measurable external background source, as they penetrate only $\mathcal{O}(10 \mu\text{m})$ in xenon, as shown in figure 2.8. However, such X-rays can also be produced by radioactive impurities in the LXe that decay via electron capture, such as the cosmogenically produced ^{127}Xe with a 36-day half-life. For xenon that has been stored underground for several months or years, this background is negligible. However, LUX has used it successfully for ER calibration at 186 eV, the lowest calibration energy in LXe achieved to date [111].

¹When these are of radioactive origin, they are sometimes still called gamma rays, despite their low energy.

2.3.3.2 Compton scattering of gamma-rays

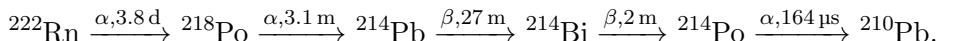
High-energy ($\mathcal{O}(\text{MeV})$) gamma rays can create low-energy electronic recoils by Compton scattering, then leave the detector (especially a small detector) without depositing the rest of their energy. As such, external gamma rays were the dominant background for earlier detectors such as XENON10 and XENON100. These detectors could use Compton scatters from external gamma-ray sources such as ^{137}Cs , ^{232}Th and ^{60}Co for calibrating their detector's electronic recoil response. Modern detectors such as XENON1T are much larger than the $< 10\text{ cm}$ penetration depths of gamma rays (see figure 2.8), so external gamma rays are not an important background in the inner detector volume. The utility of external calibration sources is likewise limited.

2.3.3.3 Beta decay and internal conversion

Beta decay produces moving electrons with arbitrarily low kinetic energy; specifically, beta decays of heavy elements all cause an effectively flat spectrum ($< 10\text{ keV}$), due to their high ($\mathcal{O}(\text{MeV})$) kinematic endpoint energies. Beta decays produce the same signal as an electronic recoil; although the electron was not set in motion by a collision, we call it an ‘electronic recoil’ signal nonetheless. Internal conversion decays and Auger electron emission are also sources of low-energy electrons in the TPC, but produce a peaked energy spectrum rather than the beta-decay continuum.

Betas from external sources are not a concern in XENON1T, since $\mathcal{O}(10\text{ keV})$ electrons are stopped in a few μm of LXe (see figure 2.5). In contrast, beta decays from internal contaminants are the dominant sources of electronic recoil background:

- The primary contribution comes from ^{214}Pb , a daughter of ^{222}Rn (part of the ^{238}U chain), which diffuses into the xenon from detector materials:



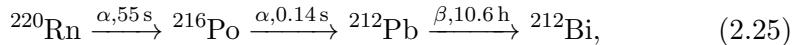
We omitted secondary branches, which occur in $\mathcal{O}(0.02\%)$ of decays. The beta from ^{214}Bi is followed so closely by the high-energy alpha of ^{214}Po that it is easy to distinguish, and the 22.3-year lifetime of ^{210}Pb is long enough that its activity is negligible compared to ^{214}Pb (for xenon that is continuously purified).

- A secondary contribution comes from contamination of the xenon with ^{85}Kr (10.76 year halflife), present in the air from which the xenon was distilled due to human nuclear activities. This can be reduced by distillation of the xenon; unless the isotope is re-introduced by an air leak, the distillation does not have to be running continuously.

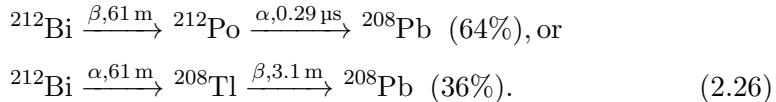
- Double-beta decay from ^{136}Xe (8.9% abundance in natural xenon) is not a significant contribution due to its very long lifetime (2.38×10^{21} y [115]).

The quantitative contribution of these backgrounds to XENON1T’s WIMP search is discussed in [68] and section 5.3.

The electronic recoil response of large LXe TPCs is usually probed by internal beta sources. Tritium (^3H) produces many useful low-energy decays due to its low endpoint energy of 18.6 keV, and has been successfully used by LUX [112]. However, its 12.3-year lifetime presents a risk for use in a functioning low-background experiment. A source of ^{222}Rn could be used, though repeated calibrations would cause accumulation of the radioactive ^{210}Pb . XENON1T instead favours ^{220}Rn (from the Thorium series):



with further decays to the stable ^{208}Pb through either:



Section 6.2 discusses ^{220}Rn calibration in more detail.

The internal conversion source $^{83\text{m}}\text{Kr}$ is also frequently used for calibration. It has a 1.8-hour decay time, and produces two photons: one of 32.2 keV followed shortly ($\tau \approx 150$ ns) by one of 9.4 keV. These energies are too high to calibrate the ER response near the WIMP region of interest. Instead, $^{83\text{m}}\text{Kr}$ is used for probing the homogeneity of the detector response, since the double photon emission provides a clear signature to distinguish $^{83\text{m}}\text{Kr}$ events from background events even near the detector’s edge.

2.3.3.4 Neutrino-electron scattering

Neutrinos can scatter off electrons in xenon atoms, yielding another source of electronic recoils. Solar neutrinos from the pp fusion process are by far the dominant contribution, though these still yield only 5% of XENON1T’s expected electronic recoil background [68]. Neutrinos are not usually considered a viable electronic recoil calibration source, since an extremely-high intensity source would be required to produce them in appreciable numbers, and easier alternatives are available.

2.3.4 Sources of nuclear recoils

Low-energy nuclear recoils in xenon can only be practically produced by neutral particles. Charged particles such as electrons, muons, and alphas can also recoil off the nucleus, but while moving through the xenon, they produce a track of ionization and excitation centers that will overwhelm the weak signal of a low-energy nuclear recoil.

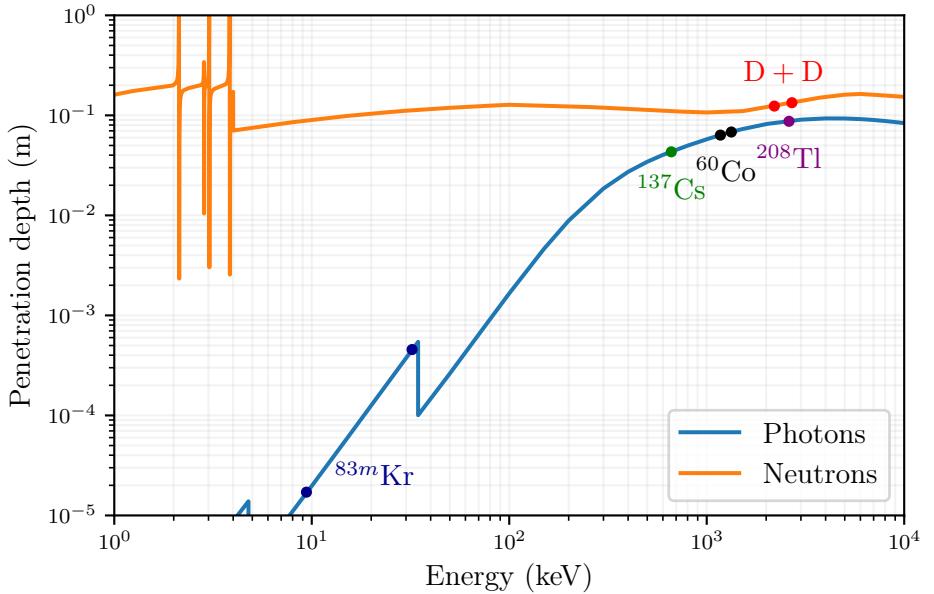
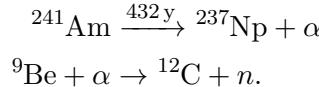


Figure 2.8: Penetration depths of photons (blue) and neutrons (orange) in LXe of 3 g/cm^3 . Peak energies of sources discussed in the text are indicated, except for $^{241}\text{AmBe}$ and ^{232}Th , which have complex multi-peaked spectra. Photon data is from [117], neutron data from [118].

2.3.4.1 Neutron scattering

Background neutrons arise from radioactivity (spontaneous fission and α -n reactions) and from spallation by cosmic muons. More importantly, however, neutrons are the only nuclear recoil calibration source used by current dark matter detectors. Calibration neutrons are usually made by an $^{241}\text{AmBe}$ source, where the Americium's alphas cause neutron emission from the Beryllium:



The neutrons are emitted with a complicated multi-peaked spectrum of energies $\lesssim 10 \text{ MeV}$. Since xenon atoms are over a hundred times heavier than neutrons, the nuclear recoil energy received in collisions is small, $\mathcal{O}(10 \text{ keV})$ for $^{241}\text{AmBe}$ (though, as mentioned, the spectrum is broad). The $^{241}\text{AmBe}$ source also produces many high-energy gamma rays, so portable deuterium-deuterium (D-D) fusion reactors are increasingly used as alternative neutron sources [109, 116].

Neutron penetration depths in LXe are $\mathcal{O}(10 \text{ cm})$ (see figure 2.8), so calibrating large detectors with neutrons is possible but challenging. Moreover, neutrons are likely to scatter multiple times in a large detector, and move too

fast ($\mathcal{O}(10 \text{ m}/\mu\text{s})$ at $\mathcal{O}(1 \text{ MeV})$) to yield easily distinguishable signals when they do, polluting the calibration sample with signals that are unlike WIMPs.

2.3.4.2 Coherent neutrino-nucleus scattering

Solar, atmospheric and diffuse supernova neutrinos can scatter coherently off the nucleus via Z-boson exchange, producing a single-scatter nuclear recoil signal. While the total cross-section for coherent neutrino-nucleus scattering (CNNS) with heavy elements is orders of magnitude larger than for neutrino-electron scattering, CNNS predominantly produces sub-keV signals that are very challenging to detect. As such, the first observation of this effect was made only recently by the COHERENT experiment, which detected neutrinos from decaying pions produced by a proton beam dump [119]. While CNNS could be used for measuring the nuclear recoil response of LXe to very low energy, calibrating large dark matter detectors with it would be challenging (even after transporting the detector to a high-power beam dump, one would face extremely high backgrounds there).

Above 0.1 keV, the dominant CNNS background comes from solar neutrinos, specifically electron neutrinos from the ${}^8\text{Be} \rightarrow {}^7\text{Be}^* + e^+ + \nu_e$ decay; at $\sim 10 \text{ keV}$, atmospheric neutrinos take over as the dominant CNNS source [4]. Since CNNS and WIMP signals are indistinguishable by a non-directional detector, this imposes a ‘neutrino floor’ on the sensitivity of dark matter scattering experiments (see e.g. figure 1.4) at cross-sections where the CNNS and WIMP scattering rates become comparable. Progress beyond this ‘floor’ will be slow and depend on accurate solar models, until directional detectors become competitive.

2.3.4.3 Photonuclear scattering

Gamma rays can produce nuclear recoils, although low-energy gamma rays are far more likely to interact with electrons than the nucleus. This can be understood in the classical Thomson limit ($E_\gamma \ll mc^2$, with E_γ the photon energy and m the mass of the scattered particle): the electron’s higher charge-to-mass ratio means it produces stronger scattered electromagnetic waves for a given incoming wave. However, when $E_\gamma \geq mc^2$, the classical approximation breaks down, and the photon-electron cross section quickly decreases with increasing photon energy. Beyond $\mathcal{O}(10 \text{ MeV})$, a photon is more likely to interact with an atomic nucleus than its electron cloud.

Still, this rarely yields observable nuclear recoils, as photon-nucleus encounters at these energies usually produce an electron-positron pair. The electron and positron produce high-energy signals in the detector which easily drown out any small signal from the nuclear recoil. Gamma rays can, however, scatter off the nucleus without producing an electron-positron pair by a variety of mechanisms such as Thomson scattering, nuclear Rayleigh scattering, and even Delbrück scattering where the photons scatter off the vacuum polarization of

the nucleus' electric field [120]. At $\mathcal{O}(10\text{ MeV})$, the dominant mechanism is nuclear resonance fluorescence [121].

Direct measurements of xenon's elastic photonuclear cross-section were not found, though methods for calculating it exist [122]. For lead, the elastic photonuclear cross section peaks at approximately 14 MeV and $\sim 0.65\text{ b}$ due to the nuclear giant dipole resonance [123]. For comparison, the total photon cross-section (including electronic and inelastic nuclear scattering modes) at 14 MeV is 19 b for lead and 9.5 b for xenon [117]. For appropriately chosen gamma energies, elastic photonuclear recoils thus make up a few percent of scattering events.

The energy deposited in an elastic photonuclear recoil is determined by kinematics, specifically Compton's formula. For $E_\gamma \ll m_N c^2$, the recoil energy is

$$E = \frac{E_\gamma^2}{m_N c^2} (1 - \cos \theta), \quad (2.27)$$

with θ the angle between the incoming and outgoing photon in the frame where the nucleus is initially at rest. The minimum E_γ to produce a photonuclear xenon recoil of energy E is therefore

$$E_\gamma \approx 7.8 \text{ MeV} \sqrt{\frac{E}{\text{keV}}}. \quad (2.28)$$

Thus, gamma rays produced by natural radioactivity ($< 3\text{ MeV}$) do not produce measurable photonuclear recoils in xenon. For example, the common 2.61 MeV ^{208}Tl gamma can produce recoils of at most 112 eV, which yields an average of ~ 0.8 electrons if the $\sim 7\text{ e/keV}$ charge yield can be extrapolated to such low energies. As such, photonuclear recoils are not a background for dark matter searches with xenon, though detectors with substantially lower energy thresholds might have to consider it.

Higher-energy gamma rays can be produced by thermal neutron activation (up to about 11 MeV, [124]) and accelerators. Indeed, photonuclear recoils from such gammas have been considered for low-energy measurement of LXe's yields [125], though the relatively small penetration depth of gammas compared to neutrons (see figure 2.8) will limit its usefulness as a calibration source for large detectors.

2.4 XENON1T and LXe TPCs

This section will focus on XENON1T, a dual-phase noble element time-projection chamber (TPC). We discuss its design (section 2.4.1), calibration (section 2.4.2) and the techniques used for discriminating dark matter signals from backgrounds (section 2.4.3). A more complete treatment of the XENON1T design, components, and performance can be found in [126], and more details on noble element TPCs in general in [99, 127]. The aim here

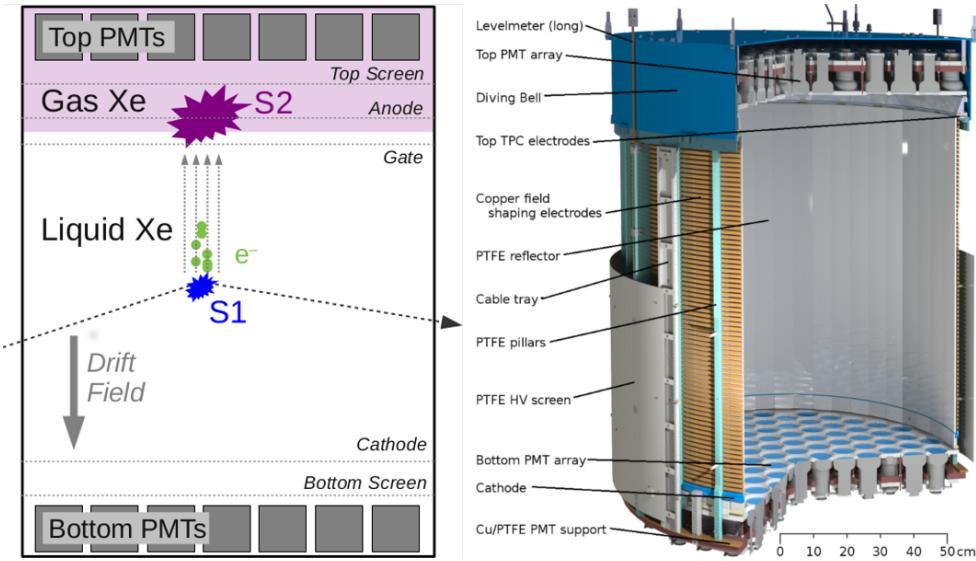


Figure 2.9: Adapted from [126]: schematic illustration of a noble liquid TPC (left) and illustration of XENON1T (right).

is only to provide a minimal context for later chapters, which will examine XENON1T’s signals (chapter 3) and analysis (chapters 5-7) in more detail.

2.4.1 Detector design

Figure 2.9 schematically shows the principle of a noble liquid TPC (left), next to a design drawing of XENON1T (right). In the left panel, a particle (dashed horizontal arrow) collides with a xenon nucleus. The resulting moving xenon ion causes excitations and ionizations. Excitation (and recombination after ionization) causes the emission of scintillation light, which is detected by the photomultipliers (PMTs) at the top and bottom of the chamber as the *S1 signal*. Ionization produces free electrons, which travel up an applied electric field to a layer of gaseous xenon (GXe). Here, they are free to be accelerated by an even stronger electric field, to energies high enough to excite xenon atoms again. This produces secondary scintillation light, which is detected by the PMTs as the *S2 signal*. The pattern of the S2 signal on the top PMTs indicates the transverse position of the interaction, while the time between S1 and S2 signals indicates the interaction depth. Chapter 3 discusses the structure of XENON1T’s events in detail.

XENON1T is the latest incarnation in the XENON program of liquid xenon TPCs, and the largest xenon TPC ever built. It uses 3.2 tonne of xenon, of which 1.2 tonne sits outside the TPC, acting (primarily) as a passive shield against external backgrounds. The GXe in the TPC sits under a diving bell, which can be separately pressurized to allow LXe shielding all around the TPC. Due to limitations of the cryogenics and purification systems, the LXe

is currently not filled above the bell, and the top of the detector is exposed to a relatively high external background.

Two arrays of 3-inch diameter Hamamatsu R11410-21 PMT tubes [128], 248 in total, observe the TPC from above and below. The TPC wall is lined with the highly reflective material PTFE, for maximal light collection. The electric fields in XENON1T are maintained by five electrode meshes (also shown in figure 2.9) and field shaping electrodes in the TPC wall. The whole TPC sits inside a stainless steel cryostat. All of XENON1T’s components in or near the TPC, from PMTs to the cryostat steel, are made from extremely-low radioactivity materials, and underwent extensive screening to measure (and when possible, select) their radioactivity levels [129, 130].

The XENON1T cryostat sits inside a cylindrical water tank of 9.6 m diameter and 10 m height. While its primary use is as a passive shield against external backgrounds (without which the event rate would be ~ 200 times higher), it is also instrumented with 84 PMTs to detect muons (which can generate spallation neutrons, see section 2.3.4) via their Cherenkov light [131].

XENON1T’s xenon is continuously circulated through a purification loop, which includes zirconium getters and a distillation column for online removal of ^{nat}Kr [132]. The column can also be inverted [133] to remove the ^{222}Rn (whose beta-decay progeny causes most of XENON1T’s low-energy background) which is outgassing into the LXe from the piping and pumps of the support system. The xenon is cooled by pulse-tube refrigerators, with a passive liquid nitrogen cooling system as a backup. The cryogenics and purification systems are located in a service building adjacent to the water tank, along with the DAQ hardware and event builder servers (see section 4.1.1).

2.4.2 Detector response calibration

To relate observed S1 and S2 signals to produced photons n_{ph} and electrons n_e , we need the mean detected S1 signal size per produced photon $g_1 = \overline{S1}/n_{\text{ph}}$ and mean S2 signal size per produced electron $g_2 = \overline{S2}/n_e$. If the light and charge yield of LXe were known accurately, these could be extracted from any sufficiently clear mono-energetic signal (e.g. a ^{137}Cs or ^{83m}Kr calibration source). Given the complications discussed in section 2.3 however, a method that makes fewer assumptions on the LXe yields is desirable, and was introduced in [134].

We can calibrate g_1 and g_2 from several mono-energetic electronic recoil calibration measurements using the fact that the sum of ER light and charge yields is constant. Substituting the definitions of g_1 and g_2 into equation 2.24 and rearranging gives:

$$\frac{\overline{S2}}{E} = \frac{g_2}{W} - \frac{g_2}{g_1} \frac{\overline{S1}}{E}. \quad (2.29)$$

Since $\overline{S2}$ and $\overline{S1}$ are measurable for any clear spectral feature, while E is known from external measurements, we can estimate g_1 and g_2 if we can

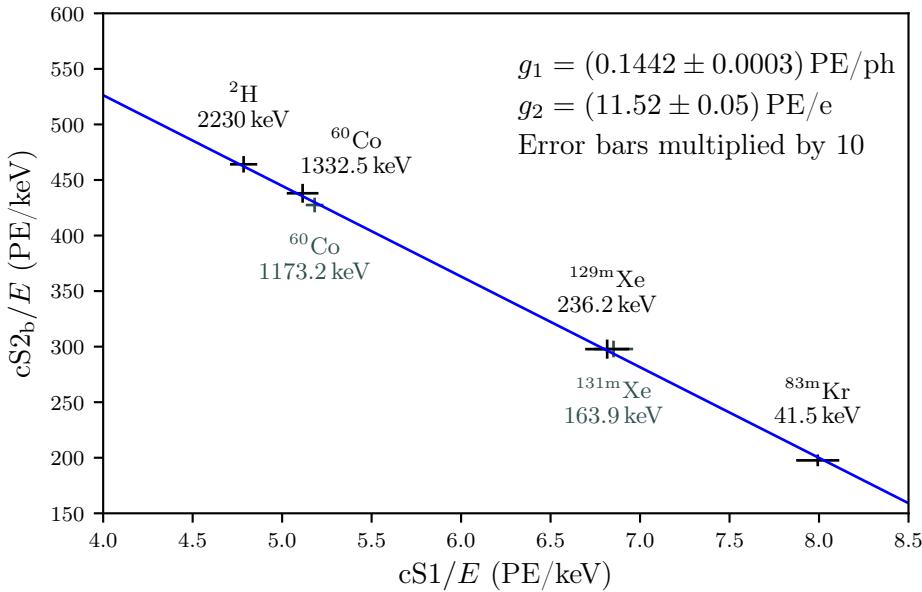


Figure 2.10: Adapted from [135]: Calibration ('Doke') plot of XENON1T. Markers correspond to the mean observed signal of monoenergetic electronic recoil calibration sources, the blue line is a least-squares straight line fit. The coefficients g_1 and g_2 are extracted using 2.29. Error bars represent statistical uncertainties, and are multiplied by 10 to make them more visible. The statistical errors on g_1 and g_2 are as listed in the figure. The ^2H datapoint was excluded from the fit, as it might be affected by PMT saturation effects (though no such effect is evident from this plot).

measure features at two or more ER energies, as shown in figure 2.10.

The detection efficiency for the light and charge signal is not perfectly homogeneous throughout XENON1T. Impurities in the LXe reduce the charge yield from deeper in the detector, and sagging of the anode mesh (see section 3.4.2) increases the electron extraction efficiency in the center of the detector relative to regions near the edge. Light yields are affected by spatial variations in light collection efficiency and variation in quantum efficiencies between PMTs.

Corrections for these inhomogeneities can be derived straightforwardly from measuring the spatial variations in the detector's response of a monoenergetic feature. For most corrections, data from the internal $^{83\text{m}}\text{Kr}$ source (see 2.3.3.3) is used, as it offers a prodigious event rate everywhere in the detector. The S1 and S2 size after corrections are denoted cS1 and cS2. In many analyses, an S2 estimate derived from correcting only the bottom-array response is used, denoted cS2_b. Since S2 signals are produced close to the top array, the top array's S2 response changes quickly near inactive top-array

PMTs.

Drift field inhomogeneity poses a more difficult problem: this affects the LXe light and charge yield both (though only slightly, especially for nuclear recoils), but more importantly, complicates interaction position reconstruction by altering the drift paths of electrons. Significant drift field inhomogeneity thus jeopardizes the derivation of all spatial corrections. To counteract this, we match electric field simulations of the detector to the observed event distribution of a homogeneously distributed calibration source, then correct the event positions appropriately for the altered drift paths. In XENON1T’s first science run, the drift field was relatively homogeneous at (116.7 ± 7.5) V/cm, so this is only a minor complication in the analysis.

Besides these detector calibrations, experiments usually try to measure the light and charge yield of LXe in-situ with their own detector, rather than relying solely on external measurements and semi-empirical models (see section 2.3). For XENON1T, this is discussed in section 6.2.

2.4.3 Background discrimination

To achieve world-leading dark matter sensitivity, merely building a large detector alone is insufficient: scrupulous care is needed to minimize all alternative sources of signals, discussed in sections 2.3.3 and 2.3.4. Much of the potential backgrounds are mitigated by the design choices discussed in section 2.4 – material screening campaign, the water shield / muon veto, and the passive shield of LXe outside the TPC. However, about ~ 5 Hz of background signals still remains, which must be discriminated from the dark matter signal by a variety of event selections.

Figure 2.11 shows the sequence of event selections we applied in XENON1T’s first science run (discussed in chapter 5) to select $\mathcal{O}(1)$ dark matter candidates (bottom left) from $\mathcal{O}(10^7)$ total events observed. The most significant reduction comes from removing high-energy signals (conventionally by their S1 signal size), which reduces the event rate by a factor ~ 500 . Substantial reductions are also made by restricting attention to the inner volume of the detector, called the fiducial volume, by removing events near the detector walls, the liquid/gas interface (the top of the detector, which as mentioned lacks a passive LXe shield) and the cathode mesh (which is a source of radioactivity and provides fringe fields that lead to unusual charge and light yields).

The other event selections shown in figure 2.11, with the exception of the last step, are data quality cuts designed to remove events that are not true single scatters in the liquid xenon. The most significant of these are cuts against ‘events’ that consist only of piled-up single electrons from delayed extraction of previous events (see section 3.4), and cuts against interactions produced in the GXe between the anode and top screening mesh, which have an unusually high fraction of S2 observed by the top array. The full list of cuts is discussed in section 5.2.5.

As the last discrimination step, we compare the remaining events against a two-dimensional ($cS1$, $cS2_b$) background model using a profile likelihood (see section 5.4.3). This efficiently separates electronic from nuclear recoils: the latter have a much higher $S1$ signal than electronic recoils with a similar $S2$ size (see figure 5.2). A slightly weaker discrimination can be achieved with a fixed cut on the $S1$ and $S2$ signal sizes ('box analysis'); for illustration, figure 2.11 shows the effect of the reference box presented in section 5.3.

Since less than one in a million of the events seen by XENON1T is a potential dark matter candidate, reliably isolating these candidates from numerous physical and pathological backgrounds requires detailed knowledge of the $S1$ and $S2$ signals that make up an event. This is the focus of the next chapter.

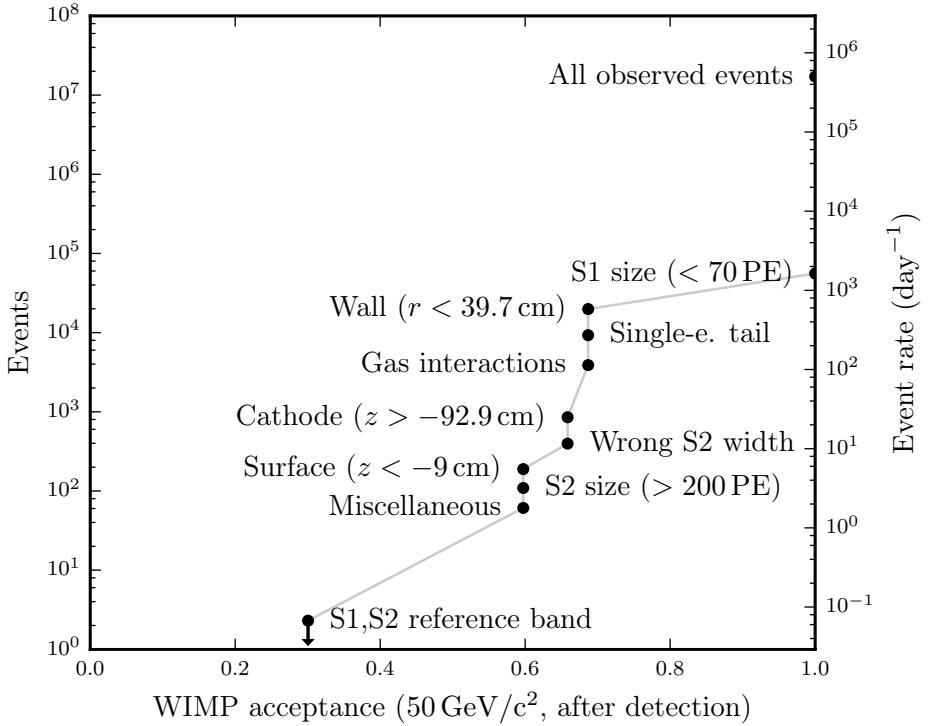


Figure 2.11: Remaining events (left y-axis), event rate (right y-axis) and WIMP acceptance (x-axis) at various stages (top right to bottom left, following the line) of signal/background discrimination during XENON1T’s first science run (see chapter 5). Each point shows an event selection, listing the selection type or primary background removed. The selection order is chosen to maximize the total event rate loss in each step. The acceptance losses of minor selections (any except volumetric cuts and the reference band) are not shown; these total to $\sim 10\%$ (see figure 5.1). The final discrimination step is a profile likelihood in $(cS1, cS2_p)$; we show a reference band cut here (see section 5.3 for the region’s definition). Since no events are observed in this region, we show a 90% confidence level upper limit on the event rate.

Chapter 3

Signals of the XENON1T detector

3.1 Introduction

This chapter is a bottom-up analysis of the signals in XENON1T, from electronic noise and single photon hits to the shape of our primary signals (S1 and S2). The goal is a description of our signals that will allow efficient processing and simulation of signals, as implemented in the XENON1T processing and simulation package `pax` [5]. A technical description of `pax` and the XENON1T waveform simulation routines is beyond the scope of this work; the interested reader can refer to [126] and the software documentation (both the software and documentation are open-source at github.com/XENON1T/pax).

3.1.1 Example event

Figure 3.1 shows an example XENON1T event from a neutron calibration. Several signals are evident: the initial short, small, relatively sharp signal at $\sim 2\text{ }\mu\text{s}$ is the S1 signal (from photons produced at the initial interaction) and the much larger peak is the S2 signal (from secondary scintillation of the ionization electrons, drifted towards and extracted into the gas phase). Notice the S2 signal is strongly peaked in the top array, while the S1 light is seen mostly in the bottom array due to total internal reflection on the liquid/gas interface – even though, as we will see in section 3.5.1, events with this short a drift time should be at most a centimeter below the gate grid. Due to the shallow depth, the S1 light that does make it to the top array is somewhat peaked at the same location as the S2.

Smaller signals are also visible:

- A secondary signal just after the S1 has roughly the size and shape of secondary scintillation due to a single extracted electron (see section 3.4) [136, 137]. It was likely created by photoionization of an S1 photon on

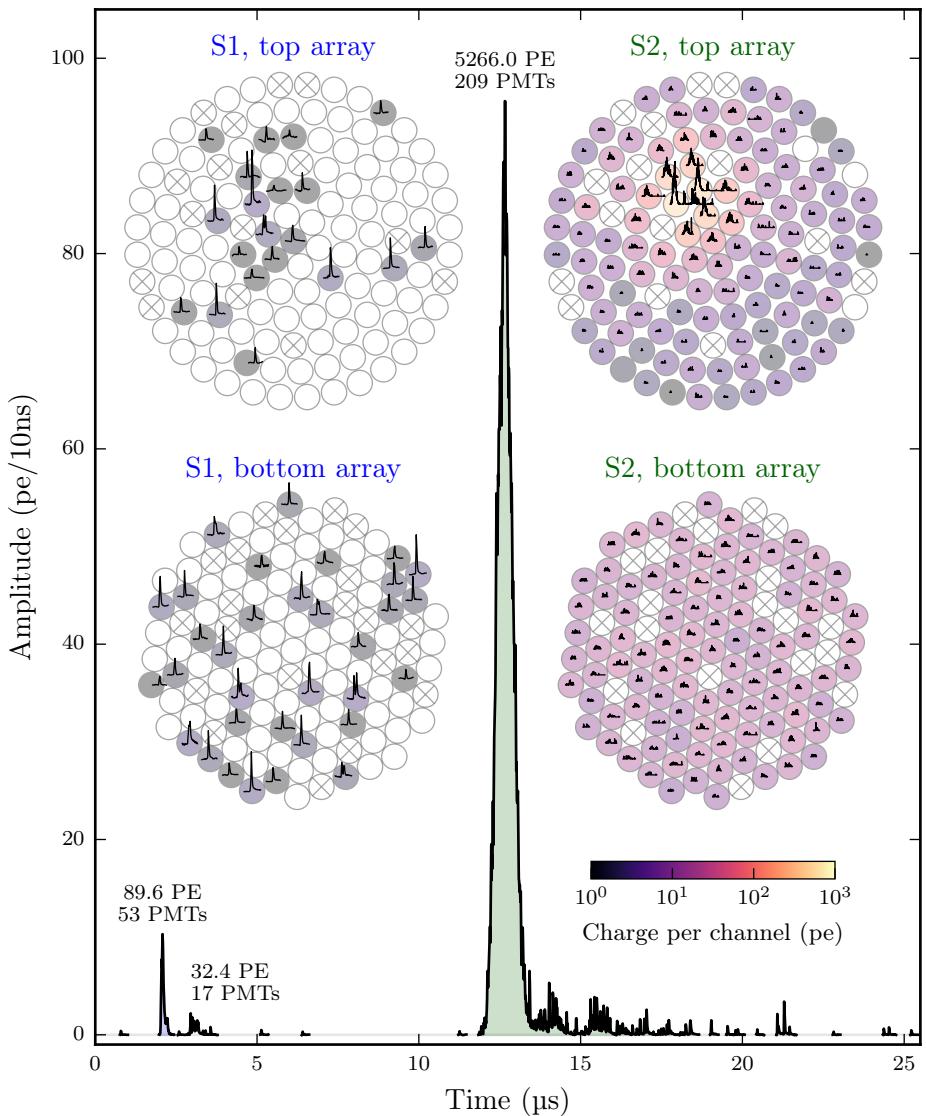


Figure 3.1: Example nuclear recoil event from XENON1T (from DD-generator calibration data), about a centimeter below the gate mesh. Main panel shows the (zero-suppressed) summed waveform of all PMTs, with S1 and S2 indicated with blue and green shading, respectively. Insets show the response of the individual PMTs to the S1 and S2, with colors indicating the charge observed per PMT. Inactive PMTs are indicated with a cross. Waveforms in the S1 insets are not to scale with those in the S2 inset. A single electron signal is seen just after the S1 (likely from photoionization). Unlabeled small signals are not coincident over PMTs, and likely consist of PMT afterpulses.

the gate mesh or an impurity in the xenon.

- Tiny signals between the S1 and S2 are seen by a single PMT, and are likely produced by gas molecules inside a PMT, ionized by either a UV photon that passed the photocathode or by a photoelectron cascade of a regular PMT hit, which gradually drift towards then strike the photocathode ('PMT afterpulses') [138].
- The tail of signals after the S2 is likely a combination of these effects. High-energy S2s are always followed by a long tail of single-electron signals caused by photoionization of the S2 light [137]. An even longer tail of single electrons may be present in some cases due to delayed extraction of a small fraction of the initial S2's charge[139, 140].

The event in figure 3.1 is not a typical event in some respects: as discussed in section 2.4.3, most events have much larger signals (and are therefore of no interest to dark matter searches). The short drift time between S1 and S2 indicates the event is close to the liquid/gas interface, which is convenient for visualization. However, dark matter events can occur anywhere in the detector.

3.1.2 Model development

Signal processing and simulation both depend on accurate signal models, creating an apparently circular dependency: an accurate signal processing program ('processor') is needed to extract unbiased models; but only an accurate simulator (informed by accurate models) can confirm the reliability of the processor. In practice, signal models and signal processing systems develop in lockstep, with several safeguards in place to ensure each step represents progress rather than a regression towards mutual biases:

1. Simulated signals contain intrinsic variability. For example, the S2 shape varies with the depth of the particle's interaction in the TPC (as discussed in section 3.5), and the shape of any small simulated signal (S1 or S2) is heavily stochastic. This forces the processor to be sensitive to a broad range of signals, as otherwise its efficiency would drop even on the simulated data (generated using the expected model) it is evaluated on.
2. Where possible, we start from physically motivated models as opposed to effective models with no physical basis (though in a few cases these could not be avoided). This ensures only biases that mimic actual physical effects can enter our signal models. Moreover, some signal model parameters can be compared to literature, providing a further consistency check.
3. Human experts regularly inspect the processor's interpretation of events, especially in the context of (proposed) changes in the processing. This

can involve examining individual events (see e.g. figure 3.1) or examining aggregate features such as the linearity of the signal response over a large signal size range (see figure 2.10).

3.2 Electronic noise

The signals of any detector sit on top of electronic noise of various kinds, from electromagnetically induced pickup to thermally induced Gaussian jitter. Since noise spikes can mimic photo-electron hits, high noise levels force the use of high thresholds in the signal finding, which reduces the sensitivity of the detector to small signals such as the few-photon S1s of low-energy WIMP interactions. To compensate, the PMT operating voltages can be increased to get higher PMT gains, but this can reduce the lifespan of the PMTs and increases the dark count rate. Hence, dark matter experiments such as XENON1T seek to eliminate, reduce, or at least cope with the various varieties of noise.

To measure noise, we configure the digitizers to take a long event window of data at an externally triggered time. Figure 3.2 shows some example waveforms from XENON100 (an earlier detector which used similar readout electronics, see [141]) and XENON1T. In both detectors, we see a high-frequency beat pattern on top of a low-frequency (~ 25 kHz) sinusoid.

To decompose the noise into individual components, we Fourier transform the long event, as shown in figure 3.3. We see most noise originates at frequencies higher than a MHz, in a series of sharp peaks occurring in harmonic progressions (e.g. the peaks at 200 kHz, 400 kHz, 600 kHz, 800 kHz, ...). It is possible to combine information from several events to measure the spectrum more accurately [142], but we will not do so here.

To investigate the origin of individual components, we examine the phase correlations between the noise peaks in different channels, as shown in figure 3.4. The 25 kHz peak shows strong phase correlation among channels on the same digitizer. We also see the 25 kHz amplitude is larger in some digitizers than in others. It is therefore likely the cause of this noise is somewhere close to or inside the digitizers themselves.

For the high-frequency peaks (see e.g. figure 3.4, bottom for the strongest peak), we find instead a strong phase and amplitude correlation between channels whose PMTs connected to the same high-voltage supply module. In XENON100, low-pass filterboxes were installed on the PMT high voltage lines to limit such noise, but in XENON1T these were initially omitted. After these noise observations, new filterboxes were designed and installed [143] to limit the high-frequency power supply noise more effectively than in XENON100, as shown by comparing the blue (before installation) and green (after installation) noise spectra in figure 3.3. The change is quite visible in the raw waveform too, see figure 3.2 (middle and bottom panels).

To consider whether this noise has a relevant impact on the analysis, we

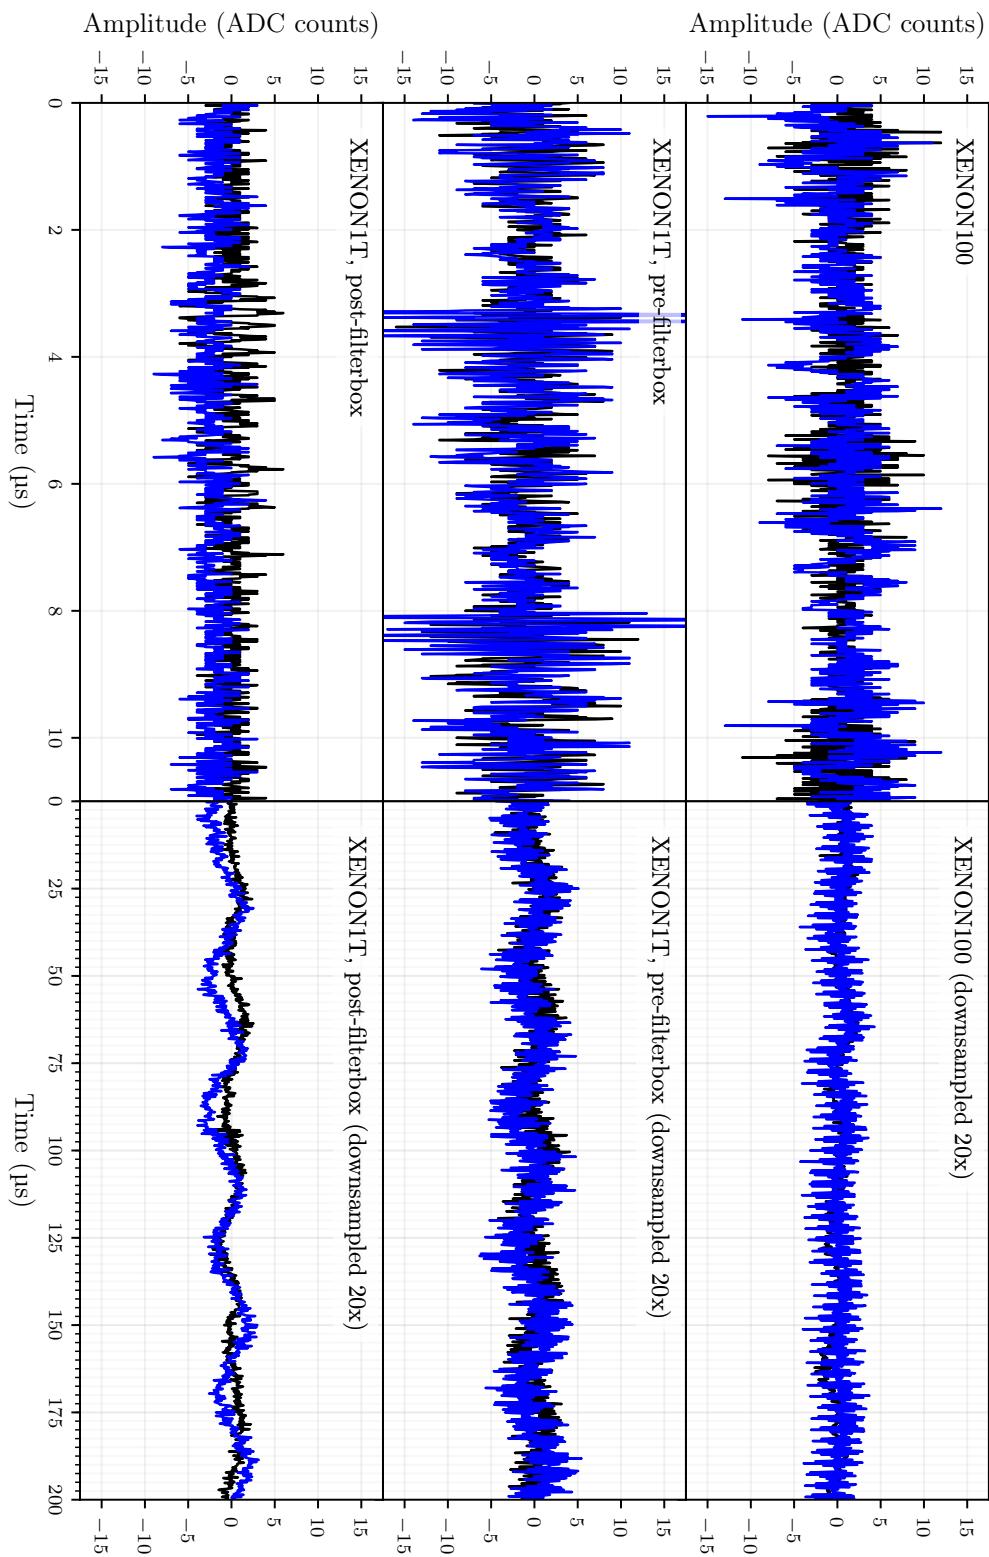


Figure 3.2: (on previous page) Externally triggered waveforms in PMT channels 20 (black) and 200 (blue) from XENON100 (top, in figure's orientation), XENON1T before (middle) and after (bottom) installation of filterboxes on the PMT high-voltage power lines. Left panels show the original waveform, right panels show the same waveform over a larger time range, downsampled to better show the low-frequency behaviour. The signal-like direction is towards more positive ADC values (i.e. the original ADC signal has been inverted after subtracting the baseline).

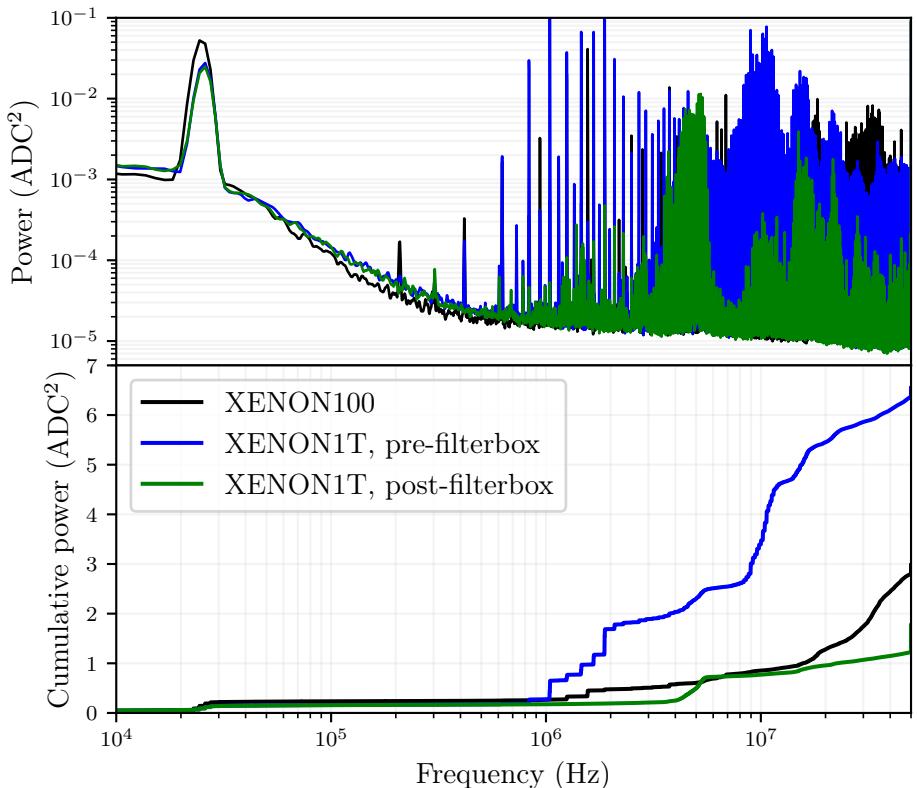


Figure 3.3: Top: noise power spectrum (averaged over all channels) in XENON100 (black), and XENON1T before (blue) and after (green) installing filterboxes on the PMT high-voltage lines. Bottom: same, cumulative sum over frequencies.

must combine it with information on the expected single-photoelectron (PE) signal shape, derived in section 3.3.3; we merely note here that the mean single-PE pulse amplitude is approximately 55 ADC counts. Due to the digitizers’ zero-suppression (see section 4.1.2), we do not ordinarily see data stretches longer than a few μs ; this means the impact of the low frequency component is minimal as long as we determine the baseline of the signal for every individual stretch of data. The high-frequency components, however, force us to record a substantial time (currently 0.5 μs) ahead of any signal of interest to do this reliably.

3.3 Single photons and photoelectrons

All the signals of interest in XENON1T are built from single-photoelectron (single-PE) PMT signals. Describing these is a fundamental prerequisite for accurate signal simulation and processing, and is the goal of this section.

3.3.1 PMT gain calibration

Due to the stochastic PMT amplification process, each photoelectron causes a signal with a different charge. The mean charge detected per photoelectron is called the PMT’s *gain*. The raw charge of a PMT signal divided by the gain is an estimate of the number of photoelectrons (PE) of which the signal consists, which can eventually be used to estimate the number of detected photons (or even the number of photons produced at the interaction)

In XENON1T, PMT gains are measured by illuminating the PMTs with flashes of 406 nm [144] LED light, guided into the TPC by optical fibers. A PMT receives zero, one or multiple photons during the flash, the probability of which depends (among other things) on the light level of the LED. The resulting distribution of charge integrals over a small duration around the time when the LED flashes can be fitted with a multi-peak functional model to extract the gain and other parameters such as the fraction of flashes yielding two detected photons [145]. Often experiments assume a Gaussian photoelectron response; more advanced models account for under-amplified hits caused by photons hitting the first PMT dynode directly, or even attempt a full simulation the photomultiplier amplification process [146].

In XENON1T, we determine the PMT gain without fitting a functional response, using the method described in [147]. First, we measure the charge integral distribution when the LED is off (noise-only) and when it is on. The ratio between the on- and off spectra at very low charge is an estimate of the fraction F_0 of LED flashes in which the PMT detects no photon. Since the hit multiplicity is Poisson distributed, this gives us the mean number of photons detected per flash $\lambda = -\ln F_0$. Then, assuming a linear response of the PMT,

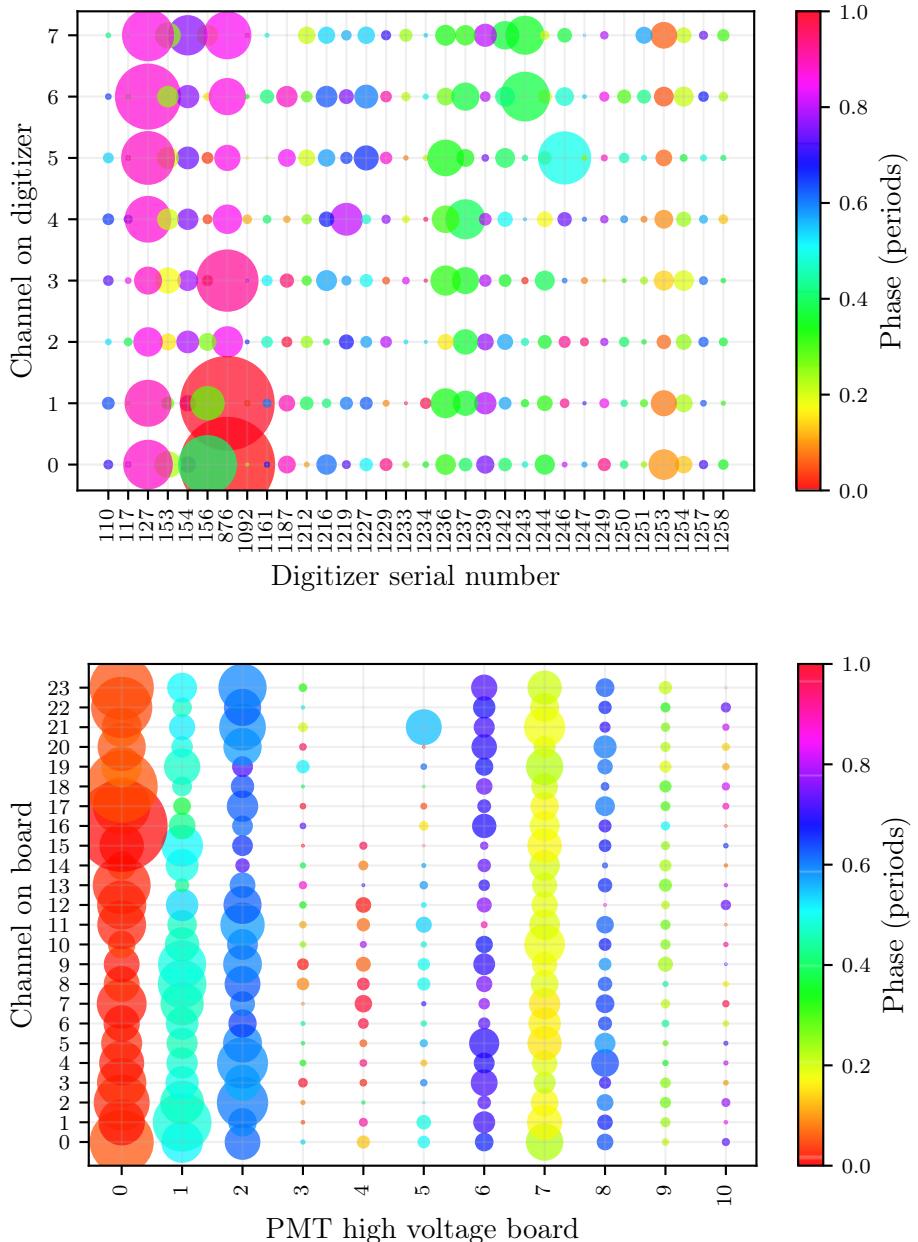


Figure 3.4: Phases and amplitudes of the 25 kHz (top) and 1.875 MHz (bottom) in the XENON1T noise, before filterbox installation. Each channel is represented by a circle whose color shows the phase of the Fourier component and whose size shows its relative strength (compared to other channels, not to scale between top and bottom figure). The x- and y-axes show the digitizer connectivity (top, 25 kHz) and PMT high voltage connectivity (bottom, 1.875 MHz) respectively.

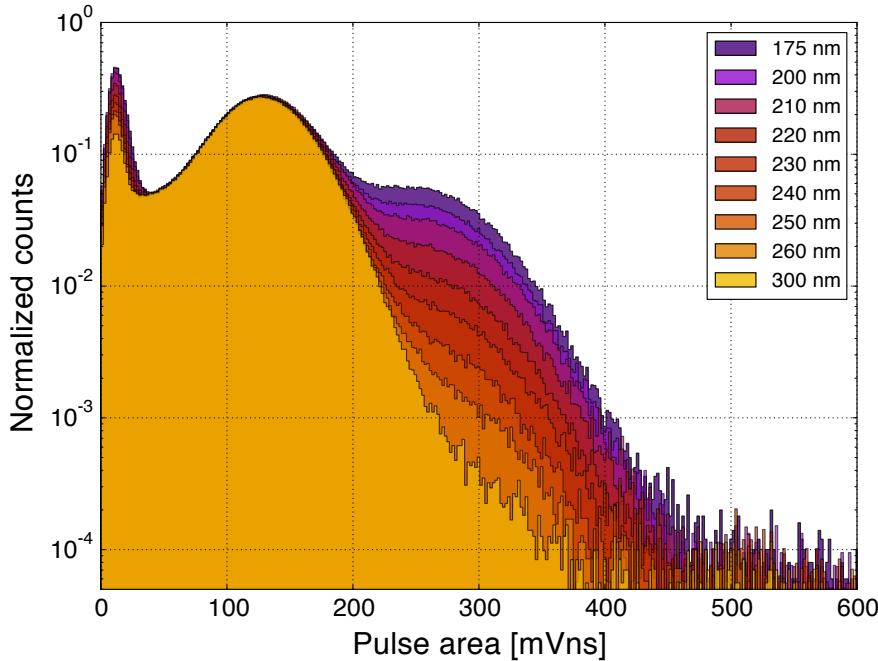


Figure 3.5: From [148]: charge signal distributions from a R11410 PMT illuminated by light of different wavelengths. The spectra are normalized in the 50 - 230 mV \times ns range for visualization purposes.

the gain μ can be recovered from directly observable quantities:

$$\mu = \frac{\mu_{\text{on}} - \mu_{\text{off}}}{\lambda} \quad (3.1)$$

where μ_{on} and μ_{off} are the observed mean charge integral with and without LED light.

3.3.2 Double photoelectron emission

When struck with a single 175 nm photon from xenon scintillation light, XENON1T’s R11410 PMTs (as well as other PMTs) emit two photoelectrons (instead of one) 18 – 24% of the time [148], as shown in figure 3.5. This probability is relevant for estimating the number of photons making up a signal from its observed charge integral. Since XENON1T lacks an externally controllable variable-wavelength light source, a precise determination of this probability is difficult. In this section we will qualitatively confirm the effect is present; for simulation purposes we use the measurements from [148].

We can study double photoelectron emission using the hit area spectrum from liquid xenon scintillation signals such as S1s (section 3.6) or single

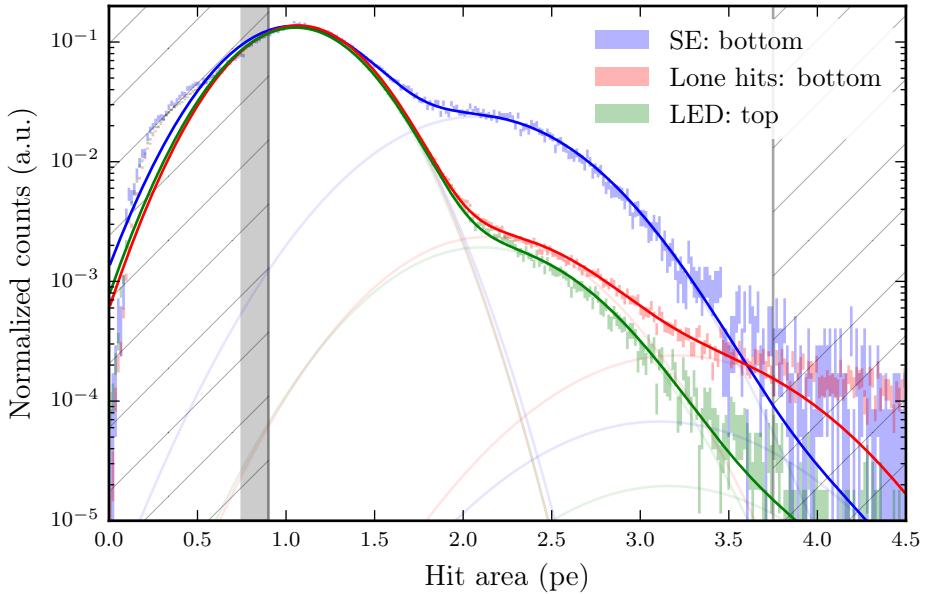


Figure 3.6: Histograms: integrated charge (in photoelectrons) spectra for 178 nm (LXe scintillation) light from single electrons hits in the bottom array (blue), for lone hits in the bottom array (red) and 406 nm LED hits in the top array (green). Bar length indicates statistical uncertainty. Solid lines show a triple Gaussian fit (equation 3.2). Thin solid lines indicate individual Gaussian components. Hatched regions (< 0.9 PE and > 3.75 PE) are excluded from the fit. The grey region between $0.75 - 0.9$ PE is used to normalize the spectra for visualization purposes.

electrons (section 3.4.2). Since we do not know in advance when these signals occur, we have to rely on the digitizer’s self-trigger (see section 4.1.2) to obtain the hits, and we will not see very small hits, such as most of the under-amplified photoelectron signals.

To distinguish double-PE emission from pile-up of single-PE hits, we should select signals that are well-spread in time over as many PMTs as possible. Single electrons fill this niche well: they are a few 100 ns wide (see section 3.4.2), and on average contribute just 8 photons to the bottom PMT array, where they are relatively homogeneously distributed over 121 PMTs. If hits arriving at the same PMT were always merged, the pile-up probability would be $\sim 7\%$, but clearly this is a rather conservative upper bound.

Figure 3.6 shows the area spectrum for hits in single electrons in the bottom PMT array. A substantial double-PE hit component is indeed seen. To analyze whether this could be due to pile-up, we fit the observed charge spectrum with a simplified model. Assuming under-amplified hits and pile-up from four or more hits can be neglected in the chosen fit range (0.9 - 3.75 PE)

| Source (array) | N (10^6) | χ^2/n_{df} | μ (PE) | σ (PE) | f_2 (%) | f_3 (%) |
|----------------|----------------|------------------------|------------|---------------|-----------|-----------|
| LED (T) | 0.84 | 1.8 | 1.053(1) | 0.328(1) | 2.02(3) | 0.025(7) |
| LED (B) | 2.41 | 3.6 | 1.057(1) | 0.3251(8) | 6.10(4) | 0.32(1) |
| Single e. (T) | 0.15 | 1.3 | 1.041(3) | 0.351(2) | 20.5(1) | 1.37(8) |
| Single e. (B) | 0.11 | 1.4 | 1.036(4) | 0.340(2) | 20.0(2) | 0.07(5) |
| Lone hits (T) | 1.42 | 2.5 | 1.068(1) | 0.3258(9) | 3.49(4) | 1.07(2) |
| Lone hits (B) | 1.50 | 2.0 | 1.068(1) | 0.3236(8) | 2.36(2) | 0.30(1) |

Table 3.1: Best-fit parameters for single-hit charge distributions for different sources and PMT arrays, with T denoting the top array and B the bottom. N denotes the number of hits in the sample, f_2 and f_3 denote the fraction of double- and triple photoelectron hits, respectively. Uncertainties are derived from covariance estimates from the least-squares fit and do not account for the approximate nature of the model.

and the charge of (fully amplified) hits is normally distributed, the spectrum will be:

$$(1 - f_2 - f_3)G(\mu, \sigma) + f_2G(2\mu, \sqrt{2}\sigma) + f_3G(3\mu, \sqrt{3}\sigma). \quad (3.2)$$

Here G is the Gaussian density function, with μ and σ the mean and spread of fully amplified single-PE areas, and f_2 and f_3 the fraction of double and triple photoelectron signals. An overall normalization constant is omitted. Note that this model is simplistic, and parameter estimates obtained from it could have substantial systematic errors. For example, if the single-PE hit area distribution has a longer tail than a Gaussian, f_2 and f_3 will be overestimated.

The best-fit function is shown as a solid line in figure 3.6, the best-fit parameters are listed in table 3.1. The $\sim 20\%$ fraction of double-PE hits is much larger than the expectation from pile-up alone, and consistent with the double-PE emission measurements from [148] if pile-up double-PE hits contribute less than 2%. A precise estimation of the double-PE fraction would require refined estimates of the pile-up rate, e.g. based on simulated single electrons (see section 3.4.2).

To show that the double-PE emission is wavelength-dependent, we can repeat the same procedure on LED data (processed in the same way to ensure a proper comparison). Due to the placement of the optical fibers guiding the LED light, the bottom array PMTs are more strongly illuminated and will show more pile-up. The charge distribution for top-array hits is shown in green in figure 3.6. Clearly double-PE emission is at most 2% for LED light, and likely less, if some of this is due to pile-up. We could verify this by reducing the LED light intensity to suppress pile-up altogether.

We can also repeat the procedure on lone hits (hits with no nearby other hits) observed in background data, as shown in red in figure 3.6. Again, no appreciable double-PE emission is seen, while the μ and σ obtained are not far

from those of other populations. This strongly suggests the dominant origin of these hits is not xenon scintillation light or electronic noise, but thermal emission of photoelectrons from the photocathode, a known component of the count rate of unilluminated PMTs [138].

3.3.3 Single photoelectron pulse shape

Electrons take ~ 100 ns to travel through XENON1T’s large PMTs [128], so even a slight spread in possible paths leads to a detectable time spread of the single-photoelectron signal. For XENON1T’s PMTs, the Hamamatsu datasheet lists a 5.5 ns PE-pulse rise time [149]. A fall time is not listed, but this is typically larger than the rise time [138]. We should thus be able to discover some structure in the single-PE pulse shape, even with our 10 ns digitizer resolution, especially considering the intermediate electronics (PMT base, pre-amplifier, digitizer) likely cause additional pulse broadening. Since S1 signals are quite sharply peaked in time (see section 3.6), simulating realistic single-PE pulse shapes is a prerequisite for simulating realistic S1s.

To find the single-PE pulse shape, we extract one million single-PE waveforms from isolated pulses in regular science data events (this takes about 5000 events, obtained in 20 minutes). As discussed above, lone hits arise predominantly from thermal emission of photoelectrons from the photocathode (see figure 3.6, red, for their charge distribution). We select pulses between 0.5 and 1.8 PE to suppress potential small contaminations from noise and double hits. Channels whose PMTs are inactive (36 in this dataset) are excluded, as are three additional channels (PMTs 31, 33 and 208) whose waveform shape deviates substantially from that of the remaining channels.

The digitizer self-trigger threshold (see section 4.1.2) could affect our extracted pulse shape if we select pulses whose amplitude occasionally falls below the threshold (for example, for a sharply peaked signal, if its charge is mostly distributed over two digitizer bins instead of one). The 15 ADC count threshold corresponds to 0.11 PE/bin at XENON1T’s median PMT gain of 2.4×10^6 electrons per photoelectron. We will see below the amplitude of SPE signals is about 40% of its integral; such signals will pass the ADC threshold if their charge is 0.26 PE or larger. Thus, signals larger than 0.5 PE essentially always pass the ADC threshold, and our analysis is not affected.

Figure 3.7 shows the average normalized waveform of the remaining 870257 pulses (aligned on the maximum sample) in solid black. A late-time component extends to about 170 ns after the maximum sample, with 16% of the charge in the mean waveform arriving more than 50 ns beyond the maximum sample. This is significant compared to the 3 ns and 24 ns decay constants of LXe scintillation (see section 2.3.1). Most channels even show a secondary spike around 110 ns. The cause is unknown, but we can exclude electronic noise or baseline errors (those would also affect the waveform before the hit) and rare outliers such as waveforms with multiple hits (by examining individual

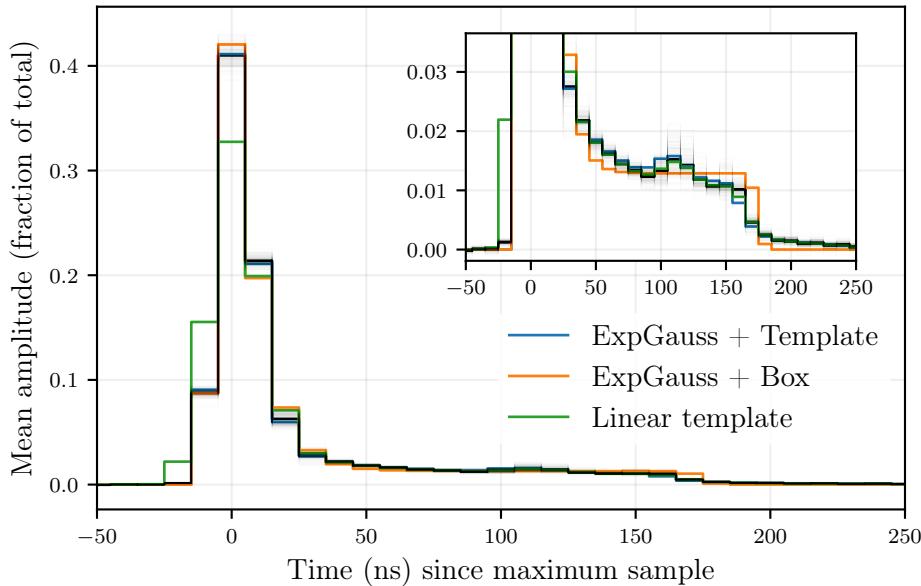


Figure 3.7: Black: Average shape of normalized lone-hit PMT waveforms, aligned on the sample with maximum amplitude. Statistical errors are smaller than the line width. Thin grey lines: same, for individual PMTs. Colors: simulated waveform shape for three single-PE pulse models (described in the text). Inset: same figure, zoomed in to low amplitudes.

waveforms, or the median instead of the mean waveform). Possible hypotheses include electrons that take a suboptimal path through the large PMTs or a signal shaping effect of the electronics (PMT bases, amplifiers, digitizers). Studying the individual hardware components in isolation could resolve this.

Per-channel average waveforms are shown as thin grey lines in figure 3.7; clearly these are roughly similar, with a few exceptions:

- The amplitude at 20 ns seems to divide the channels into two groups (too small to clearly see in figure 3.7), corresponding roughly (but not exactly) to the top (low amplitude) and bottom PMT array (high amplitude).
- Channels vary in the extent to which they show a secondary spike at ~ 110 ns. No obvious correlation between the variations and the cabling map was found.
- At 150 ns, all channels on two amplifiers (serial numbers 24678 and 24681) show a sharp (but small) third spike. A possible explanation is a mismatch between the load resistance of these amplifiers and the PMTs, allowing reflected signals to travel between the PMTs and amplifiers.

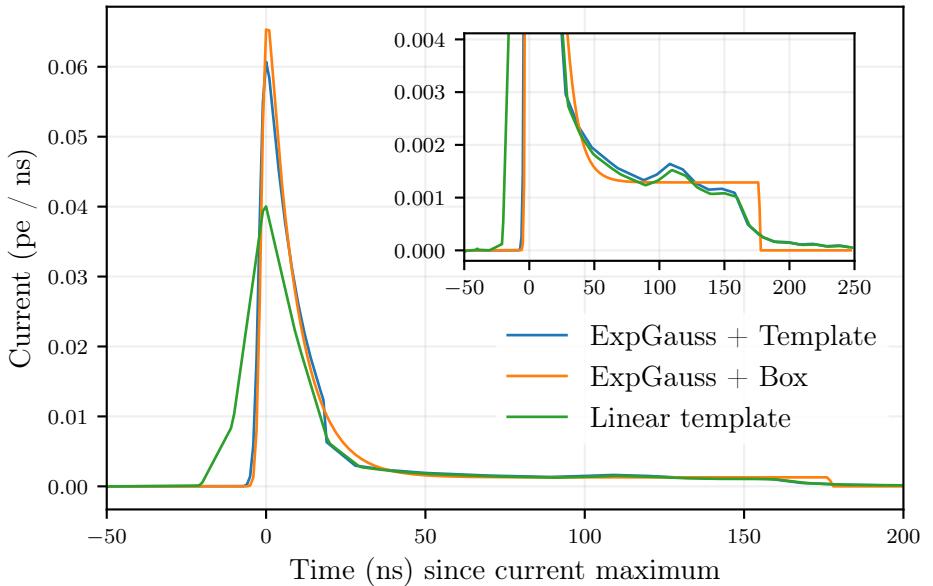


Figure 3.8: Current vs. time for single-PE models described in the text.

Since most channels contribute about 4000 pulses each, the average of per-channel averages is extremely close to the average waveform computed on all pulses.

For accurate simulation, we must reconstruct the original analog current signal corresponding to a single photoelectron. To do this, we use a simulation of the digitizer (assumed to be a perfect integrator, then see which current model best reproduces the measured averaged digitized waveform of figure 3.7. This does not have a unique solution: a series of Dirac delta functions separated by 10 ns will reproduce the digitized waveform (if their integrals are chosen appropriately), but is clearly unrealistic. We must strike a balance between obtaining a good fit and choosing a reasonable model.

Three models are shown in figure 3.8: linearly interpolating the average waveform (green), an exponentially modified Gaussian plus a boxcart function to describe the late component (orange), and an exponentially modified Gaussian that switches to the interpolated template at later times (blue). The resulting average waveforms are shown in figure 3.7 in the same colors. For the latter two models, free parameters were optimized by minimizing the Kolmogorov-Smirnov (KS) distance between the observed averaged waveform (black) and that of the model.

We see the linearly interpolated model (green) is unsuccessful near the peak, but quite accurate for the late-time component. It results in a smoother waveform than observed. The functional model (orange) shows the converse performance: accurate in the peak, but a poor description of the tail. The

mixed model (blue) is most successful, achieving a KS-distance of 3.8×10^{-3} . For comparison, observing this value in a $N = 0.96 \times 10^5$ two-sample KS-test would yield a 50%-confidence rejection of the null hypothesis that both samples come from the same distribution. This number should not be directly compared to the number of waveforms used, as a waveform supplies much more information than a single random variable (it takes a value in each time bin).

The best-fit parameters are a 11.2 ns decay time and a 1.2 ns Gaussian smearing, with the template interpolation taking over at 20.5 ns (a clear discontinuity is visible in figure 3.8). Setting meaningful confidence intervals on these parameters would require a detailed study on the sources of systematic error; due the number of waveforms used statistical errors are extremely small. A more complex fit function might achieve a better fit or remove the discontinuity, but as we will see in section 3.6, this model is sufficient to obtain an accurate S1 waveform model.

3.4 Single electrons

Single electron (SE) signals are produced in copious amounts following large signals, primarily by photoionization of the initial scintillation light on exposed metal components in the detector and on impurities in the liquid [136, 137]. A longer-lived population of single electrons can occasionally be seen after large S2s; this is hypothesized to be due to delayed extraction of electrons from under the liquid-gas interface [140].

Regardless of their origin, SEs are of great use in calibrating the detector. The average charge integral of single electron μ_{SE} is called the *secondary scintillation gain*; combined with measurements of S2 gain g_2 from known mono-energetic peaks (see 2.4.2), this allows a measurement of the electron extraction efficiency $\eta = g_2/\mu_{\text{SE}}$. As discussed below, the width of single electron signals gives a sensitive measurement of the width of the layer of gas under the anode. This can be used to level the detector and detect warping of the anode mesh. Finally, since S2 signals consist of several single electron signals close together, the waveform shape of single electrons is a fundamental input for the waveform simulator.

3.4.1 Single electron gain

3.4.1.1 Data selection

Single electrons are extremely common, which makes it straightforward to select a sample of single electrons that is relatively pure (few contaminating backgrounds) and free of bias (e.g. no preference for large or small single electrons). Two significant backgrounds have to be avoided: pile-up of two single-electron S2s, arising from high-density single electron tails after large

S2s, and random clusters of single-PMT hits (e.g. from PMT dark rate or afterpulses). To this end, we select peaks satisfying three conditions:

1. No other peaks (other than lone hits) within 200 μ s (almost a third of the maximum drift time in XENON1T) on either side. This avoids selecting SE candidates from tails of high-energy S2s with a large amount of pile-up.
2. Number of hits in [12, 29]. The lower bound suppresses random hit clusters, the upper bound further suppresses double-SE pileup. This naturally biases the average single-electron size, but this can be accounted for in fitting.
3. Width of the signal between 30 and 500 ns, as measured by the interquartile range (see section 3.5.3). This is a mild cut that removes some outlying signals, such as when a large PMT afterpulse is merged with an ordinary SE.

Note no other selection criteria were used; in particular, we do not rely on `pax`'s classification of the peaks.

Figure 3.9 shows the peaks from a single one-hour run that pass the first condition. The single electron population is clearly visible at ~ 20 hits. The effect of the other two conditions is shown by the white box.

3.4.1.2 SE spectrum model

Figure 3.10 shows a histogram of the number of hits per single electron in an example run. We can extract the mean number of hits per single electron by fitting this distribution. In a perfect detector this would simply be a Poisson distribution, since secondary scintillation is a Poisson process. In a real detector, the detection efficiency and other factors affecting the single electron gain vary with position (see section 2.4.2). Thus, we should expect a Poisson distribution convolved with the distribution describing this variation of the SE gain in the detector. Here we take this distribution to be a Gaussian. Thus, we fit a Poisson convolved with a Gaussian, i.e.:

$$\text{PMF}(n|\mu, \sigma) \propto \int_0^{\infty} d\lambda \text{PoissonPMF}(n|\lambda) \text{NormalPDF}(\lambda|\mu, \sigma) \quad (3.3)$$

with n the number of hits and μ and σ the two parameters of the model. Notice this has just two free parameters, just like a Gaussian – and unlike the Gaussian, this model is physically motivated (as explained above). The scale of the model is fitted as an additional free parameter.

The fit of eq. 3.3 to the number of hits histogram is also shown in figure 3.10. This fit has $\chi^2/\text{ndf} = 16.0/15$, which has a p-value of 0.38, indicating

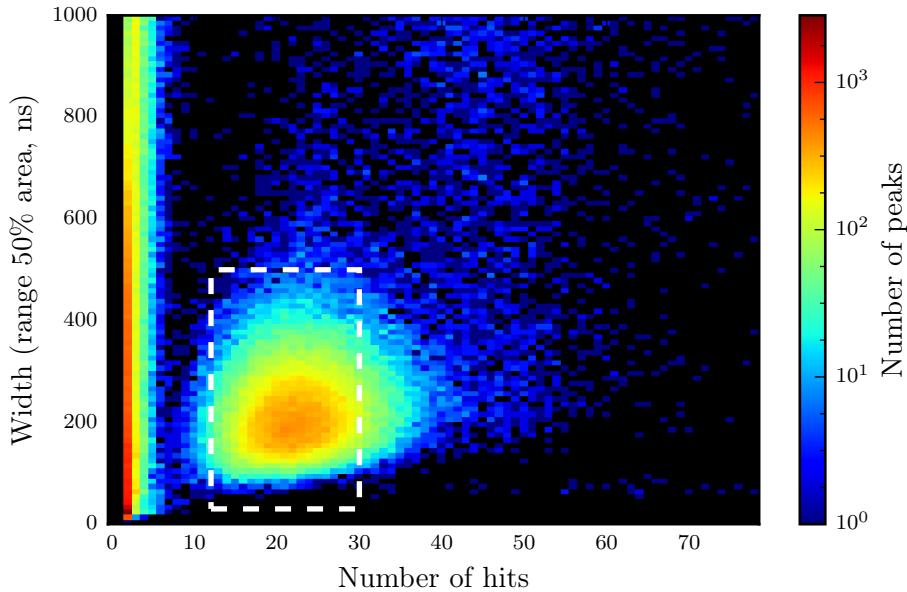


Figure 3.9: Histogram of isolated peaks (other than lone hits) during a one-hour run. The dashed box shows the remaining cuts used to select single electrons.

the fit is statistically limited. Using a simple Gaussian instead of eq. 3.3 gives $\chi^2/\text{ndf} = 121/15$, a considerably worse fit.

Figure 3.11 shows the result of applying this fit to all runs during and surrounding the XENON1T first science run. On average, the mean number of hits per single electron was 22.0, stable within 0.92 hits/SE (RMS). We clearly see a correlation with the gas gap size: the larger the gas gap in which the electrons luminesce, the larger the single electron size (see also section 3.4.2). A few jumps in the single electron gain are uncorrelated to gas gap size and can instead be related to known changes in the data acquisition settings or PMT gains.

Depending on the analysis application, we might want to convert this to an average observed single electron size in PE, using the mean area per hit in the single electrons. For analyzing the extraction efficiency, the mean observed charge in the bottom array is most relevant, which is 11.4 PE. Comparing this with $g_2 = (11.52 \pm 0.05)\text{PE}$ from section 2.4.2, we see XENON1T’s extraction efficiency is close to perfect; a more careful treatment accounting for the non-linear response of the data acquisition and processing systems (due primarily to the digitizer self-trigger threshold) finds an extraction efficiency of $(96.1 \pm 4.6)\%$ [135].

We might also use the observed variation in single electron size to correct S2 areas from different parts of the science run. In practice, the stability of

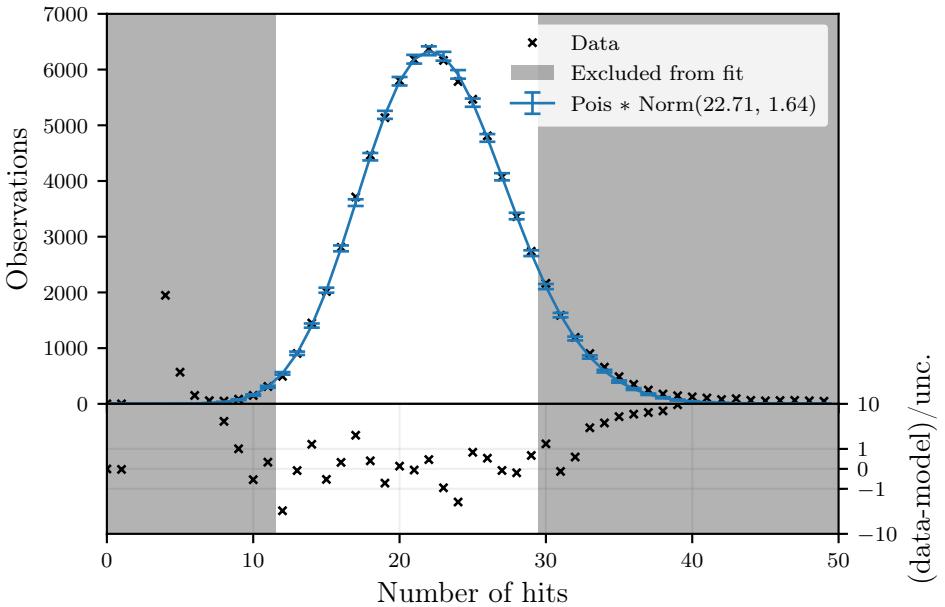


Figure 3.10: Black crosses: histogram of the 71739 single electrons from a one-hour background data run. Grey areas show the same histogram when the number of hits cut is relaxed; the points for 2 and 3 hits are off-scale. Blue curve: best-fit of eq. 3.3 to the data. Error bars show Poisson \sqrt{n} uncertainties. Lower panel: residuals, i.e. $(\text{data-model})/\text{uncertainty}$, for the fit shown. The y-scale is linear from 1 to -1 and logarithmic beyond.

the SE gain is such that this would have little effect on the energy resolution, especially at the low S2 sizes relevant for dark matter searches.

3.4.2 Single electron waveform shape

Single electron signals are created by the scintillation light produced as an electron travels the ~ 2.5 mm between the liquid/gas interface and the anode mesh wires, alternately accelerated by the strong electric field and decelerated as it collides with (and excites) the xenon atoms.

Scintillation light in GXe is produced by the same mechanism as LXe (see 2.3.1), but the intermediate state lifetimes are different. According to [100], the singlet lifetime is $\tau_1 = 5.88 \pm 5.5\text{ns}$, short enough that its actual value is a weak effect on the S2 shape. For the triplet lifetime [100] quotes $\tau_3 = (100 \pm 8)\text{ns}$ as an error-weighted world average.

The electric field in the gas affects single electron scintillation in two ways: through the motion of the electron and through changing the scintillation yield per unit length. Within the range of conditions relevant to XENON1T, the mean electron drift velocity v is approximately (from fitting the measurements

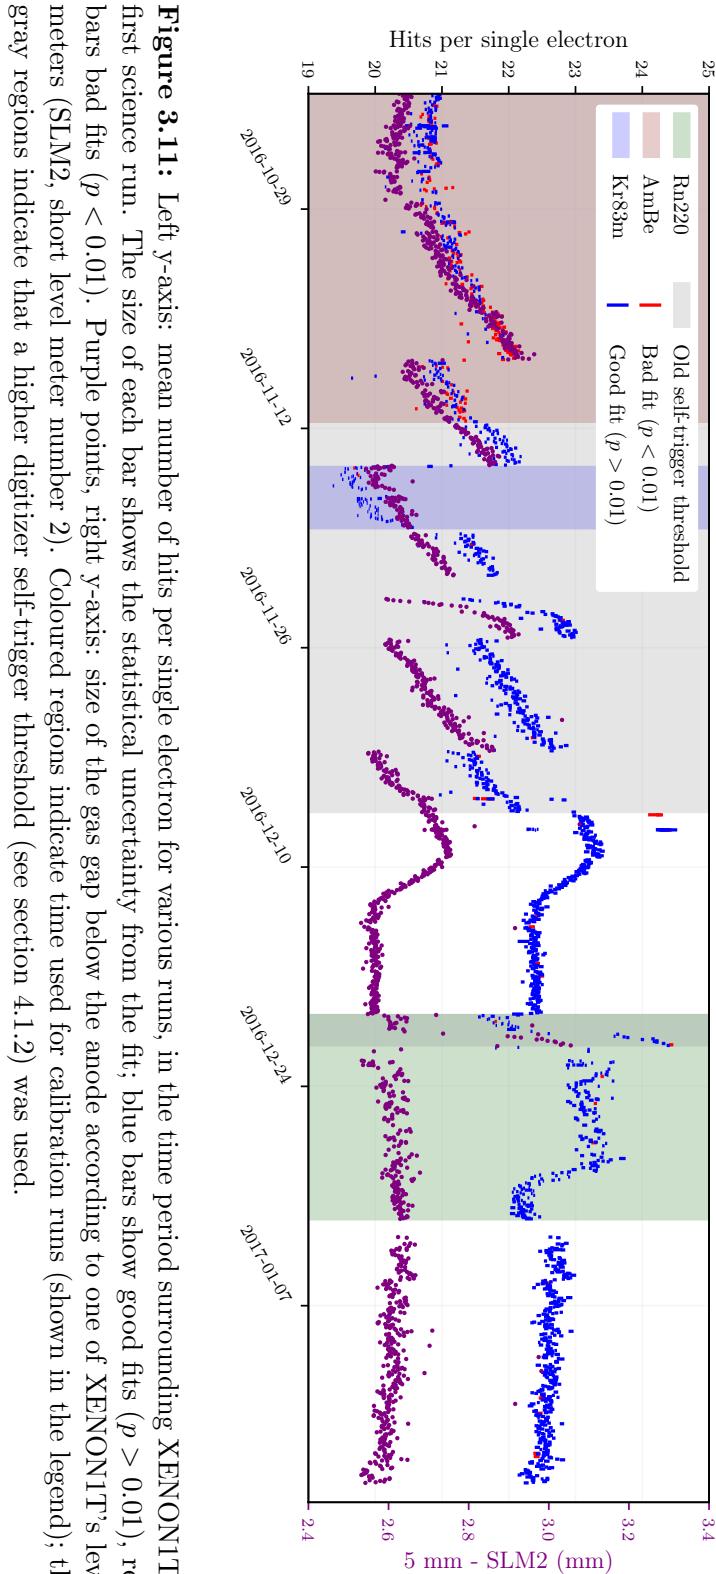


Figure 3.11: Left y-axis: mean number of hits per single electron for various runs, in the time period surrounding XENON1T’s first science run. The size of each bar shows the statistical uncertainty from the fit; blue bars show good fits ($p > 0.01$), red bars bad fits ($p < 0.01$). Purple points, right y-axis: size of the gas gap below the anode according to one of XENON1T’s level meters (SLM2, short level meter number 2). Coloured regions indicate time used for calibration runs (shown in the legend); the gray regions indicate that a higher digitizer self-trigger threshold (see section 4.1.2) was used.

of [150]) given by:

$$v = \alpha E_g / N, \quad (3.4)$$

with E_g the electric field strength in the gas, N the number density, and with $\alpha = 0.54 \text{ mm}/(\mu\text{s Td})$, where Td (Townesend) is a unit of reduced electric field E/N equal to 10^{-17} V/cm^2 . The scintillation yield Y is proportional to:

$$Y \propto \left(\frac{E_g}{p} - b \right) p d_g, \quad (3.5)$$

with p the gas pressure and d_g the gas gap length. While the proportionality constant a is under some debate (since it can only be measured by detectors with known absolute light detection efficiencies), most studies agree $b \approx 0.8 \text{ kV bar}^{-1} \text{ cm}^{-1}$ [151]. This threshold for secondary scintillation to take place is easily exceeded under normal conditions in XENON1T ($E_g > 10 \text{ kV/cm}$, $p = 1.9 \text{ bar}$, see section 5.1.3).

As a first approximation, we take the meshes as parallel plates separated by d_g of GXe on top of d_l of LXe. The electric field E_g strength in the gas is then constant:

$$E_g = \frac{V}{d_g + d_l / \epsilon}, \quad (3.6)$$

with V the anode-gate voltage difference and $\epsilon = 1.96$ the permittivity of LXe. With XENON1T's 4 kV anode voltage, and the design configuration ($d_l = d_g = 2.5 \text{ mm}$), this amounts to a 11 kV/cm field.

Since the scintillation density depends on E_g (see eq. 3.5), a more refined model accounts for the increase in electric field as the distance r between the electron and anode wire becomes sufficiently small. We will model the electric field as constant E_0 until some critical distance $r = r_A$ is reached, at which point the field becomes $1/r$ (characteristic of an infinitely long wire):

$$E_g(r) = E_0 \begin{cases} 1 & r_A \leq r \leq d_g \\ r_A/r & r_W \leq r \leq r_A \\ 0 & \text{otherwise} \end{cases}, \quad (3.7)$$

where $r_W = 118 \mu\text{m}$ is the anode wire radius. Integrating eq. 3.7, we can express E_0 in terms of the potential difference between the liquid-gas interface and anode, which we approximate with the value from integrating the large-scale field model (eq. 3.6). Combining this with eq. 3.4, we can find the electron's motion by integration.

To extract the XENON1T single electron waveform shape, we select single electrons as described in section 3.4.1. While the spatial distribution of single electrons is relatively homogeneous (as expected from photoionization on impurities), the mean width of single electrons shows a striking pattern shown in figure 3.12. This is well-described by a parabola centered on the origin plus a small linear increase towards the bottom right.

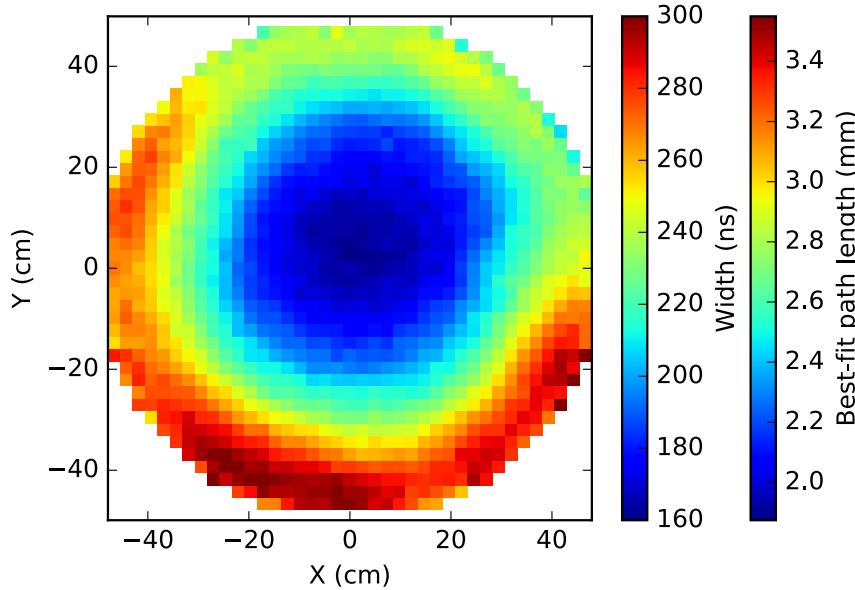


Figure 3.12: Mean width of single electron signals in XENON1T, as a function of observed (x,y) position. Second color scale shows the estimated gas gap size from fits to the single electron waveform shape.

The mean width of single electrons is a sensitive measurement of the gas gap length d_g . Under the constant field (and thus constant drift velocity) model of eq. 3.6, the width of the signal is proportional to

$$\frac{d_g}{v} \propto \frac{d_g}{E_g} = \frac{d_g(d_g + d_l/\epsilon)}{V_A}. \quad (3.8)$$

For $d_l \approx d_g$, this is about four times as sensitive to small changes in d_g than d_l :

$$\frac{\partial \text{width}}{\partial d_l} / \frac{\partial \text{width}}{\partial d_g} = \frac{1}{2\epsilon + d_l/d_g} < \frac{1}{2\epsilon} \approx \frac{1}{4}. \quad (3.9)$$

The most likely explanation of the pattern is therefore a downward sagging of the anode mesh (the parabolic component), combined with a small tilt of the TPC (the linear component). Level meters fixed to the sides of the TPC confirm the detector is indeed tilted as predicted from the linear component [152].

For single electron shape extraction, we divide the TPC into ten regions according to the average width of single electrons – that is, take the darkest blue bins in figure 3.12 as one region, the slightly less dark blue bins as another, etc. The normalized average single electron waveforms, aligned on the center of gravity of the signal waveform (i.e. $t_c = \int t w(t) dt$ with w the normalized waveform), in each region are shown in figure 3.13.

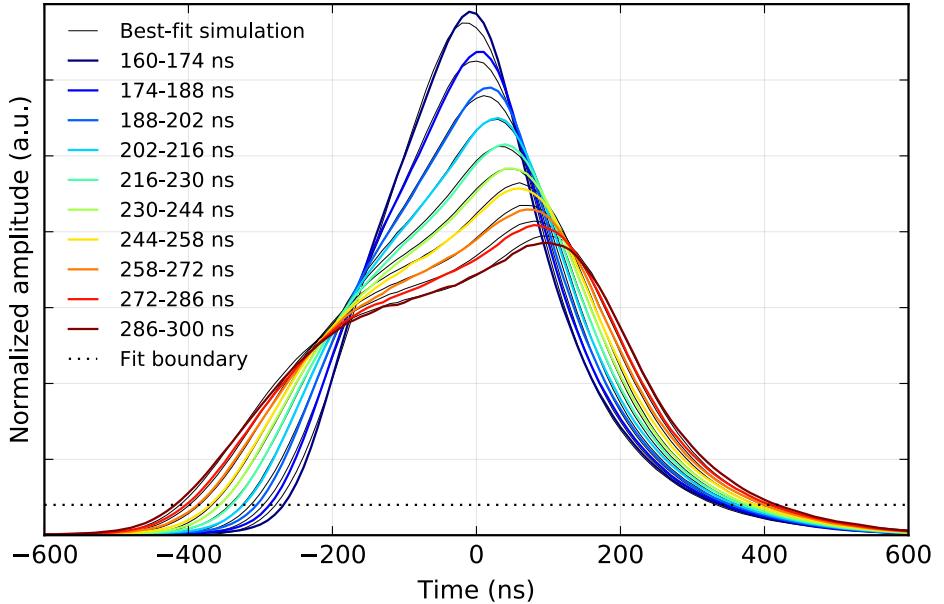


Figure 3.13: Colored lines: average waveforms of single electrons (normalized), aligned on the center of gravity of the waveform, in different lateral regions of the TPC, grouped by the mean single electron width (see figure 3.12). Thin black lines: best-fit average normalized simulated single electron waveforms aligned by the same procedure. The dashed horizontal line indicates the amplitude below which samples are not considered in the fit.

We clearly see differences between the different regions in the TPC, as expected from the width. In the high-width regions (reddish colors) we see the waveform first plateaus (consistent with a uniform field) then rises again for a second peak (consistent with a field increase near the anode).

Figure 3.13 also shows best-fit simulated single electron waveform shapes, using the model discussed above, with r_A , τ_3 and the fraction of singlets excimers f_s as free parameters (constrained to be the same in each region), and d_g free to vary in each region. An extra parameter (the same for all regions) σ is added to describe an additional Gaussian spread of the photon detection times, to account for PMT transition time spread, photon travel time, and imperfections in the model (in particular, the assumption of a delta-function PMT pulse shape, as opposed to the shape derived in section 3.3.3). Samples at the far left and right of the waveform are excluded from the fit, as they are likely dominated by effects not included in the model (PMT afterpulses and accidentally coincident dark rate pulses).

The best-fit values are listed in table 3.2. Note the value of τ_3 is compatible with the literature value. From the fits in different regions, we find d_g increases approximately linearly with the mean width w , so it is simple to convert

| Parameter | Best-fit value | Reference value |
|-----------|------------------------|------------------------|
| d_g | See figure 3.12 | 2.5 mm (design) |
| r_A | (0.372 ± 0.002) mm | |
| τ_3 | (113 ± 10) ns | (100 ± 8) ns [100] |
| f_s | (2.5 ± 1.7) % | |
| σ | (42 ± 4) ns | |

Table 3.2: Best-fit parameters of the single electron shape fit in XENON1T. The parameter meanings are described in the text. The d_g parameter is different for different lateral regions; the result is shown in the second color scale in figure 3.12. A reference value is shown for comparison when available.

the measured mean single electron width into an estimate of d_g , shown as a second color scale in figure 3.12. This would imply the anode is sagging by about 1 – 1.5mm. For LXe TPCs larger than XENON1T, an anode mesh with greater mechanical stiffness should be considered.

3.5 S2s

S2s consist of several single electron signals generated close together. Understanding the drift and diffusion of electron clouds in the TPC will allow us to extend our understanding of single electrons (section 3.4) to full S2s.

3.5.1 Drift

The electron drift velocity v in liquid xenon depends primarily on the drift field strength, but also on temperature and density [127, p. 44-46]. Experiments usually use drift velocities measured in their own detector, with literature values (insofar as they are available) serving only as approximate estimates.

We can estimate v from the drift times of events originating from the cathode and gate electrodes:

$$v = \frac{z_{\text{cathode}} - z_{\text{gate}}}{t_{\text{cathode}} - t_{\text{gate}}}, \quad (3.10)$$

with z the depth below the surface (or any coplanar reference point) and t the drift time. Since we are interested in the average motion of electron cloud, we must measure t until the center of the S2. The numerator of equation 3.10 is known up to construction tolerances of the TPC, and uncertainties in estimating its contraction due to cooling; the denominator has to be inferred from data. To locate an electrode, we can look for an excess of events localized in t , or a change in the mean S2/S1 ratio indicating a change in recombination probability – a higher field suppresses recombination of electron/ion pairs more effectively, strengthening the S2 at the cost of S1 (see section 2.3.2).

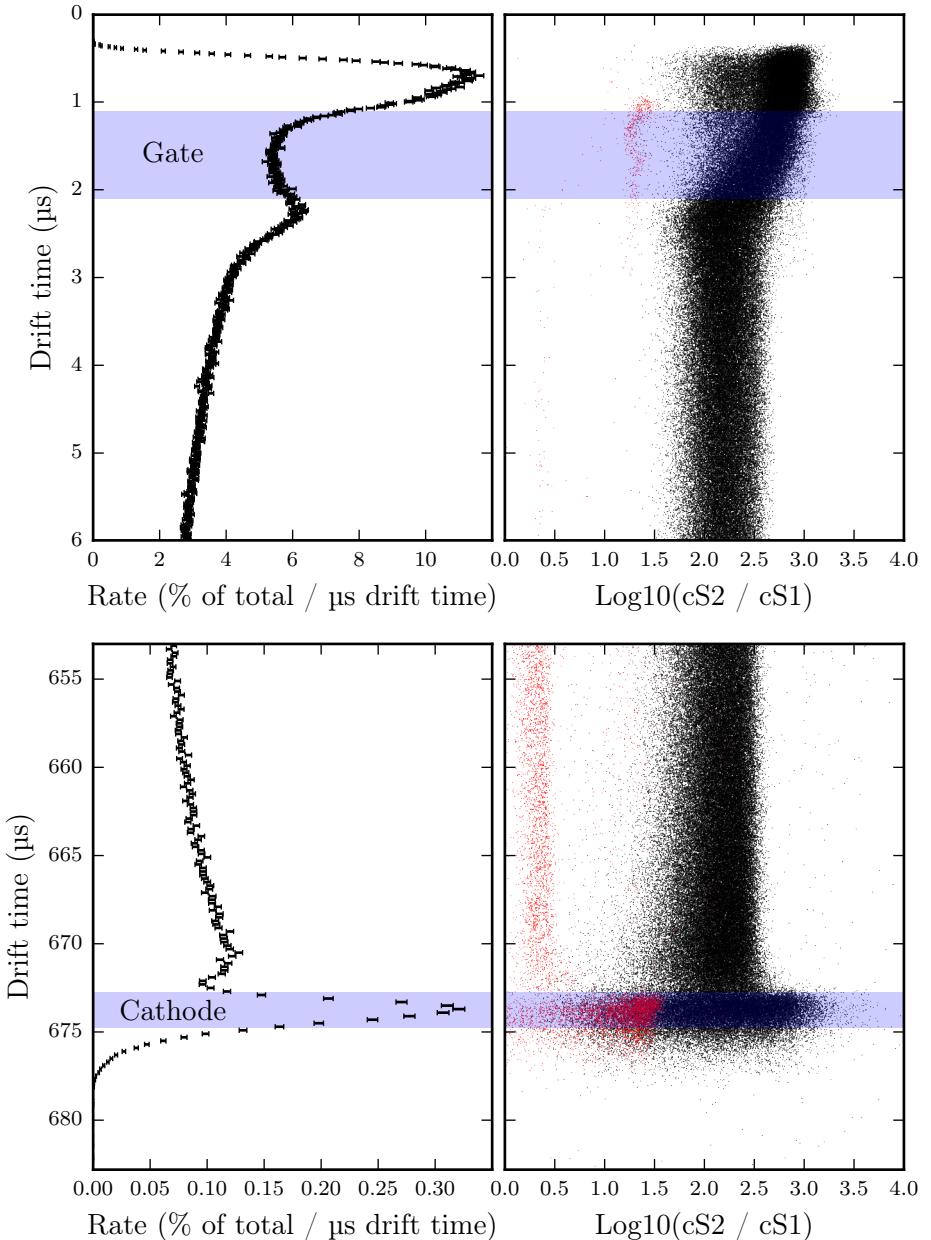


Figure 3.14: Event rates (left column) and S2/S1 ratios (right column) as a function of drift time, for events in the top (top row) and bottom (bottom row) of the TPC in XENON1T’s first science run (34.2 days of data). Estimated drift times corresponding to the locations of the gate and cathode are indicated in blue. Events with $cS1 > 2 \times 10^4$ PE are coloured red. To aid visualization, only one out of ten events is drawn in the top right scatter plot.

Figure 3.14 shows event rates and S2/S1 ratios of events as a function of drift time in XENON1T’s top and bottom. Low-energy S2s ($cS2 < 10^4$ pe), which require more data quality cuts to obtain a clean sample, are excluded from this plot, as are events reconstructed near the TPC’s edge ($r > 37$ cm, after correcting for the field distortion effects discussed in section 2.4.2). Events with too large signals to be ordinary electronic recoils are shaded red; these also show a lower S2/S1 ratio, indicative of alpha decays. The extremely low S2/S1 ratios seen in the bottom right plot are likely from alpha decays of ^{210}Pb in the PTFE walls, where charge collection is imperfect.

The cathode is clearly identified by the local maximum in the event rate at $t_{\text{cathode}} = (674 \pm 1)$ μs , and a simultaneous broadening of the S2/S1 distribution due to the wires’ fringe fields. The event rate gradually increases with t until the cathode is reached, since less LXe shields from external gamma rays lower in the TPC. The cathode appears as a significant overdensity, indicating the wires themselves are also a source of events, including a population of alpha decays from impurities in the wires or adsorbed radon daughters (see section 2.3.3.3 and [153]). Few events are seen with $t > t_{\text{cathode}}$, since no S2 electrons can be collected from below the cathode. That we do see a few such events might indicate small drift field inhomogeneities.

Locating the gate mesh is more challenging. The gate contains about 2-3 times as much metal as the cathode (its optical transparency is 92.7% compared to the cathode’s 97.2%, while it has a similar wire thickness), but the event rate near the top of the TPC is about 50x larger than that near the bottom (since the top of XENON1T lacks LXe self-shielding, as explained in section 2.4.1).

We can still, however, see the increase in S2/S1 ratio between 2.5 μs and 1 μs , as we move from the 124 V/cm drift field to the much stronger (see section 3.4.2) extraction field above the gate. The gate should be located where the change is fastest: around (1.7 ± 0.5) μs . This is also where we see a population of alpha decays, as we saw on the cathode. The local maximum in the event rate at slightly higher drift times could be related to the bending of electron paths near the gate. The much smaller local maximum and minimum seen just before the cathode peak might have a similar explanation.

The resulting drift velocity estimates at different drift fields are shown in table 3.3. The uncertainty is about 0.2%, from adding the uncertainties on the mesh positions in quadrature. Figure 3.15 (top) compares the results with those from XENON10 and XENON100. We can clearly see the drift velocity increases with drift field (as expected) approximately according to a power law. The relevant value for XENON1T’s first science run (nominal field 124 V/cm) is (1.440 ± 0.003) mm/ μm .

After this study was performed, it was discovered the XENON1T electric field during the first science run varied spatially by $\sim 15\%$ within its the fiducial volume, defined by a $r < 37$ cm radial cut (as used here) and a longitudinal cut corresponding to the drift time range of $[-6, -640]$ μs (much wider than

| Drift field (V/cm) | v (mm / μ s) | D (cm 2 /s) |
|--------------------|--------------------|------------------|
| 155 | 1.482 ± 0.003 | 19.2 ± 0.7 |
| 134 | 1.456 ± 0.003 | 21.9 ± 0.8 |
| 124 | 1.440 ± 0.003 | 22.8 ± 0.8 |
| 93 | 1.371 ± 0.003 | 26.8 ± 0.9 |

Table 3.3: Drift velocity v and diffusion constant D in XENON1T, at different nominal drift fields. Detector conditions are equal or close to the nominal conditions of the first XENON1T science run, see chapter 5.

used here). The variation will be larger at lower nominal drift fields, though this has not been quantified for all fields used in the study. The drift field strength listed in this section is always the nominal field, inferred from the electrode voltages and the xenon permittivity. Quantifying the effects of the field distortion on the measured drift velocity and diffusion constant is beyond the scope of this work.

3.5.2 Delayed extraction

The different dielectric constants of liquid and gaseous noble elements cause drifting electrons to encounter a potential well and barrier at the liquid-gas interface [127, p.60]. In xenon (but not argon) this well is deep enough to strongly suppress emission of thermalized electrons, but not the hot electrons of a drifting electron cloud [127, p.62]. When they arrive at the interface, electrons are emitted “without delay” or are thermalized within 1 ps; in the latter case, they will stay trapped below the liquid-gas interface, and can eventually be extracted after moving to parts of the liquid/gas interface with stronger fields, as happens in a tilted TPC [158].

Mock et al. [100] claim electrons exiting the liquid/gas interface are delayed by an exponentially distributed electron trapping time, with a mean of approximately 100 ns for the XENON10 detector. Though they also cite [127] this seems in tension with the statements referenced above. Even if [100] is correct though, an $O(100\text{ ns})$ exponential smearing would be a minor effect on the several μs S2 widths observed in XENON1T; since independent contributions to widths add approximately in quadrature, delayed extraction will be ignored here.

3.5.3 Diffusion

Since XENON1T is large, the shape of an S2-forming electron cloud is significantly affected by diffusion during its drift. There is not much literature on diffusion in liquid xenon, and with the exception of a measurement from XENON10 [157], most of it concerns transverse diffusion (orthogonal to the drift field). We will only consider longitudinal diffusion here, as this affects

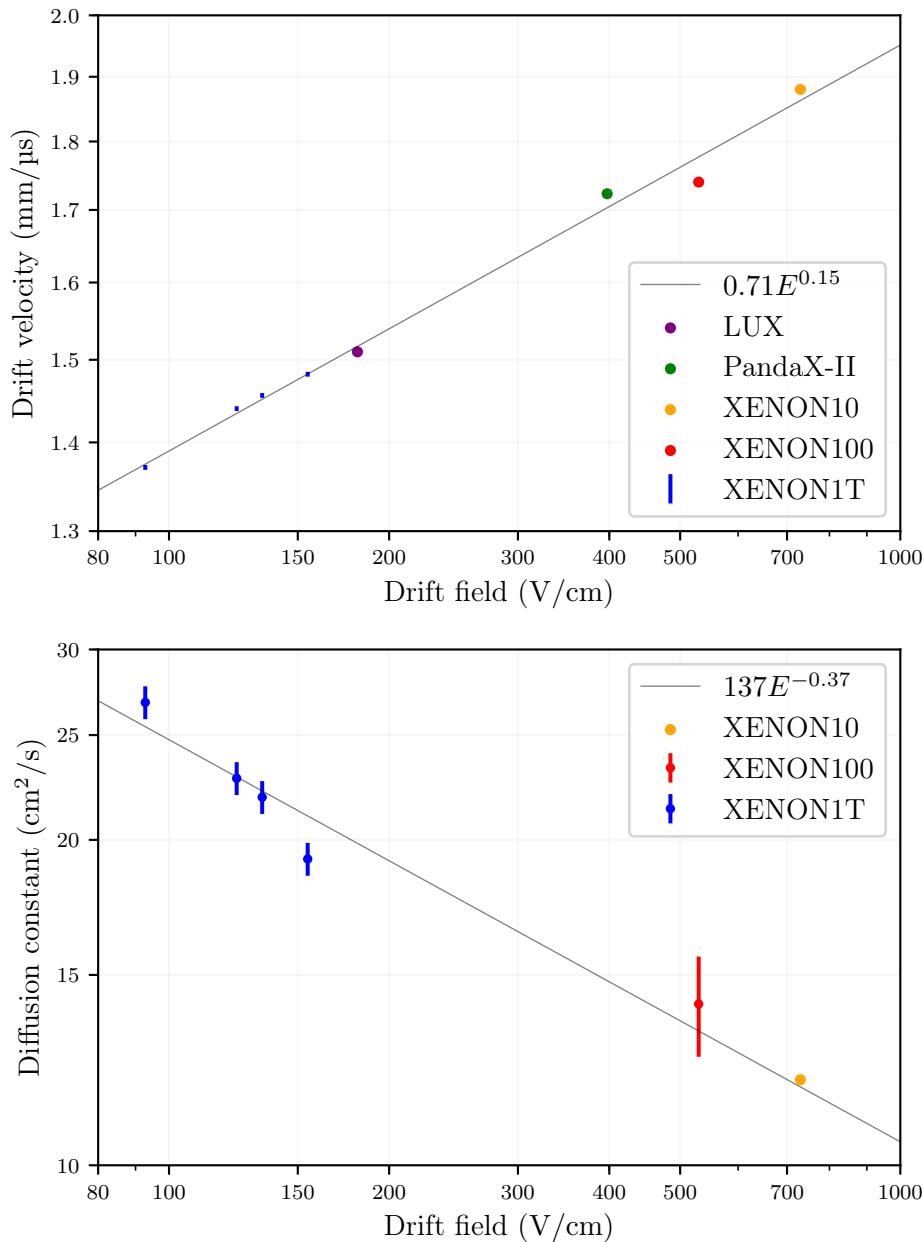


Figure 3.15: Drift velocity and diffusion constant at different fields, in different experiments. Note the log-log scale on both figures. No uncertainties are known for several of the datapoints. A power-law fit to all data shown (assuming the datapoints all have the same uncertainty) is shown in thin gray. ZEPLIN-III achieved a drift field of 3.9 kV/cm (well beyond the scale of this plot) and measured a drift velocity of 2.6 mm/μs [154], which agrees with the power-law fit shown here (within two significant figures). PandaX-II data is from [155], LUX data from [44], XENON100 data from [156] and XENON10 data from [157].

the shape of S2s.

An initially pointlike source of particles diffusing in a fluid (undergoing standard Brownian motion) broadens to a normal distribution with standard deviation Δt_σ proportional to \sqrt{t} . For LXe, the dependence is usually written as:

$$\Delta t_\sigma = \frac{\sqrt{2Dt}}{v}, \quad (3.11)$$

with t the (mean) drift time of the electron cloud, v the (mean) drift velocity, and D the *diffusion constant*, which varies with drift field and in principle also the temperature and density of the medium.

To measure D for XENON1T, we will examine the evolution of S2 widths with depth. The width of very shallow S2s is not zero (as eq. 3.11 would give) due to the scintillation of the single electrons that make up the S2s, and to a lesser extent by the initial width of the electron cloud (except for high-energy ERs, where this is the dominant effect, see figure 2.5).

For sufficiently long drift times, diffusion ensures the S2 signal is approximately Gaussian regardless of the single electron shape, and scintillation of single electrons merely broadens it. The observed S2 width is then:

$$\Delta t_\sigma = \sqrt{\frac{2Dt}{v^2} + \sigma_0^2}, \quad (3.12)$$

where σ_0 is related to the single electron width. If the single electron waveform shape were Gaussian, equation 3.12 would be valid for all drift times, and σ_0 would be the (standard-deviation) width of single electrons.

In XENON1T we do not measure the S2 width by the standard deviation Δt_σ , but by the interquartile range $\Delta t_{50\%}$ (or ‘range of 50% area’) of the signal, that is, the time between the moments when 25% and 75% of the signal’s charge has arrived. This width measure is more robust than the standard deviation to failures of the signal processing to separate the S2 from its single electron photoionization tail (see section 3.4). For a Gaussian waveform shape, $\Delta t_{50\%} \approx 1.349\sigma$, so eq. 3.12 is adjusted to:

$$\Delta t_{50\%} \approx \sqrt{\frac{3.64Dt}{v^2} + w_0^2} \quad (3.13)$$

with $w_0 \approx 1.82\sigma_0$.

For selecting a range of S2 sizes to fit equation 3.14 to, we have to account for two non-physical effects on the S2 width:

1. *Statistical bias of the width estimator.* Many estimators of a distribution’s spread are biased even if they are asymptotically unbiased, such as the sample standard deviation (especially without Bessel’s correction). As long as the S2 is much broader than a single electron (which diffusion usually ensures) the number of electrons, not the number of photons, determines the stochasticity of the S2 shape. Hence we can expect

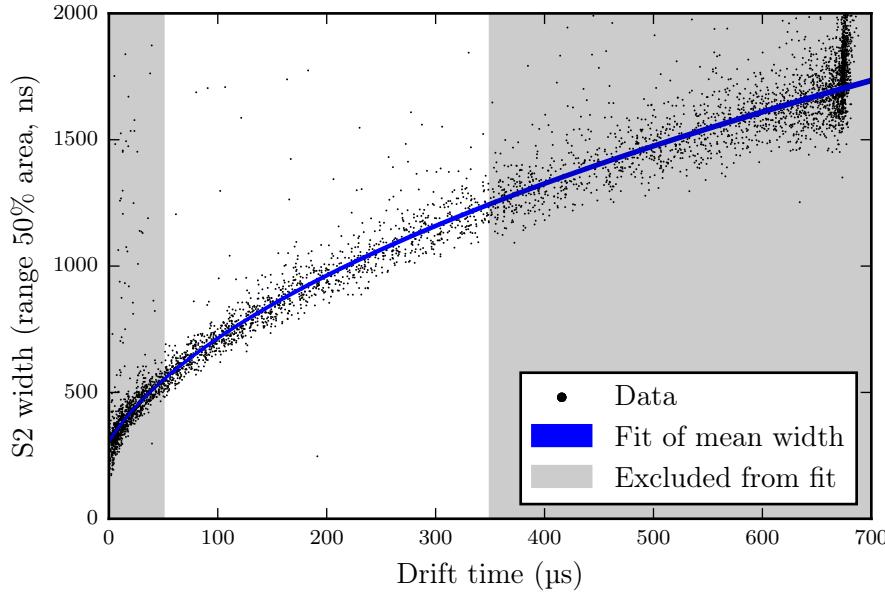


Figure 3.16: Determination of D from S2 widths. Dots: S2 width versus drift time for observed events. Blue line: mean S2 width according to the fitted S2 width model (eq. 3.13); the line width is the 1 sigma fitter error. Gray regions are excluded from the fit.

some statistical bias to affect the S2 widths for $S2 \ll 1 \times 10^4$ PE (500 electrons).

2. *ADC saturation.* From $S2 \gg 1 \times 10^5$ PE, some PMTs in the top array produce signals outside the dynamic range of the digitizer. This removes some charge from the center of the S2 (where the amplitude is highest), causing a flattened shape which has an artificially increased width. Saturation of other parts of the electronics chain (amplifiers, PMT bases) could also distort the waveform shape of high-energy S2s.

Here we use S2 signals with an observed charge in $[1, 2] \times 10^4$ PE. Rather than quantifying the aforementioned biases in greater detail (which would bring in significantly more assumptions), we include changes in the result for D due to making different selections ($[0, 1] \times 10^4$ PE or $[2, 3] \times 10^4$ PE) in the systematic error on the result.

Note the selection is on the observed S2 size, so the energy of the selected events will increase with drift time. At the surface, 2×10^4 PE corresponds to the observed size of the 32.2 keV ^{83m}Kr peak in XENON1T; near the cathode, only 22% of electrons survive the drift (with the 452 μs mean electron lifetime during XENON1T's first science run), so a 2×10^4 PE S2 indicates a ~ 150 keV interaction. From figure 2.5 we see this still has a track length about three

times smaller than the electron cloud spread from diffusion, but it might just be enough to bias the measurement of D . Hence, we limit the maximum drift time included in the fit to 350 μ s. Events with $t < 50 \mu$ s are also excluded from the fit, since they are more contaminated by gas scintillation events, and they may not have a Gaussian shape, since diffusion is no longer the dominant effect determining their shape. The effect of small (50 μ s) variations of the drift time selections are included in the systematic error estimate on D .

Figure 3.16 shows how 3.13 compares to data for 124 V/cm. The best-fit model is obtained by maximizing a standard unbinned maximum likelihood:

$$L = \sum_i^{\text{events}} \frac{1}{\sigma\sqrt{2\pi}} e^{-\left(W_i - \Delta t_{50\%}(t_i)\right)^2/(2\sigma^2)}. \quad (3.14)$$

Here W_i and t_i are the observed S2 width and drift time of each event i , and $\Delta t_{50\%}$ is eq. 3.13 (with w_0 and D as free parameters). The extra free parameter σ controls the fluctuation of S2 widths, which is assumed to be constant for the range of drift times included in the fit. To prevent outliers from e.g. double scatters with unresolved S2s from unduly influencing the fit, the $-\ln L$ that each event contributes is capped to a maximum of 15.

The results of the fit for several drift fields are shown in table 3.3 and figure 3.15, together with results from XENON10 and XENON100. Clearly D decreases with drift field, with approximately a power-law dependence. The uncertainty on D includes the fitter's covariance estimate of 0.3 cm²/s (slightly different for each fit), mostly due to anti-correlation between w_0 and D (which both broaden the S2s if they increase), and systematic effects from the event selections discussed above.

Since S2s are much broader than single electrons in most of XENON1T, a few-electron S2 appears as a series of separate peaks rather than a single large signal. This has important implications for the trigger efficiency, discussed in section 4.2.5. It also forces the data processor to be very conservative in separating nearby signals, which limits our ability to resolve low-energy double scatters. On the other hand, the large variation in S2 widths with drift time enables us to estimate the depth of an interaction from the S2 alone, which we will exploit in the S2-only analysis of chapter 7. For events with an S1, consistency between S2 width and drift time is a powerful selection criterion against accidentally coincident S1s and S2s.

3.6 S1s

The shape of S1 pulses is determined by several factors:

- Decay of the LXe excimers (see section 2.3) generating the S1 signal. The relative amount of fast (3.1 ± 0.7) ns and slow (24 ± 1) ns excimers produced is different for different sources [100], which means the S1 pulse

shape could be used to discriminate electronic and nuclear recoils. In practice, this difference is very challenging to even detect for the small S1s of dark matter candidates. Recent measurements have also shown the excimer decay distribution has little if any energy dependence in the sub-50 keV region [101].

- For very low fields, late-time recombination of excitons and ions. At the $\sim 1.4 \mu\text{m}/\text{ns}$ drift velocity of XENON1T (see table 3.3), the electron and ion clouds are quickly separated at the interaction site (at least for dense tracks, such as the $\mathcal{O}(100 \text{ nm})$ tracks of nuclear recoils), and recombination does not meaningfully affect the S1 time distribution.
- Photon travel time spread. The time for scintillation light to cross the $L \approx 1 \text{ m}$ XENON1T TPC is about $t = Ln/c \approx 5.6 \text{ ns}$, with $n = 1.69$ the refractive index [159] of LXe. Since many light paths lead from the interaction site to a PMT, including via reflections on the TPC wall, the photon travel time can be considerably spread out. The full time distribution can be derived by ray-tracing simulations (see e.g. [160]), but depends sensitively on uncertain parameters such as the UV-reflectivity of the TPC wall material PTFE.
- Transit time spread in the PMT: the path of electrons through the PMT depends on where the photocathode is hit, inducing a 9 ns spread in the time between the arrival of a photon and creation of to the PMT datasheet [149].
- The pulse shape of single photons. This is discussed in section 3.3.3, and can include effects from the analog electronics carrying the PMT signal.

An ab-initio model including all influences on the S1 waveform shape would clearly be quite complex, and have several mostly degenerate parameters. In this section we instead describe an effective model which suffices to simulate realistic S1 waveforms. An in-depth investigation into the S1 signal shape is better done in dedicated small detectors, where optical effects can be neglected, and the S1 pulse shape can be measured with higher-resolution digitizers.

We will only model the pulse shape of electronic recoil S1s. Nuclear recoil S1s are narrower if they are at all detectably different [161, 162, 101], and thus easier to detect and separate from S2s. Detection efficiencies for dark matter estimated with the S1 pulse shape derived below will be conservative.

To measure the S1 shape of XENON1T, we select about 10^4 S1s from $^{83\text{m}}\text{Kr}$ calibration data events (see section 2.3.3.3). This requires 10 hours of calibration data; while much more is available, this should suffice for our purposes. To get a clean sample of $^{83\text{m}}\text{Kr}$ events, we require:

- Interaction depth between 5 – 95 cm, to avoid backgrounds near the surface and cathode.

- Time between S1s of 600 to 2000 ns. The lower bound ensures the two ^{83m}Kr S1s are far enough apart to essentially never be confused by the data processor; almost no events are observed beyond the upper bound due to the ~ 150 ns lifetime of the intermediate ^{83m}Kr state.
- S1 areas roughly consistent with the values expected (and observed) for the ^{83m}Kr source. This mainly removes events in which the smaller ‘S1’ in the event is actually a misidentified single electron rather than the signal of the 9.4 keV decay. In such events the primary S1 consists of both 32.2 keV and 9.4 keV ^{83m}Kr S1s, merged together by the processor.

Figure 3.17 shows the resulting average normalized sum waveforms for the primary (32.2 keV) and secondary (9.4 keV) ^{83m}Kr S1s, aligned on the sample closest to the time where 10% of each S1’s charge has been deposited. Data taken at two different drift fields is shown. Observe:

- There is essentially no difference between the S1 shapes at different voltages. This indicates late-time recombination probably has at most a limited influence on the S1 shape.
- The primary (32.2 keV) and secondary (9.4 keV) S1 look very much alike. This indicates that ^{83m}Kr S1s are large enough that statistical effects (e.g. variation in number of photons per S1 and the detected charge per photoelectron) play only a minor role in determining the average shape. Simulating these may still be relevant to reproduce the variation in S1 shapes.
- At early and late times, the secondary S1’s average waveform flattens off. This indicates it is affected by outliers for which the alignment has failed. Inspection shows several examples where the secondary S1 is grouped together with PMT afterpulses from the primary S1.
- After 300 ns, even the primary S1s’s average waveform flattens off. The likely cause is again PMT afterpulses, which contribute to the (far) tail of S1s but are not included in the waveform simulation routine used here. PMT afterpulses occur $8.6 \pm 2.2\%$ of the time a photon is detected in a XENON1T PMT (spread out over several μs) [128]. Excluding times beyond 300 ns from the fit below hardly changes the results.

Below we focus on the primary S1 waveform for the 81 V/cm dataset, since the secondary S1 is likely contaminated with afterpulses from the primary S1.

To obtain an effective model for the S1 shape, we assume that the photon detection times are distributed according to an exponentially modified Gaussian, then simulate waveforms using the single-PE pulse shape model derived in section 3.3.3. The scale and shape of the distribution are chosen to minimize the Kolmogorov-Smirnov distance between the simulated and observed average waveform, as in section 3.3.3. The best-fit model is shown in

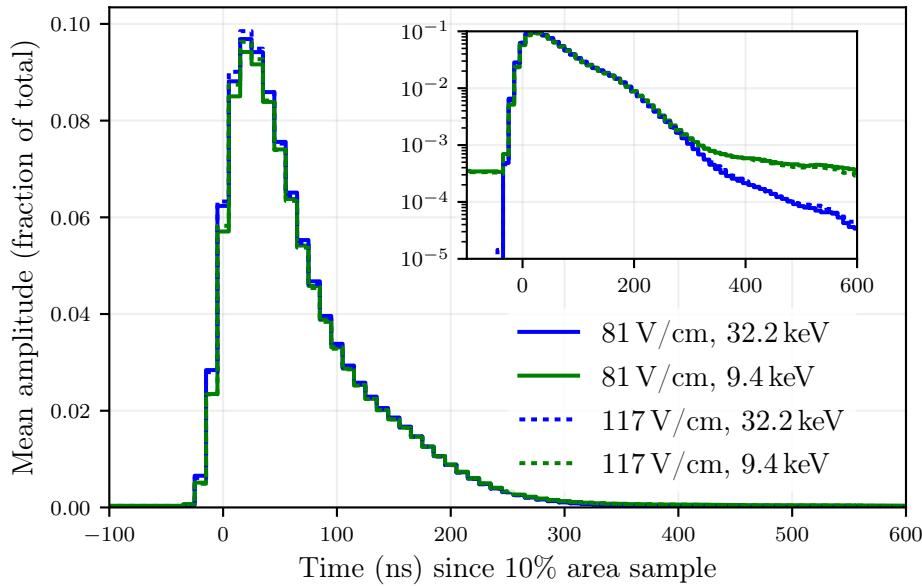


Figure 3.17: Average S1 waveform for the primary (blue) and secondary (green) ^{83m}Kr S1s, observed at 117 V/cm (dashed) and 81 V/cm drift field (solid). Inset: same, with a logarithmic y-axis.

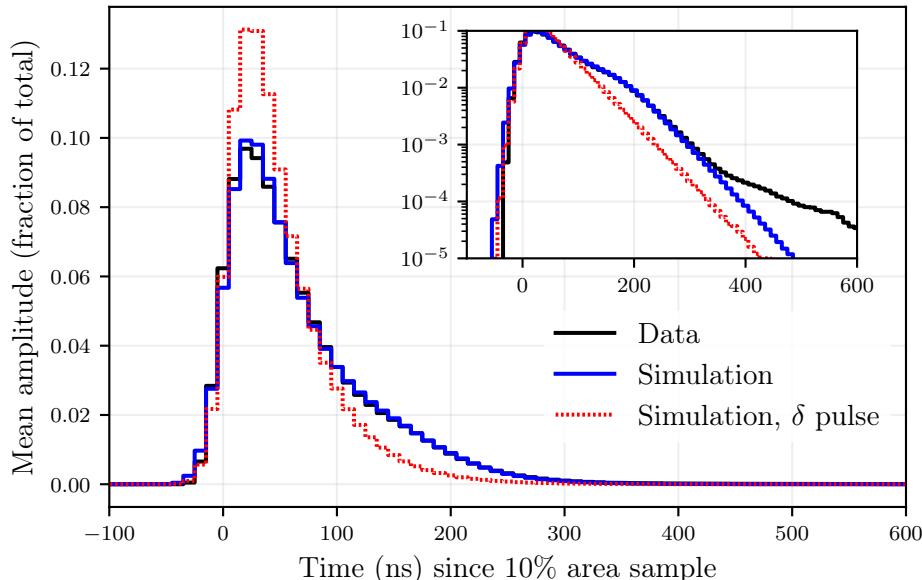


Figure 3.18: Average S1 waveform for 32.2 keV ^{83m}Kr S1s observed under a 81 V/cm drift field (black), and best-fit model (blue). Red dashed line shows the same model, but with the single-PE pulse (figure 3.8, blue) replaced with a delta function. Inset: same, with a logarithmic y-axis.

figure 3.18, achieving a KS-distance of 5.6×10^{-3} . For comparison, observing this value in a $N = 4.4 \times 10^4$ two-sample KS-test would just yield a 50%-confidence rejection of the null hypothesis that both samples come from the same distribution. As before, this number should not be directly compared to the number of S1s used, as a single waveform supplies information at each sample point.

The best-fit parameters are a 43.0 ns decay time and a 14 ns Gaussian spread. The decay time is significantly larger than the LXe excimer decay times (see above), indicating photon travel time (or other effects causing a similar broadening of the S1) might play a significant role. Replacing the single-PE pulse current with a delta function (figure 3.18, red dashed), we see the effect of the late-time component in the single-PE pulse current (see figure 3.8) is also substantial.

To obtain even stronger evidence that we can successfully model the S1 waveform shape, we should compare simulation and data for an observable other than the average waveform (without changing the fit). Moreover, the mean waveform in figure 3.18 is just an average property; we can check if we can reproduce variations in the data as well. To this end, figure 3.19 compares the widths of the observed 32.2 keV $^{83\text{m}}\text{Kr}$ S1s with those of simulated S1s, for different choices of width measure.

The mean widths of real and simulated S1s agree very well, to within $\sim 1\%$ for the 50% area width (the width measure used in classification). The one striking difference between data and model is in the width of the central 90% area: real S1s are slightly wider than simulated, but more importantly, show a larger variation in widths. This could be related to the poor matching at very late times seen in figure 3.18, for which we mentioned PMT afterpulses as the likely cause.

The small remaining discrepancies are very unlikely to have an impact on the dark matter search, since this relies on much smaller S1s, in the range 3-70 PE. Even at 70 detected photons (according to the best-fit waveform model) the spread of the 50% area width due to statistical effects is about twice as large as for the ~ 250 PE $^{83\text{m}}\text{Kr}$ S1s studied here.

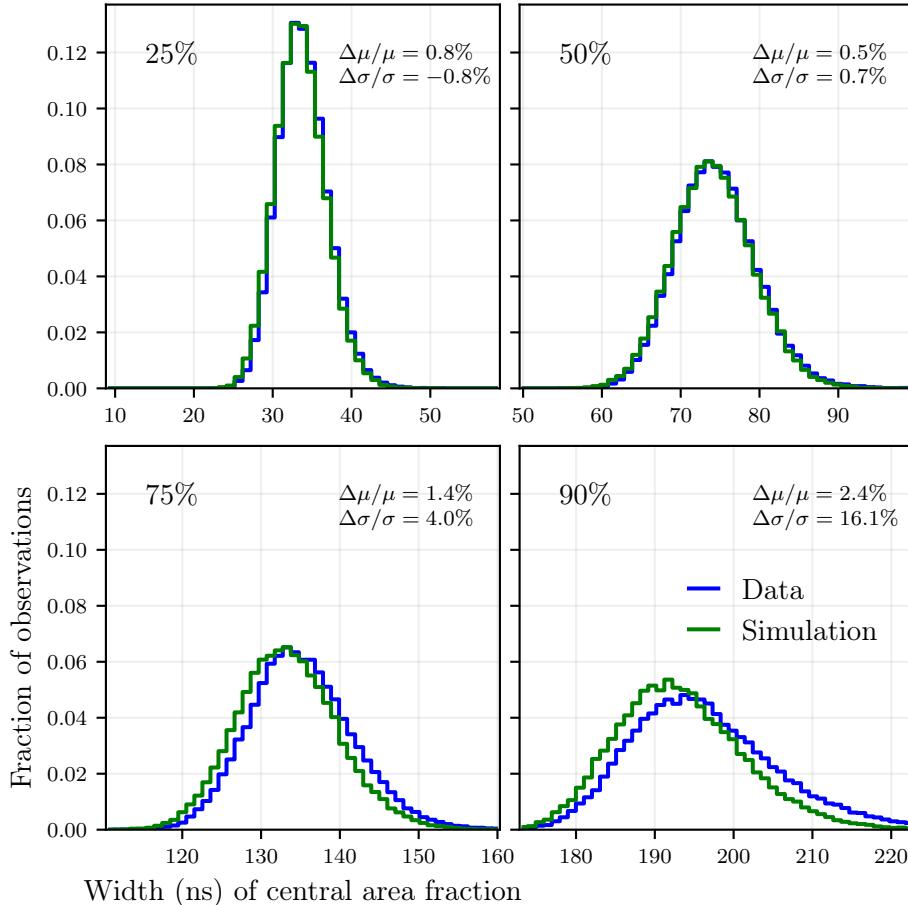


Figure 3.19: Comparison of 32.2 keV $^{83\text{m}}\text{Kr}$ S1s widths observed in data (blue) and in simulated waveforms (green) according to the model shown in figure 3.18. Percentages in the top left corner indicate the central area fraction whose width is measured (with e.g. 50% the usual interquartile range). $\Delta\mu/\mu$ and $\Delta\sigma/\sigma$ show the relative differences in means and standard deviations, where the latter is estimated as half the difference between the 15.8th and 84.1th percentile for increased robustness to outliers. Positive values indicate a larger value is observed in data than in simulation.

Chapter 4

The XENON1T event builder

This chapter details the design and performance of the XENON1T trigger and event builder. In data processing for experimental physics, the *trigger* is the data acquisition subsystem that decides which parts of the datastream are saved or discarded. Since the initial data stream volume is large and continuous, trigger decisions are generally taken before storing the data permanently, hence these decisions are final. The *event builder*'s role is to correctly segment, then encode, the data stream into events that can be processed independently. Delayed extraction tails of single electrons (see section 3.4), double scatters, and decay sequences such as $^{214}\text{BiPo}$ and $^{83\text{m}}\text{Kr}$ (see section 2.3.3.3) make decisions in the trigger and eventbuilder nontrivial. The trigger and event builder roles overlap, and in XENON1T, the trigger is implemented as a subsystem of the event builder.

Trigger decisions in the XENON1T DAQ are taken completely in software, which enabled quick development and the addition of various monitoring features. The data from before the trigger decision was taken ('untriggered' data) remains available for minutes to hours (depending on the data rate) after the data was originally acquired. This data can be saved permanently for diagnostic purposes, or on externally provided triggers, such as alerts for galactic supernovae.

This chapter starts with an overview of the XENON1T DAQ in section 4.1. The trigger design and its performance is discussed in section 4.2. Finally, section 4.3 discusses the broader event builder software implementation, highlighting features such as the interprocess communication pipeline supporting distributed event building, and the supernova neutrino detection system.

4.1 The XENON1T DAQ

4.1.1 System overview

The XENON1T DAQ consists of several subsystems, illustrated in figure 4.1:

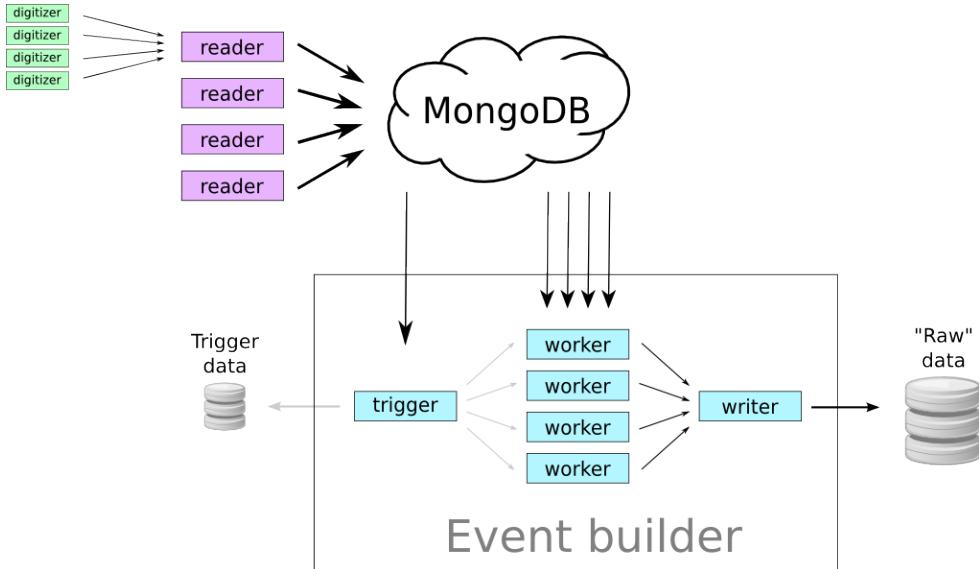


Figure 4.1: Illustration of the XENON1T DAQ systems and data flow. Arrow thickness is a qualitative data transfer volume indicator.

- The *digitizers* send pulses ($\sim 1\text{ }\mu\text{s}$ stretches of single-PMT data, see section 4.1.2 below) to the eight *reader PCs*. Kodiaq [163], the software running on the readers, controls the readout process, computes basic quantities for each pulse (such as the baseline or the integral) and stores these, along with the pulse’s actual data, in MongoDB.
- *MongoDB* is a multi-host open-source database system [164]. Three separate MongoDB instances running on separate machines store pulse data coming from the readers: each reader is assigned to one of the databases. The database ensures the consistency of simultaneous write (from the reader) and read (from the event builder) operations.
- The *event builder* is software consisting of several processes on different machines. It reads the untriggered data from MongoDB, looks (usually) for S2 candidates, then saves pulses in some window around them as events to a “raw data” file (see section 4.3.2). This is a three-stage process:
 - *Input and trigger decision.* The event builder first reads out summary information of each pulse from MongoDB: the start time, channel, and (if available) the charge corresponding to the pulse’s integrated amplitude. Based on this, we decide which events to build. This stage also produces a stream of monitoring data (see section 4.2.3).
 - *Workers: Data reading and compression.* On several cores on

several machines, processes are waiting to receive event ranges to build. When a worker receives a range, it pulls the full pulse data from MongoDB, then encodes and compresses it into our raw data format.

- *Writing*: A single process receives the compressed events from the workers. It restores the sort order (earliest event first), then writes the events to disk. During high-rate calibration, this stage also deletes old pulse data from MongoDB to maintain sufficient capacity. For the low-rate dark matter data taking, deletion is managed separately: see section 4.3.3.

Some components integral to the XENON1T DAQ are omitted from figure 1:

- The *high-energy veto* is an FPGA board that monitors the data flow in real time. It can throttle the digitizer’s readout during high-rate calibration conditions, offering a more efficient deadtime control than the digitizers’ automatic BUSY signal (triggered when their internal memory buffer is full).
- The *acquisition monitor* is a single digitizer that stores the analog sum waveform of the bottom-array PMTs, as well as digital signals from the BUSY, high-energy veto, and muon veto detector (see section 2.4.1) trigger. The event builder stores acquisition monitor signals to a separate file, regardless of whether they occur inside events.
- The *RabbitMQ broker* [165] serves as the event builder’s backend for interprocess communication, see section 4.3.1.
- The *runs database* stores metadata about data taking conditions and settings, facilitates state communication between the various DAQ subsystems, and has been various other uses in XENON1T (such as tracking the location of raw and processed data across analysis and processing facilities).
- The *DAQ website* exposes the control interface of the DAQ and contains visualizations from various monitoring systems.

4.1.2 Data structure

The fundamental unit of data in the XENON1T DAQ is the *pulse*, a short (roughly 1 μ s) block of samples from a single digitizer channel, connected to a single PMT. The digitizers report a pulse whenever a configurable amplitude threshold is exceeded in a channel. This is usually called *zero-length encoding* or *pulse-only digitization*. Most DAQs in low-background experiments actually store such pulses only when instructed to do so by a realtime hardware trigger.

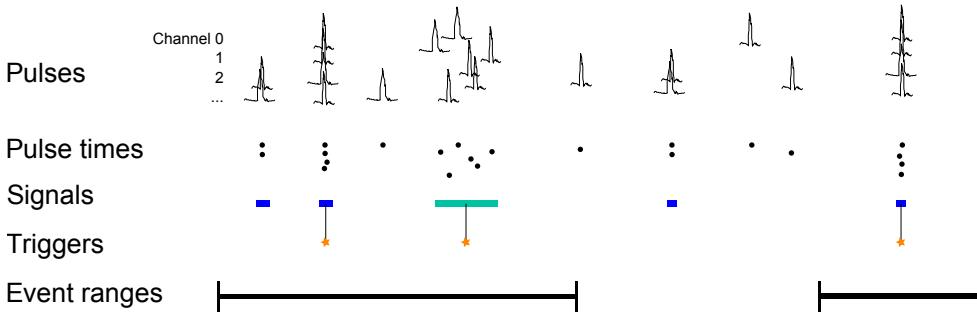


Figure 4.2: Illustration of the XENON1T trigger architecture. The trigger defines successive abstractions on the data, from pulses to signals (which can be labelled as triggers) and event ranges. Colors of the signal ranges indicate the trigger’s guess for the signal type: blue for S1, green for S2.

In contrast, XENON1T’s digitizers output pulses whenever a threshold crossing occurs, a property called *self-triggering* or *triggerless readout*.

In the analysis after the DAQ, the fundamental unit of data is the *event*, a collection of pulses from several channels that will be processed independently. XENON1T’s event processing is stateless between events: processing an event individually (e.g. to visualize its waveform) gives the same result as processing it in the middle of a dataset. Hence a time period allocated to an event must be large enough to include all signals that could be relevant for the interpretation of the main physical interaction(s) than constitute the event, and in particular, to decide on the pairing of each S1 and S2.

4.2 Trigger design

4.2.1 Architecture

The trigger has a ‘pipeline’ architecture: data proceeds through several stages (implemented as ‘plugins’) which each add some information. Initial stages deal with the data on a basic level (e.g. individual per-channel pulse times); later stages deal with larger abstractions (e.g. groups of pulses across channels, or even entire event ranges).

The stages in the trigger decision are shown schematically in figure 4.2:

- The trigger’s input is summary information about each pulse from the event builder – most importantly, the start time of each pulse. The actual raw waveforms of the pulses are not yet requested from the database (until a later stage in the event builder, see 4.3.1).
- The pulse times are grouped into *signals* consisting of closely spaced pulses. A signal ends when the difference between the start time of two successive pulses exceeds a threshold, currently set at $1\text{ }\mu\text{s}$. This is too

large to split close S2s or $^{83\text{m}}\text{Kr}$ S1s, but that can be done during later event processing. What matters here is that we do not split small deep S2s, which can be several μs wide (see section 3.5), into fragments that are too small to trigger individually.

- Next, the trigger computes basic properties of the signals, such as the total number of contributing channels and the standard deviation of the pulse start times. Based on these properties, the signals are classified as S1 candidates (blue in figure 4.2), S2 candidates (green in figure 4.2), or ‘other’. Any signal of 10 pulses or more is considered an S1 or S2 candidate. This should be small enough to include even ~ 22 -hit single-electron S2s, while large enough not to avoid flagging accidentally coincident PMT dark rate (see figure 3.9). XENON1T is sensitive to interactions with smaller S1s too, but these will be triggered by the S2 (and the small S1 will still be found during event processing).
- Based on the signal properties and classification, we flag some signals as *triggers*. Currently, S1 candidates of 50 pulses or more and S2 candidates of 60 pulses or more cause a trigger. These settings are chosen to avoid constant triggers due to pile-up of several single electrons during the delayed extraction tails of long events (see section 3.4), while retaining an acceptably low S2-threshold (see section 4.2.5).
- Finally, we send out an event range around each trigger, grouping nearby triggers together. To be precise, we save 1 ms (slightly more than our maximum drift time, see section 3.5.1) of data on either side of each trigger, and group events together if their ranges would overlap as a result of this.

The efficiency of these settings will be discussed in section 4.2.5.

During high-rate calibration data taking, the trigger can be configured to store only a fraction of high-energy events – which take the vast majority of disk space and processing time, but are seldom the goal of the calibration. The trigger signal information in the affected events will still be saved, but not the raw pulse data required for full processing. This functionality is complementary to data reduction by the hardware high-energy veto (see section 4.1.1).

4.2.2 Batch edge handling

The trigger’s interpretation of a single pulse depends on the pulses before and after it – we need to look for nearby pulses to decide if the pulse is part of a signal, or just random noise or PMT dark rate. Similarly, a single trigger cannot be converted to an event time range without looking if additional triggers are nearby, since these may have to be grouped into the same event to avoid shortened events or duplicated data regions.

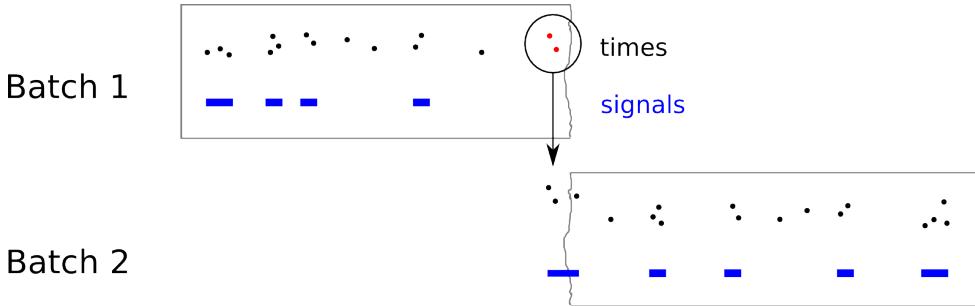


Figure 4.3: Illustration of batch edge handling in the XENON1T trigger for pulse times: we postpone processing of any pulse that could be grouped into a signal with pulses in the next batch until that next batch has arrived. Edge handling for (trigger) signals could be illustrated by an identical figure, with (times, signals) replaced by (signals, events).

This causes a problem. We read out the pulse data in batches from the database, as opposed to waiting for a data acquisition run to finish and reading everything at once. However, we cannot process each batch independently, since information from other batches is needed to interpret the pulses at the batch edges. We must find a way to carry over information from one batch to another, so that the result is *invariant* to how the data is broken up into batches. The solution implemented in XENON1T is (partially) illustrated in figure 4.3:

- *Edge handling for times*: after receiving a batch of pulses, we determine which pulses could belong to signals extending into the next batch. This is easy to check, as any break larger than the $1\mu\text{s}$ signal separation threshold guarantees a new signal. These times are hidden temporarily, and prepended to the times in the next batch when we receive it.
- We build signals from the remaining pulses, classify the signals and decide which of these cause triggers.
- *Edge handling for signals*: we check which signals might belong to events extending partially into the next batch. This is again easy to check, as any time without triggers larger than a window (currently 2 ms) guarantees a new event. We again hide these signals and prepend them to the signals found in the next batch.
- Finally, we decide which event ranges to record based on the remaining signals.
- If a batch of data is the last in the run, the edge handling steps will not hide away any times or signals.

This system ensures each stage in the trigger receives a ‘cleaned’ batch in which all features can be interpreted independently of the data that comes before or after it. Since all pulses are eventually considered, the raw data produced by the trigger does not change if batch breaks are inserted at different points in the datastream.

Parallelizing the trigger decision over several processes would require a considerably more complicated system. In this case, interpreting the data on a batch edge cannot be delayed until the next batch has arrived: that batch has to be processed in parallel, or may already have been processed in a different process or even a different machine. Fortunately, trigger decision parallelization (unlike event builder worker parallelization) is not needed for XENON1T, as discussed in section 4.2.4.

The XENON1T DAQ was eventually configured to use very large query batches of 21.47 s each (a single digitizer clock cycle), so in hindsight an alternative to this edge handling would be to simply accept that fragmented events can occur for $\mathcal{O}(10\text{ ms})$ around batch edges. However, such fragmented events can still cause confusion downstream in the analysis, unless they are carefully tracked and removed.

4.2.3 Live monitoring

Since the trigger is implemented in software, it was natural to add a variety of live monitoring capabilities. The trigger is uniquely positioned: it is the last system in the DAQ to see the entire untriggered data stream, and the first to know which pulses belong to signals and events and which are merely lone pulses.

- *Lone and total pulse rate per PMT.* For each PMT, for each second of data taking, the trigger stores the rate of pulses and the rate of pulses outside signals (lone pulses). This has proven extremely useful during DAQ commissioning, and for detecting transient technical issues with individual PMTs, as illustrated in figure 4.4.
- *Two-PMT coincidence rate.* Correlated noise between channels or light emission from PMTs can produce fake S1s in excess of those predicted from pileup of the PMT dark rate. To study this, the trigger tallies signals consisting of pulses in just two PMTs, and saves the count per PMT pair every minute. Figure 4.5 shows an example of this data.
- *Trigger signals.* Inside events, all the trigger’s signals are kept whether they triggered or not – though at least one signal triggered, else there would be no event. This facilitates data-driven trigger efficiency studies, briefly discussed in section 4.2.5. Trigger signals outside of events can be saved in the trigger monitor data, but this is rarely done due to the large data volume (in comparison to other trigger monitor data) required.

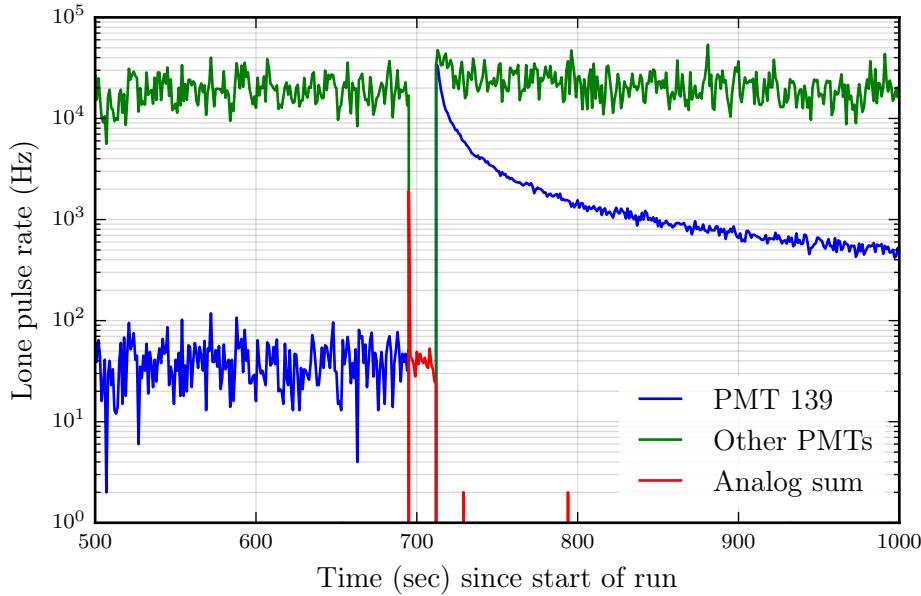


Figure 4.4: Example of trigger monitor data during a run with a PMT problem. At 695 seconds into the run, we stop receiving pulses from the PMTs. Since the analog sum waveform shows (lone) pulses, we know the PMTs remain functional. At 712 seconds, PMT signals return, and the cause becomes clear: a problem in PMT 139 caused emission of large amounts of light, saturating the digitizer-reader communication. Separately digitized BUSY-signals from the digitizer (not shown here) confirm this.

Instead, these signals are tallied in a histogram of (number of pulses, standard deviation of pulse times), and saved every minute during data taking.

4.2.4 Processing speed

The XENON1T trigger is a single-threaded python application, but it must process data live during the most strenuous calibration data runs. The digitizer-reader optical links were designed to transfer up to 300 MB/s data, so if the trigger can handle this, it will never be the bottleneck in the DAQ.

Before XENON1T started data-taking, we defined two test datasets to measure the trigger speed:

- *Modified XENON100 calibration data.* One minute of XENON100’s nuclear recoil ($^{241}\text{AmBe}$) calibration data takes roughly 300 MB uncompressed (depending on the data format used), a mere 5 MB/sec. This is, of course, because the data is stored after it has passed through

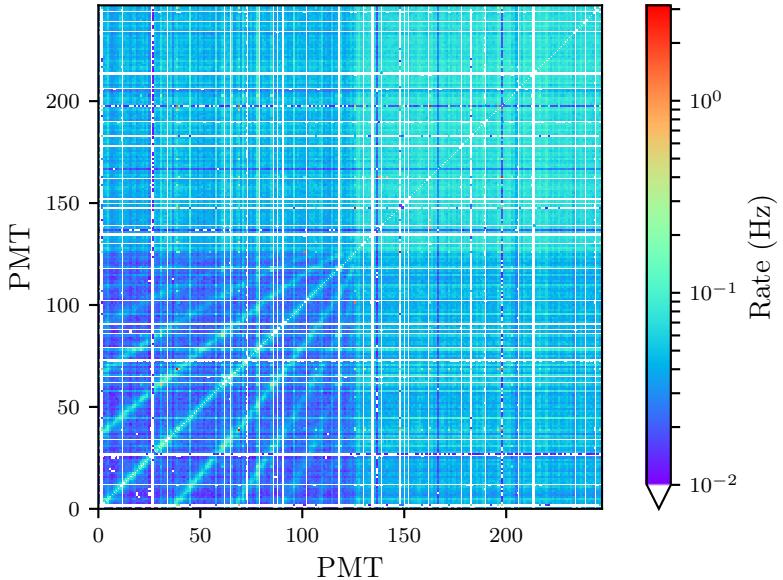


Figure 4.5: Rate of 2-PMT multiplicity signals for each TPC PMT pair during an ordinary one-hour background run. The plot is symmetric across the identity diagonal by definition. Mostly-empty rows and columns correspond to PMTs that have been turned off; remaining pulses in those channels are due to electronic noise. The band structure in the lower left corner correspond to PMTs in the top array (0-126) which neighbour each other (see [126] for the numbering scheme). This feature is not seen in the bottom PMT array (127-247), likely due to the higher signal rates due to its direct contact with the LXe scintillator.

XENON100’s trigger. We convert it to a denser (but obviously corrupted) dataset by shuffling the events closer together (by applying a modulus to all event start times). In the most extreme condition tested, 50 minutes of XENON100 $^{241}\text{AmBe}$ were compressed into 3 seconds; the full data rate for this would be 5 GB/sec.

- *Random pulse times.* Fully synthetic data, consisting of uniformly distributed random pulses with the minimum pulse size (200 bytes: 1 μs of one 16-bit sample per 10 ns). This represents a worst-case scenario for the trigger: its workload depends on the number of pulses (not their length), and it is optimized for groups of well-separated pulses rather than a continuous stream.

The trigger’s processing speed on both datasets are shown in figure 4.6. We see the trigger processing rate is comfortably fast for the XENON100 $^{241}\text{AmBe}$ dataset, but suffers some degradation under worst-case conditions

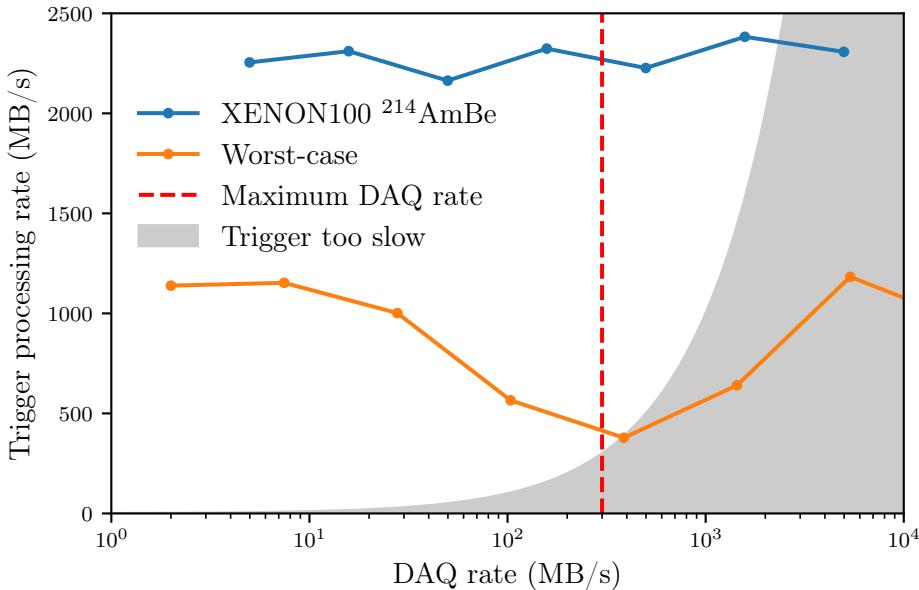


Figure 4.6: Data processing rate of the XENON1T trigger for two data sources (see text), at different data output rates from the DAQ. The shaded region indicates conditions in which the DAQ would produce data faster than the trigger can process. The red line indicates the maximum data rate the XENON1T digitizers are designed to output.

around data rates of 400 MB/s. At this rate, the signal grouping algorithm likely produces a constant rate of signals, which causes extra work for the subsequent stages. The workload decreases at higher rates, when the pulse density is high enough to cause larger and larger signal groups; at the highest rates, the trigger groups all pulses into one signal, and the data processing rate (in MB/sec) is the same as for the lowest rates (where no signals are generated at all).

Clearly the trigger is fast enough to process data at our design rate of 300 MB/s under realistic conditions. Experience with XENON1T data confirms the trigger is comfortable fast – we have never produced data at rates the trigger could not process live.

4.2.5 Trigger efficiency

XENON1T can trigger on both S1s and S2s, but since the S2 signal is much larger, S2s essentially always provide the trigger for low-events. Thus, measuring the S2 trigger efficiency is the most important. We can do this by running the trigger on simulated S2 waveforms, or by selecting events that triggered by their S1 but in which the event processor finds a valid S2, then checking if the

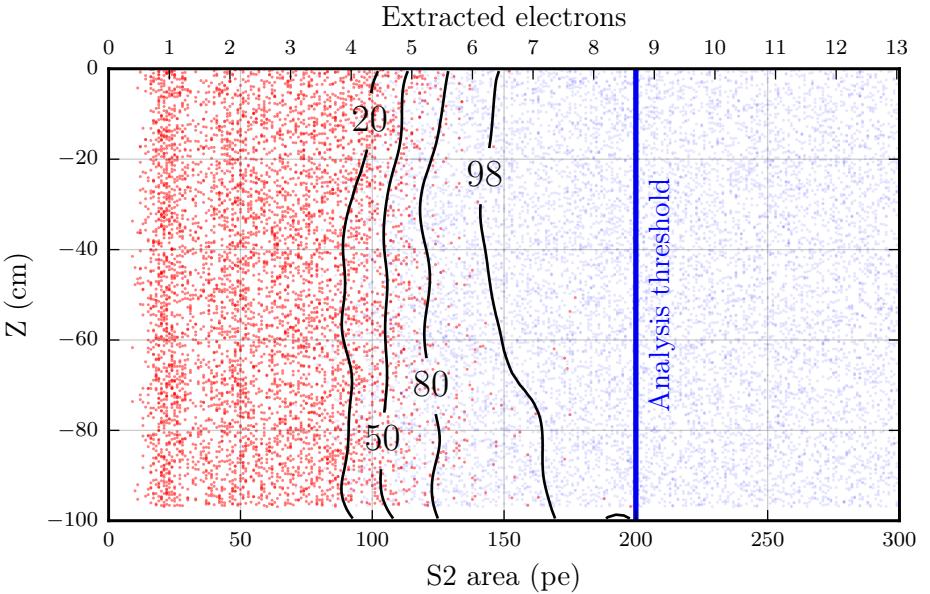


Figure 4.7: Simulated S2 signals in $(S2, z)$, with trigger status indicated by color: light blue events were found, red events were missed. The bottom y-axis lists the observed S2 area in PE; the top y-axis lists the equivalent number of extracted electrons using the single electron gain from section 3.4.1. Contours indicate levels of 20, 50, 80 and 98% trigger efficiency, estimated by ratio of integrated kernel density estimates for found and missed events. The main analysis' S2 threshold of $S2=200$ pe is indicated by the thick blue line.

trigger also found that $S2$ (using the saved trigger signals, see section 4.2.3). The former technique relies on the correctness of the waveform simulator, while the latter relies on correct accounting for biases and efficiencies in the event processing and data selection. The simulation-driven method is discussed here, the data-driven method is discussed in [166] (and yields similar results).

We simulate 10000 S2s uniformly distributed inside the TPC. The S2 signal shape simulation is discussed in chapter 3. Each S2 is independently fed to the trigger algorithm. Figure 4.7 shows the simulated events in $(S2, z)$, with the color indicating whether the S2 caused a trigger or not. Clearly, the trigger efficiency is essentially perfect above the main analysis' S2 threshold at 200 pe.

Figure 4.7 also hints at a small depth dependence of the trigger efficiency: around 150 PE, deep S2s are occasionally missed by the trigger. Examination of the simulated waveforms shows these S2s are split over two signals by the trigger, and neither fragment crosses the 60-pulse threshold. This can be mitigated by increasing the signal separation threshold (see section 4.2.1), at the risk of causing more false events during bursts of signals caused by PMT

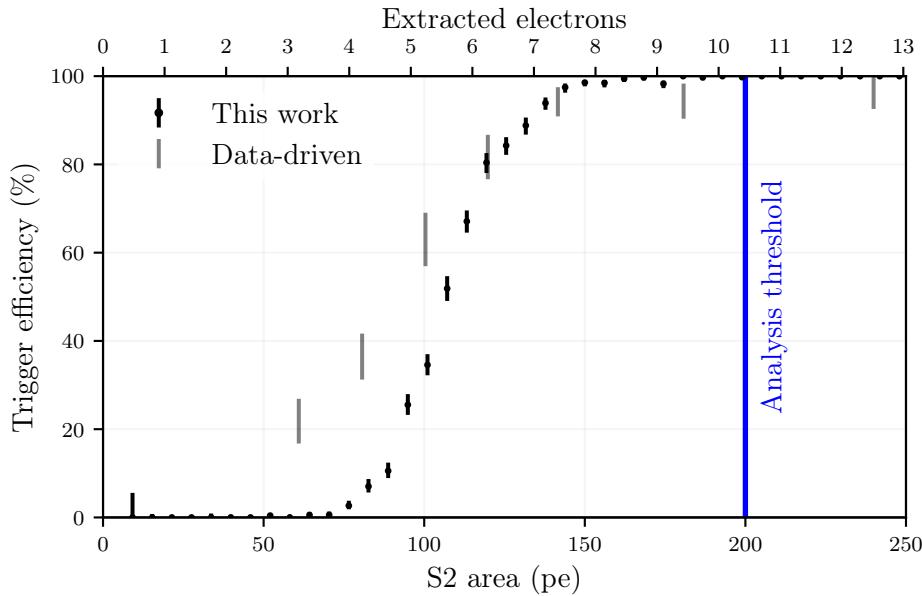


Figure 4.8: Trigger efficiency as a function of S2, derived from simulated events shown in figure 4.7. Error bars are statistical (68%).

malfunctions or electronic noise. Simply lowering the 60-pulse threshold would increase the efficiency for all S2s (at the cost of more false triggers), but would not change the depth dependence.

Neglecting the possible depth dependence, we histogram the simulated events in S2 and obtain the efficiency shown in figure 4.8. Again, it is clear the trigger efficiency has no impact on the dark matter analysis. The data-driven study [166] agrees in this conclusion, though it shows a higher trigger efficiency of S2s substantially below the threshold. This is likely an artifact of the data-driven study, which requires events with S1s greater than ~ 50 PE, who normally have S2s of thousands of PE (see e.g. figure 6.3, which even shows only the S2 area in the bottom array). Thus, many of the ‘events’ with low S2s it finds are likely to be artifacts, whose S2s are actually pile-up of PMT afterpulses and/or single electrons, on which the trigger efficiency is unpredictable. Moreover, no data quality cuts were used to potentially remove these.

The current trigger settings were set well before XENON1T’s first science run started, when detector conditions were substantially different. In particular, there used to be a large rate of single electron emission from the cathode, which would have caused many additional events (and lengthened existing events) had the trigger threshold been lowered further. Since this emission has been removed by lowering the cathode voltage, lowering the trigger threshold could be investigated.

As mentioned, the main dark matter analysis has an S2 area threshold of 200 PE; below this, accidental coincidences and wall leakage events (see chapter 5) are expected to cause a large and uncertain background, and the dark matter detection efficiency is already very limited due to requiring an S1 signal. The S2-only analysis (see chapter 7), however, is designed to handle unknown backgrounds, and by definition requires no S1. Thus, lowering the trigger threshold below the current value is still desirable.

4.3 Event builder

4.3.1 Architecture

The event builder is the software responsible for the entire process of converting pulse data into encoded raw data files (see also section 4.1.1). The main design challenge is combining operations that are intrinsically serial (e.g. ordered writing of events to a file) or very hard to parallelize (triggering a batch of data) with CPU-intensive query and encoding stages which must be parallelized to support high data rates.

The event builder architecture is built around several instances of the XENON1T event processor **pax** [5], as illustrated in figure 4.9. An input process produces events and puts them on a processing queue. Several workers fetch these events, process them, and put them on an output queue. This is the most CPU-intensive stage, which necessitates a multiprocessing system. Finally, an output process fetches from this queue, puts the events back in the original order, then writes them to a file.

For local multiprocessing (not shown in figure 4.9), the queues are maintained by a single master process in shared memory with the **pax** child processes. The master process additionally shares memory to enable state communication, e.g. to indicate the input source is exhausted, the worker stage is complete, or a crash has occurred in some process. While simple to implement, local multiprocessing limits the number of useful workers to the number of cores available on a single machine.

To achieve higher throughput rates, we need a distributed multiprocessing system, and must replace all shared memory operations with network services. Specifically, we use a message broker system called RabbitMQ [165], which provides remotely accessible queues and ‘fanouts’ (queues that deliver a single message to multiple endpoints). Processes in the same chain use these remote queues to pass blocks of events to each other. The blocks are numbered when they are created, to allow reconstruction of the event order at the output stage. Each **pax** process runs inside a host process; one type (**event-builder**) for hosting input and output processes, another (**paxmaker**) for hosting workers. Host processes stay alive after a run is processed; **event-builder** will start processing a new run (or wait until one is available from DAQ) while the **paxmaker**’s correspondingly wait for new requests to start workers. Host

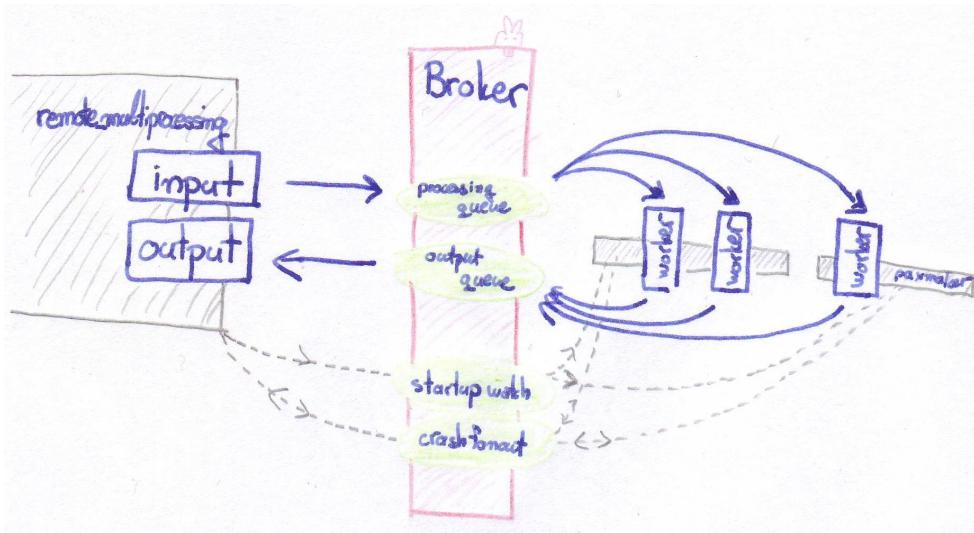


Figure 4.9: Distributed multiprocessing architecture of the XENON1T event builder. Each blue box is a process running `pax`, each gray box indicates a host process. Arrows indicate the flow of events (dashed arrows, other messages); green circles indicate RabbitMQ communication channels.

| Process name | Function | Number | Persistent |
|---------------|--|--------|------------|
| Input | Run trigger algorithm | 1 | No |
| Worker | Query and encode event data | 70 | No |
| Output | Write events to disk | 1 | No |
| event-builder | Host input and output, check for new runs to trigger | 1 | Yes |
| paxmaker | Host workers | 3 | Yes |

Table 4.1: Processes in the XENON1T event builder in remote multiprocessing mode. ‘Number’ shows how many processes of each type are active during processing of a single run; ‘persistent’ indicates if the process survives after processing a run.

| Channel name | Writers | Readers | Messages contain |
|------------------|------------------------------|------------|-----------------------|
| Processing queue | Input | Workers | Event time ranges |
| Output queue | Workers | Output | Encoded events |
| Startup watch | event-builder | paxmaker’s | Start worker requests |
| Crash fanout | event-builder and paxmaker’s | | Crash notifications |

Table 4.2: RabbitMQ communication channels used by the XENON1T event builder in distributed multiprocessing mode. Writers and readers list which processes (see table 4.1) communicate on the channel. The worker processes exceptionally write to the processing queue when passing on the `NO_MORE_EVENTS` message (described in the text).

| XENON1T- aaaaa - bbbbbbbb - cccccccc - ddddddddd .zip | | | |
|---|-------------|------------|-------------|
| Run ID | First event | Last event | Event count |

Table 4.3: Filename convention for XENON1T raw data archives. The top row shows the file name template. Repeated bold characters are replaced by numbers describing the contents of the archive, with meanings as indicated by the second row. All fields except the DAQ run id are zero-padded to 9 digits to ensure lexical filename sorting coincidences with the trigger’s event numbering (which coincides with the temporal order of events in a run).

processes communicate on two fanout channels: one to exchange requests to start **pax** workers when a processing chain starts, and one to share news of crashes. To avoid infinite hangs, a crash in one process of a chain causes the host processes to terminate all **pax** processes in the same chain.

If processing ends regularly, this must also be properly communicated throughout the chain. When the input process determines the source of events has dried up, it sends the message **NO_MORE_EVENTS** to the processing queue. A worker process reading this message will shut down after passing the message back on the processing queue (so the information eventually reaches all workers) and notifying the output process via the output queue. If the output process sees all workers have shut down, all events have been written to disk, and it too can shut down.

4.3.2 Raw data format

A XENON1T raw data directory consists of several ZIP archives using the naming convention illustrated in table 4.3. Generally, an individual archive contains 100 events, striking a compromise between file size (approximately 1.5 MB/event) and number of files.

Inside an archive, each event is stored in a single file, whose name is simply the event number. The files are compressed with zlib rather than the (less optimal) compression of standard ZIP archives. The event files contain event data in the same format used internally by the XENON1T processor **pax**, encoded using Python’s **pickle** module. This format showed the highest encode speed among the formats we tested. Other formats (such as MessagePack or BSON) would result in slightly more compact files and would make our data readable without installing the XENON1T processor (or even in languages other than python). Since the encoding is done in the worker stage, which has been parallelized (see above), encoding speed is no longer a serious concern, and an event format switch could be considered.

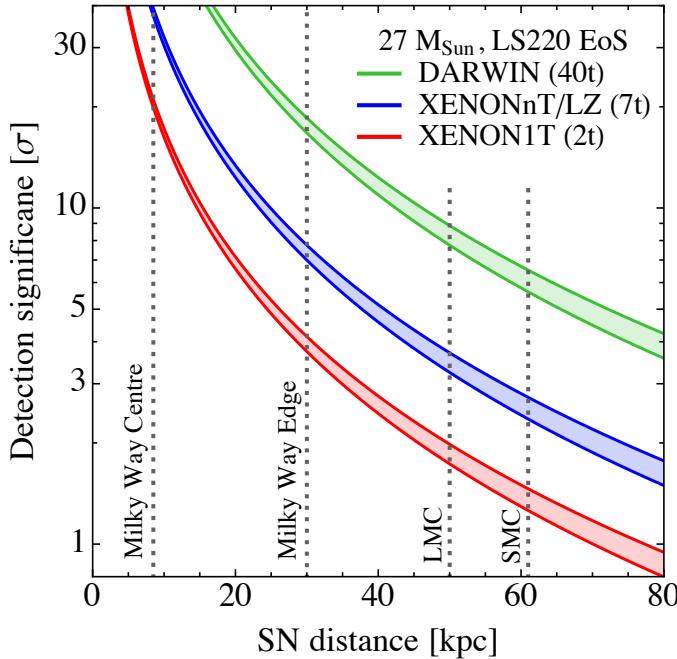


Figure 4.10: From [170]: detection significance for supernovas at different distances from earth, for XENON1T (red), its upgrade XENONnT (blue), and the proposed 40-ton LXe detector DARWIN (green). The band width reflects theoretical uncertainties. Vertical lines mark distances corresponding to the galactic center, edge, and the large and small Magellanic clouds (LMC, SMC, respectively)

4.3.3 Supernova neutrino detection

Supernovas emit neutrino bursts, which offer a unique window for studying the physics of collapsing stellar cores. Dark matter detectors such as XENON1T can detect these neutrinos via coherent neutrino-nucleus scattering (CNNS)[167, 168], especially in an S2-only search. Current and planned future LXe detectors can detect supernovas in or near the Milky Way galaxy (see figure 4.10), which occur roughly twice per century [169, 170]. The expected neutrino event rate in LXe dark matter detectors from such supernovas is far lower than that of dedicated neutrino telescopes such as Icecube [170]. However, CNNS gives dark matter detectors sensitivity to neutrinos and antineutrinos of all flavours, unlike e.g. the muon production channel accessed by a Cerenkov detector. Thus, should a galactic supernova occur, it will be interesting to look for its neutrino signal in XENON1T.

While an S2-only search using the normal data stream is feasible (see chapter 7), it would be challenging to detect a supernova signal independently. Instead, we can look in a short window around supernova signals observed in other detectors. This greatly reduces the impact of detector backgrounds, and

makes the trigger threshold a limiting factor.

Thanks to triggerless readout and the supernova early warning system SNEWS [171], we can save the full untriggered data when a galactic supernova occurs, then re-analyze the data with an arbitrarily low threshold. This involves two programs called *deleter* and *delete decider*, which ensure untriggered data is kept in MongoDB for some time after it is acquired, and eventually deleted unless an SNEWS trigger occurs.

Table 4.4 illustrates how this works for a regular XENON1T run. The *delete decider* periodically checks for new supernova triggers, and if it finds one, marks a run for preservation. It also searches for runs that have stopped two hours ago and were not marked for preservation, then marks these for deletion by setting the `trigger_status` field in the run document to `deleteme`. The *deleter* is a simple script that constantly watches for runs with the `deleteme` status and deletes them. If either of the scripts crash (e.g. because of an error in the rarely-reached parts of the code that actually handles supernova triggers) or their host machines silently drop off the network, data deletion is halted. This will eventually crash the DAQ, but the valuable supernova data will be preserved.

For high-rate calibration data runs, this procedure would be undesirable, as it leaves untriggered data in the MongoDB buffer for two to three hours. Moreover, it is unlikely we would be able to extract a reliable supernova signal while the detector is being calibrated. Hence, during calibration runs, the event builder itself continuously deletes data from the MongoDB buffer; the deleter and delete decider are not involved, and we are not sensitive to supernovae.

Table 4.4: Life-cycle of a normal XENON1T run

| Time | trigger_status | DAQ | Event builder | Deleter | Delete decider |
|-----------------|-------------------------|--------------------|-----------------|-----------------|--------------------|
| Run start | waiting_to_be_processed | Starts acquisition | | | |
| Run stop | staging | | Starts building | | |
| Two hours later | processed | Finished | Finished | | |
| | deleting | | | Starts deletion | |
| | deleted | | | Finished | Marks for deletion |

Chapter 5

First dark matter search results

This chapter has been published in [6], co-authored by the XENON1T collaboration and co-written by the author with Christopher Tunnell and Guillaume Plante. Section headers have been added for clarity.

5.1 Introduction

5.1.1 Dark matter

Modern cosmology precisely describes observational data from the galactic to cosmological scale with the Λ cold dark matter model [172, 173]. This model requires a nonrelativistic nonbaryonic component of the Universe called *dark matter*, with an energy density of $\Omega_c h^2 = 0.1197 \pm 0.0022$ as measured by Planck [1]. Theories beyond the Standard Model of particle physics (*e.g.*, supersymmetry [7]) often attribute this energy density to weakly interacting massive particles (WIMPs) that may be detectable by underground detectors [174, 49].

5.1.2 XENON1T

The XENON1T experiment is designed primarily for detecting nuclear recoils (NRs) from WIMP-nucleus scattering, continuing the XENON program [175, 141] that employs dual-phase (liquid-gas) xenon time projection chambers (TPCs) [141, 68]. With a total mass of \sim 3200 kg of ultra-pure liquid xenon – more than two orders of magnitude larger than the initial detector of the XENON project [175] – XENON1T is the first detector of such scale realized to date. It is located at the Laboratori Nazionali del Gran Sasso (LNGS) in Italy, at an average depth of 3600 m water equivalent. The approximately 97-cm long by 96-cm wide cylindrical TPC encloses (2004 ± 5) kg of liquid xenon (LXe), while another \sim 1200 kg provides additional shielding. The TPC

is mounted at the center of a 9.6-m diameter, 10-m tall water tank to shield it from ambient radioactivity. An adjacent service building houses the xenon storage, cryogenics plant, data acquisition, and slow control system. The water tank is mounted with 84 photomultiplier tubes (PMTs) as part of a Cherenkov muon veto [131]. The TPC is instrumented with 248 3" Hamamatsu R11410-21 PMTs arranged in two arrays above and below the LXe target [129, 128]. Interactions in the target produce scintillation photons (S1) and ionization electrons. The electrons drift in a (116.7 ± 7.5) V/cm electric field towards the liquid-gas interface at the top of the TPC. They are extracted into the gas by an electric field $E_{\text{gas}} > 10$ kV/cm where, via electroluminescence, they produce a proportional scintillation signal (S2). This charge-to-light amplification allows for the detection of single electrons [136, 137]. The ratio of the S2 to S1 signals is determined by both the ratio of ionization to excitation in the initial interaction and subsequent partial recombination of the ionization, with lower S2/S1 for NR signals than electronic recoils (ERs) from γ and β radiation.

5.1.3 First science run

Here we report on 34.2 live days of blinded dark matter search data from the first science run of the experiment. The run started on November 22, 2016, and ended on January 18, 2017, when an earthquake temporarily interrupted detector operations. The detector's temperature, pressure, and liquid level remained stable at (177.08 ± 0.04) K, (1.934 ± 0.001) bar, and (2.5 ± 0.2) mm respectively, where the liquid level was measured above the grounded electrode separating the drift and extraction field regions. While the PMT high voltage remained stable during the run, 27 PMTs were turned off for the dark matter search and 8 were masked in the analysis due to low single-photoelectron (PE) detection efficiency. The PMT response was calibrated periodically using pulsed LED data [147]. The xenon was continuously purified in the gas phase through hot metal getters, leading to an increase in the electron lifetime from 350 to 500 μs , with an average of 452 μs ; 673 μs is the drift time over the length of the TPC. Using cryogenic distillation [132], the ${}^{\text{nat}}\text{Kr}$ concentration in the LXe was reduced while the TPC was in operation, from (2.60 ± 0.05) ppt [mol/mol] at the beginning of the science run to (0.36 ± 0.06) ppt one month after the end of the science run, as measured by rare-gas mass spectrometry [176] on samples extracted from the detector. The ${}^{214}\text{Pb}$ event rate was $(0.8 - 1.9) \times 10^{-4}$ events/(kg \times day \times keV_{ee}) in the low-energy range of interest for WIMP searches, where the bounds are set using *in-situ* α -spectroscopy on ${}^{218}\text{Po}$ and ${}^{214}\text{Po}$. The ${}^{222}\text{Rn}$ concentration was reduced by $\sim 20\%$ relative to the equilibrium value using the krypton distillation column in inverse mode [133].

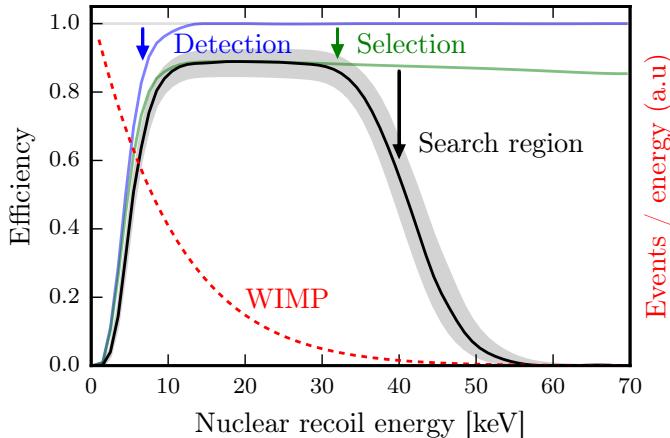


Figure 5.1: NR detection efficiency in the fiducial mass at successive analysis stages as a function of recoil energy. At low energy, the detection efficiency (blue) dominates. At 20 keV, the efficiency is 89% primarily due to event selection losses (green). At high energies, the effect of restricting our data to the search region described in the text (black) is dominant. The black line is our final NR efficiency, with uncertainties shown in gray. The NR energy spectrum shape of a 50-GeV/c² WIMP (in a.u.) is shown in red for reference.

5.2 Data analysis

5.2.1 DAQ

The data acquisition (DAQ) system continuously recorded individual PMT signals. The efficiency for recording single-PE pulses was 92% on average during the science run, and stable to within 2%. A software trigger analyzed the PMT pulses in real-time, allowing for continuous monitoring of the PMTs. The trigger detected S2s larger than 200 PE with 99% efficiency, and saved 1 ms before and after these to ensure small S1s were captured. An analog-sum waveform was separately digitized together with a signal recording when any of the digitizers were inhibited. The average DAQ live time was 92% during the science run.

5.2.2 Data processing

Physical signals are reconstructed from raw data by finding photon hits in each PMT channel, then clustering and classifying groups of hits as S1 or S2 using the PAX software. For S1s, we require that hits from three or more PMTs occur within 50 ns. To tune the signal reconstruction algorithms and compute their efficiency for detecting NRs – shown in blue in Fig. 5.1 – we used a Monte Carlo code that reproduces the shapes of S1s and S2s as determined by the interaction physics, light propagation, and detector-electronics chain.

This was validated against $^{83\text{m}}\text{Kr}$ and ^{220}Rn calibration data.

5.2.3 Position reconstruction and corrections

The interaction position is reconstructed from the top-array PMT hit pattern of the S2 (for the transverse position) and the time difference between S1 and S2 (for depth). The S2 transverse position is given by maximizing a likelihood based on an optical simulation of the photons produced in the S2 amplification region. The simulation-derived transverse resolution is ~ 2 cm at our S2 analysis threshold of 200 PE (uncorrected). The interaction position is corrected for drift field nonuniformities derived from a finite element simulation, which is validated using $^{83\text{m}}\text{Kr}$ calibration data. We correct S2s for electron losses during drift, and both S1s and S2s for spatial variations of up to 30% and 15%, respectively, inferred from $^{83\text{m}}\text{Kr}$ calibration data. These spatial variations are mostly due to geometric light collection effects. The resulting corrected quantities are called cS1 and cS2. As the bottom PMT array has a more homogeneous response to S2 light than the top, this analysis uses cS2_b, a quantity similar to cS2 based on the S2 signal seen only by the bottom PMTs.

5.2.4 Calibration data

To calibrate XENON1T, we acquired 3.0 days of data with ^{220}Rn injected into the LXe (for low-energy ERs), 3.3 days with $^{83\text{m}}\text{Kr}$ injected into the LXe (for the spatial response) and 16.3 days with an external $^{241}\text{AmBe}$ source (for low-energy NRs). The data from the ^{220}Rn [177] and $^{241}\text{AmBe}$ calibrations is shown in Fig. 5.2 (a) and (b), respectively. Following the method described in [178] with a W -value of 13.7 eV, we extracted the photon gain $g_1 = (0.144 \pm 0.007)$ PE per photon and electron gain $g_2 = (11.5 \pm 0.8)$ PE (in the bottom array, 2.86 times lower than if both arrays are used) per electron in the fiducial mass by fitting the anti-correlation of cS2_b and cS1 for signals with known energy from $^{83\text{m}}\text{Kr}$ (41.5 keV), ^{60}Co from detector materials (1.173 and 1.332 MeV), and from decays of metastable $^{131\text{m}}\text{Xe}$ (164 keV) and $^{129\text{m}}\text{Xe}$ (236 keV) produced during the $^{241}\text{AmBe}$ calibration. The cS1 and cS2_b yields are stable in time within 0.77% and 1.2% respectively, as determined by $^{83\text{m}}\text{Kr}$ calibrations.

5.2.5 Event selection

WIMPs are expected to induce low-energy single-scatter NRs. Events that are not single scatters in the LXe are removed by several event-selection cuts: (i) a single S2 above 200 PE must be present and any other S2s must be compatible with single electrons from photoionization of impurities or delayed extraction; (ii) an event must not closely follow a high-energy event (*e.g.*, within 8 ms after a 3×10^5 PE S2), which can cause long tails of single electrons; (iii) the S2 signal's duration must be consistent with the depth of the interaction as

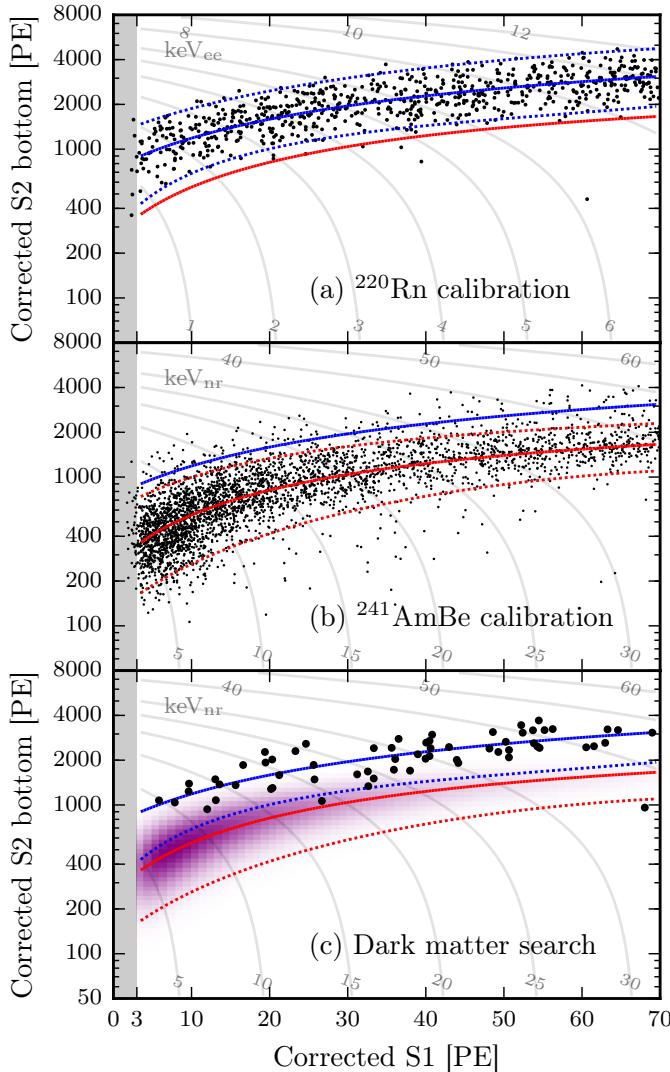


Figure 5.2: Observed data in $cS2_b$ vs. $cS1$ for (a) ^{220}Rn ER calibration, (b) $^{241}\text{AmBe}$ NR calibration, and (c) the 34.2-day dark matter search. Lines indicate the median (solid) and $\pm 2\sigma$ (dotted) quantiles of simulated event distributions (with the simulation fitted to calibration data). Red lines show NR (fitted to $^{241}\text{AmBe}$) and blue ER (fitted to ^{220}Rn). In (c), the purple distribution indicates the signal model of a $50\text{ GeV}/c^2$ WIMP. Thin gray lines and labels indicate contours of constant combined energy scale in keV for ER (a) and NR (b, c). Data below $cS1 = 3$ PE (gray region) is not in our search region and shown only for completeness.

inferred from the drift time; (iv) the S1 and S2 hit patterns must be consistent with the reconstructed position at which these signals were produced; (v) no more than 300 PE of uncorrelated single electrons and PMT dark counts must appear in the region before the S2. Single scatter NR events within the [5, 40] keV_{nr} energy range pass these selections with >82% probability, as determined using simulated events or control samples derived from calibration, and shown in green in Fig. 5.1.

The dark matter search uses a cylindrical (1042 ± 12) kg fiducial mass, which was defined before unblinding using the reconstructed spatial distribution of ERs in the dark matter search data and the energy distribution of ERs from ^{220}Rn . We restrict the search to $\text{cS1} \in [3, 70]$ PE and $\text{cS2}_b \in [50, 8000]$ PE, which causes little additional loss of WIMP signals, as shown in black in Fig. 5.1.

5.3 Background model

Table 5.1 lists the six sources of background we consider inside the fiducial mass and inside the search region. For illustration, we also list the expected rate in a reference region between the NR median and -2σ quantile in cS2_b (*i.e.*, between the red lines in Fig. 5.2c), for which Fig. 5.3 shows the background model projected onto cS1 . This reference region would contain about half of the WIMP candidate events, while excluding 99.6% of the ER background. The WIMP search likelihood analysis uses the full search region. Below we describe each background component in more detail: all event rates are understood to be inside the fiducial mass and the full search region.

5.3.1 Electronic recoils

First, our background model includes ERs, primarily from β decays of ^{85}Kr and the intrinsic ^{222}Rn -progeny ^{214}Pb , which cause a flat energy spectrum in the energy range of interest [68]. The ER background model is based on a simulation of the detector response. We use a model similar to [112] to convert the energy deposition from ERs into scintillation photons and ionization electrons, which we fit to ^{220}Rn calibration data in $(\text{cS1}, \text{cS2}_b)$ space (Fig. 5.2a). The best-fit photon yield and recombination fluctuations are comparable to those of [112]. The model accounts for uncertainties of g_1 , g_2 , spatial variations of the S1 and S2 light-collection efficiencies, the electron-extraction efficiency, reconstruction and event-selection efficiency, and time dependence of the electron lifetime. The rate of ERs is not constrained in the likelihood analysis, even though we have independent concentration measurements for ^{214}Pb and ^{85}Kr , since the most stringent constraint comes from the search data itself.

Table 5.1: Expected number of events for each background component in the fiducial mass; in the full $cS1 \in [3, 70]$ PE, $cS2_b \in [50, 8000]$ PE search region and in a reference region between the NR median and -2σ quantile in $cS2_b$. Uncertainties < 0.005 events are omitted. The ER rate is unconstrained in the likelihood; for illustration, we list the best-fit values to the data in parentheses.

| | Full | Reference |
|--|------------------------|--------------------------|
| Electronic recoils (<i>ER</i>) | (62 ± 8) | $(0.26^{+0.11}_{-0.07})$ |
| Radiogenic neutrons (<i>n</i>) | 0.05 ± 0.01 | 0.02 |
| CNNS (ν) | 0.02 | 0.01 |
| Accidental coincidences (<i>acc</i>) | 0.22 ± 0.01 | 0.06 |
| Wall leakage (<i>wall</i>) | 0.52 ± 0.3 | 0.01 |
| Anomalous (<i>anom</i>) | $0.10^{+0.10}_{-0.07}$ | 0.01 ± 0.01 |
| Total background | 63 ± 8 | $0.36^{+0.11}_{-0.07}$ |
| $50 \text{ GeV}/c^2, 10^{-46}\text{cm}^2$ WIMP | 1.66 ± 0.01 | 0.82 ± 0.06 |

5.3.2 Nuclear recoils

Second and third, our background model includes two sources of NRs: radiogenic neutrons contribute (0.05 ± 0.01) events, and coherent neutrino-nucleus scattering (CNNS) ~ 0.02 events. Cosmogenically produced neutrons are estimated to contribute $\mathcal{O}(10^{-3})$ events even without muon-veto tagging. The NR background model is built from a detector response simulation that shares the same detector parameters and associated systematic uncertainties as the ER background model above. The main difference is the energy-conversion model, where we use the model and parametrization from NEST [110]. We obtain the XENON1T response to NRs by fitting the $^{241}\text{AmBe}$ calibration data (Fig. 5.2b) with the light and charge yields from [179] as priors. Our NR response model is therefore constrained by the global fit of external data. It is also used to predict the WIMP signal models in $(cS1, cS2_b)$ space.

5.3.3 Accidental coincidences

Fourth, accidental coincidences of uncorrelated S1s and S2s are expected to contribute (0.22 ± 0.01) background events. We estimated their rate and $(cS1, cS2_b)$ distribution using isolated S1 and S2 signals, which are observed to be at (0.78 ± 0.01) Hz and (3.23 ± 0.03) mHz, respectively, before applying S2-selections. The effect of our event selection on the accidental coincidence rate is included, similar to [180]. Isolated S1s may arise from interactions in regions of the detector with poor charge collection, such as below the

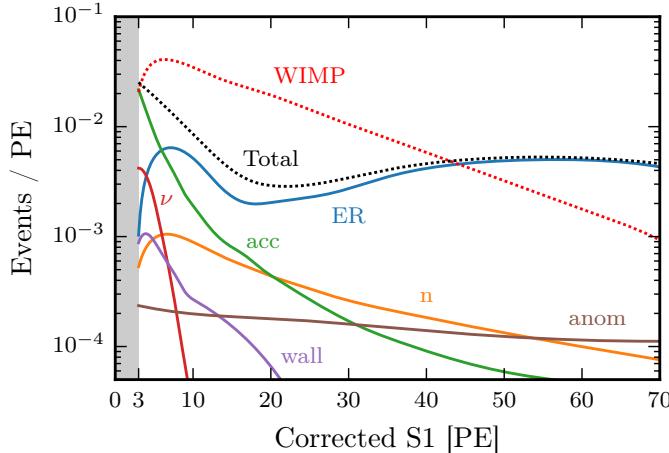


Figure 5.3: Background model in the fiducial mass in a reference region between the NR median and -2σ quantile in $cS2_b$, projected onto $cS1$. Solid lines show that the expected number of events from individual components listed in Table 5.1; the labels match the abbreviations shown in the table. The dotted black line *Total* shows the total background model, the dotted red line *WIMP* shows an $m = 50$ GeV/c 2 , $\sigma = 10^{-46}$ cm 2 WIMP signal for comparison.

cathode, suppressing an associated $cS2$ signal. Isolated $S2$ s might arise from photoionization at the electrodes, regions with poor light collection, or from delayed extraction [140]. Most accidental events are expected at low $cS1$ and at lower $cS2_b$ than typical NRs.

5.3.4 Wall leakage

Fifth, inward-reconstructed events from near the TPC’s PTFE wall are expected to contribute (0.5 ± 0.3) events, with the rate and $(cS1, cS2_b)$ spectrum extrapolated from events outside the fiducial mass. Most of these events would appear at unusually low $cS2_b$ due to charge losses near the wall. The inward reconstruction is due to limited position reconstruction resolution, especially limited for small $S2$ s, near the 5 (out of 36) top PMTs in the outermost ring that are unavailable in this analysis.

5.3.5 Anomalous background

Sixth and last, we add a small uniform background in the $(cS1, \log cS2_b)$ space for ER events with an anomalous $cS2_b$. Such *anomalous leakage* beyond accidental coincidences has been observed in XENON100 [180], and a few such events are seen in the ^{220}Rn calibration data (Fig. 5.2a). If these were not ^{220}Rn -induced events, their rate would scale with exposure and we would see numerous such events in the WIMP search data. We do not observe this, and

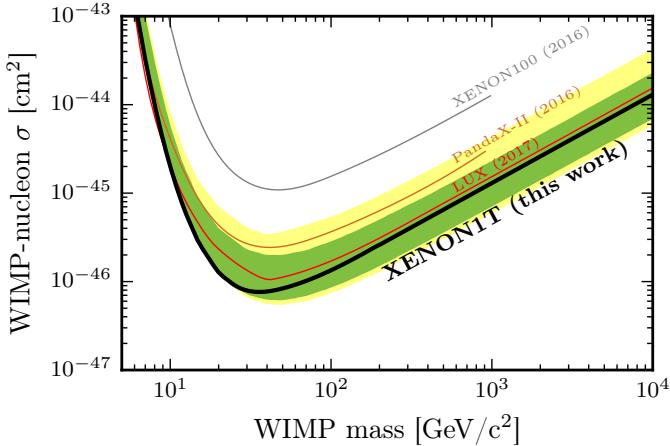


Figure 5.4: The spin-independent WIMP-nucleon cross section limits as a function of WIMP mass at 90% confidence level (black) for this run of XENON1T. In green and yellow are the 1- and 2 σ sensitivity bands. Results from LUX [181] (red), PandaX-II [155] (brown), and XENON100 [180] (gray) are shown for reference.

therefore assume their rate is proportional to the ER rate, at $0.10^{+0.10}_{-0.07}$ events based on the outliers observed in the ^{220}Rn calibration data. The physical origin of these events is under investigation.

5.4 Results

5.4.1 Blinding

The WIMP search data in a predefined signal box was blinded (99% of ERs were accessible) until the event selection and the fiducial mass boundaries were finalized. We performed a staged unblinding, starting with an exposure of 4 live days distributed evenly throughout the search period. No changes to the event selection were made at any stage.

5.4.2 Observed events

A total of 63 events in the 34.2-day dark matter search data pass the selection criteria and are within the $\text{cS1} \in [3, 70]$ PE, $\text{cS2}_b \in [50, 8000]$ PE search region used in the likelihood analysis (Fig. 5.2c). None are within 10 ms of a muon veto trigger. The data is compatible with the ER energy spectrum in [68] and implies an ER rate of $(1.93 \pm 0.25) \times 10^{-4}$ events/(kg \times day \times keV_{ee}), compatible with our prediction of $(2.3 \pm 0.2) \times 10^{-4}$ events/(kg \times day \times keV_{ee}) [68] updated with the lower Kr concentration measured in the current science run. This is the lowest ER background ever achieved in such a dark matter

experiment. A single event far from the bulk distribution was observed at $cS1 = 68.0$ PE in the initial 4-day unblinding stage. This appears to be a *bona fide* event, though its location in $(cS1, cS2_b)$ (see Fig. 5.2c) is extreme for all WIMP signal models and background models other than anomalous leakage and accidental coincidence. One event at $cS1 = 26.7$ PE is at the -2.4σ ER quantile.

5.4.3 Statistical interpretation

For the statistical interpretation of the results, we use an extended unbinned profile likelihood test statistic in $(cS1, cS2b)$. We propagate the uncertainties on the most significant shape parameters (two for NR, two for ER) inferred from the posteriors of the calibration fits to the likelihood. The uncertainties on the rate of each background component mentioned above are also included. The likelihood ratio distribution is approximated by its asymptotic distribution [182]; preliminary toy Monte Carlo checks show the effect on the exclusion significance of this conventional approximation is well within the result's statistical and systematic uncertainties. To account for mismodeling of the ER background, we also calculated the limit using the procedure in [183], which yields a similar result.

The data is consistent with the background-only hypothesis. Fig. 5.4 shows the 90% confidence level upper limit on the spin-independent WIMP-nucleon cross section, power constrained at the -1σ level of the sensitivity band [184]. The final limit is within 10% of the unconstrained limit for all WIMP masses. For the WIMP energy spectrum we assume a standard isothermal WIMP halo with $v_0 = 220$ km/s, $\rho_{\text{DM}} = 0.3$ GeV/cm 3 , $v_{\text{esc}} = 544$ km/s, and the Helm form factor for the nuclear cross section [72]. No light and charge emission is assumed for WIMPs below 1 keV recoil energy. For all WIMP masses, the background-only hypothesis provides the best fit, with none of the nuisance parameters representing the uncertainties discussed above deviating appreciably from their nominal values. Our results improve upon the previously strongest spin-independent WIMP limit for masses above 10 GeV/c 2 . Our strongest exclusion limit is for 35-GeV/c 2 WIMPs, at 7.7×10^{-47} cm 2 .

5.5 Outlook

These first results demonstrate that XENON1T has the lowest low-energy background level ever achieved by a dark matter experiment. The sensitivity of XENON1T is the best to date above 20 GeV/c 2 , up to twice the LUX sensitivity above 100 GeV/c 2 , and continues to improve with more data. The experiment resumed operation shortly after the January 18, 2017 earthquake and continues to record data.

Chapter 6

Statistical data analysis

This chapter focuses on the statistical methods used for the first XENON1T results. We start with a general description of a likelihood analysis framework for XENON1T in section 6.1, and show its application to background modelling in section 6.2. In section 6.3 we examine three topics related to the dark matter limit setting itself: non-asymptotic behaviour of the likelihood ratio test statistic, the impact of (potential) outlier events, and the effect of various small changes and uncertainties on the assumptions used in the result. We end in section 6.4 with recommendations for further science runs of XENON1T and future dark matter experiments.

6.1 Methodology

For some analyses it suffices to have a simple analytical model of a feature in the data (for example, determining the electron lifetime from a mono-energetic feature at high-energy). For others, a Monte Carlo simulation of parts or all of the signal generation process in the detector is needed (for example, determining the distribution of the electronic recoil background). This latter type of analysis faces several technical challenges:

1. We need to quantify the similarity between two sets of events (simulation and data) rather than a single set of events and a model. Even restricting attention to likelihood-based methods, there are several sensible prescriptions – based on binning, kernel density estimation, nearest-neighbour methods, etc. – and it is not always clear which is superior in an individual case.
2. Statistical errors increase in importance near a minimum of the test statistic. As differences between successive runs of the simulations become smaller, statistical variation due to the finite (and usually fixed) Monte Carlo sample size will dominate the difference between successive iterations. This means the test statistic will be a noisy, non-deterministic

function of parameter space, especially in the region we are most interested in.

3. Repeated simulation runs to explore the parameter space can quickly escalate the computation requirements. However, many iterations are usually required due to the complexity of the model. A large number of uncertain and possibly degenerate parameters, e.g. to capture systematic uncertainties, is often the reason to use a simulation-based model in the first place. However, this results in a high-dimensional space of possible models with multiple local minima and saddle points.

For XENON1T, we address the first challenge by constructing probability density function (PDF) estimates of the various event sources from histograms of the results of high-statistics simulations ($10^4 - 10^6$ times the number of expected events), then compare data against these with an extended unbinned likelihood ratio test statistic (see section 6.1.1). The second and third challenge are addressed to some extent by interpolating PDF estimates on a precomputed grid in the parameter space of interest (also known as ‘vertical template morphing’ [185]) rather than re-running the simulation at every parameter space point during optimization (see section 6.1.2).

6.1.1 Likelihood function and inference method

Rare-event searches work with sparse science data, where the precise location of one or two events in parameter space can play a decisive role. Thus, the XENON1T collaboration chose to work primarily with an extended unbinned likelihood:

$$L = \text{PoissonPDF}(N_{\text{tot}}|\mu) \prod_i^{\text{events}} \left(\sum_j^{\text{sources}} \frac{\mu^{(j)}}{\mu} \text{pdf}^{(j)}(\text{event } i) \right) \quad (6.1)$$

where N_{tot} is the total number of detected events, μ the expected number of events, $\mu^{(j)}$ the expected events from source j , and $\text{pdf}^{(j)}$ the probability density function in the analysis parameter space (cS1, cS2_b for XENON1T’s main analysis) for events from source j . The μ ’s and PDF’s are all functions of the likelihood’s parameters, which depend on the application. In some cases, additional constraints are added to the likelihood (by multiplication), e.g. for uncertainties on the rates of various backgrounds. In practice we work with the logarithm of the likelihood, which is

$$\log L = -\mu + \sum_i^{\text{events}} \log \left(\sum_j^{\text{sources}} \mu^{(j)} \text{pdf}^j(\text{event}_i) \right), \quad (6.2)$$

up to a multiplicative (after log, additive) factor independent of the inference parameters.

For setting confidence intervals on parameters, we must choose between the frequentist and Bayesian statistical paradigms. Rare-event searches in physics almost universally use frequentist statistics, as Bayesian methods require specification of a prior on the rate (or cross-section) of the sought-for process. If the observed data strongly points to a particular non-zero rate, the resulting interval will not depend much on the prior; but upper limits from null results generally do: a prior uniform within some range of the logarithm of the cross-section results in upper limits that depend sensitively on the chosen range [186].

For frequentist inference, we must choose a test statistic that quantifies the extremeness of the observed data under several models. When computing upper limits on the WIMP-nucleon cross-section σ , we use a modified profile likelihood ratio test statistic:

$$t(\sigma) = \begin{cases} -2 \log \frac{L(\text{best fit with } \sigma \text{ fixed})}{L(\text{best fit})} & \text{best-fit } \sigma < \sigma \\ 0 & \text{otherwise} \end{cases}. \quad (6.3)$$

Here the denominator in the logarithm is the likelihood of the global best fit (in all parameters considered; parameters of L beyond the variable of immediate interest, σ , are called *nuisance parameters*) and the numerator is the likelihood of the best fit among all considered parameters, under the condition that the WIMP-nucleon cross-section is fixed to σ . The statistic is set to 0 if the best-fit σ is larger than the tested σ , and is positive otherwise. The larger the value of $t(\sigma)$, the stronger the evidence against a dark matter model.

To quantify the strength of evidence against a model, we must know the distribution of $t(\sigma)$ (for the case that σ is the true WIMP cross-section). This is straightforward in the asymptotic case of infinitely many observed events and no nuisance parameters; the best fit σ will be infinitely close to the true σ , just above it in 50% of cases (so $t = 0$), just below it in the other 50%. By an extension of Wilks' theorem, it is distributed as a chi-squared distribution with one degree of freedom in the latter case; thus [182]:

$$\text{PDF}(t) = \frac{1}{2} \delta(0) + \frac{1}{2} \frac{\exp(-t/2)}{\sqrt{2\pi t}} \quad (6.4)$$

where δ is the Dirac delta function. As stated in section 5.4.3, we assumed this was approximately true in the XENON1T first results. We examine the validity of this assumption in section 6.3.2.

It is possible to observe data that is extreme under any signal hypothesis, especially in experiments with significant backgrounds. For example, a counting experiment with an expected background of 100 events might see only 50 events, which is extreme under any positive signal expectation. Correctly interpreted, such a result is strong evidence against the combined hypothesis of a signal *and* the experiment's background model. However, it translates

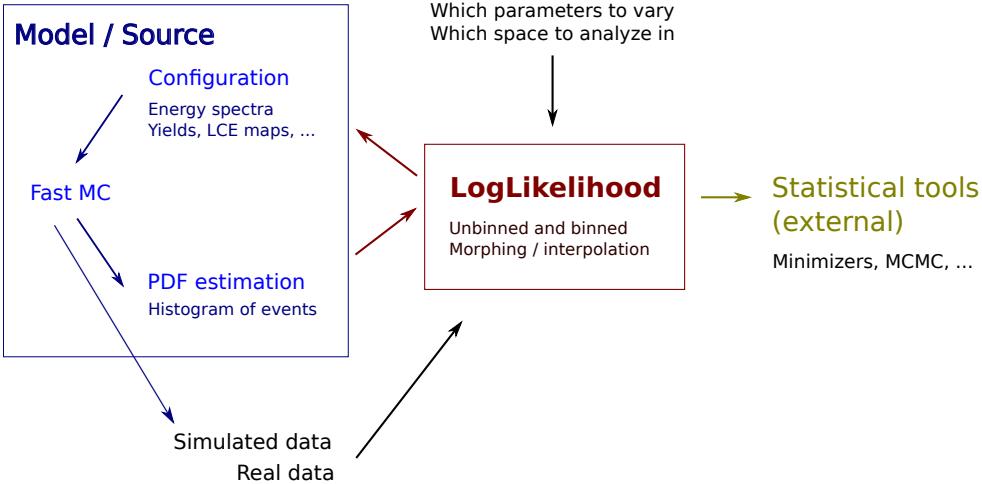


Figure 6.1: Schematic architecture of the `blueice` software; described in the text.

to a very strong upper limit – $n_{\text{signal}} = 0$ in this case, apparently excluding the existence of WIMPs altogether. To limit the impact of such background underfluctuations on the limit, whether they arise by chance or mismodelling, The XENON1T collaboration chose to constrain the limit to the 15.8th percentile (-1σ) of the distribution of limits in the background-free case. This is a simplified form of the power-constrained limits approach advocated in [184]. The unconstrained limit from chapter 5 is within 10% of this constraint for $m_\chi < 30 \text{ GeV}/c^2$ – whether the constraint is actually binding depends on details of the interpolation of the limit curve between different m_χ points.

6.1.2 PDF determination

Likelihood inference with equation 6.1 requires probability density functions of all considered signal and background models. We developed a python package named `blueice` to streamline this task within XENON1T’s software environment. Figure 6.1 illustrates the architecture of `blueice`:

- The user interacts with a `LogLikelihood` object, representing the likelihood function to specify the observed dataset, which of the Monte Carlo’s settings become parameters in the analysis (e.g. the rate of WIMPs, electron lifetime or photon detection efficiency), and which data dimensions become the space for comparison (e.g. `cS1` and `cS2b`).
- The `LogLikelihood` then behaves as a regular python function, which can be passed to (external) minimizers, plotted, fed to a Markov chain sampler, etc. Some glue code for standard minimizers (`scipy.optimize` and `iminuit`) and the Markov Chain Monte Carlo package `emcee` [187]

is included for convenience. The `scipy` optimizers (specifically Powell’s method) are used most often in practice.

- Under the hood, the `LogLikelihood` function constructs `Models`, each representing a single combination of parameters. A model contains several `Sources`, which are usually interfaces to a Monte Carlo simulator, but can be created from analytical or numerical PDFs. If the former case, we must estimate the PDF from a large simulated dataset before we can use the likelihood. Currently this is done by histogramming; although kernel density estimation is generally more accurate for a fixed number of events, this takes additional computation time, which can also be used to simply simulate more events and employ a finer binning (assuming the simulation is reasonably fast).

During likelihood inference, a minimizer will request the value of the likelihood function at many different parameter space points. To avoid having to run a new simulation every time, we precompute the PDF on a grid (of user-specified size) of ‘anchor points’ in parameter space, then interpolate the PDFs at parameter space points in between. Rather than interpolating the entire grid of PDF values, we interpolate only the PDF scores of the observed events themselves. This is efficient with few events and a finely sampled PDF, which is the more common case in XENON1T. Interpolating the PDFs (as we do) is very different from interpolating the likelihood function evaluated at the anchor points: in the latter case, a minimum would always be at an anchor point.

Since parameters often change the histograms in nonlinear ways, accurate results often require a relatively dense anchor grid (5-7 anchor points not too far from the best-fit value); moreover, the number of models to compute increases exponentially as more parameters are included in the inference. This means we still need substantial computational power, and face constraints on the number of parameters we can practically consider per individual background source. In the future, PDF estimation methods more advanced than linear interpolation might relieve some of this burden.

6.2 ER response fitting

Electronic recoils are XENON1T’s most numerous and most important background. Although the $(cS1, cS2_b)$ distribution of ERs only partially overlaps with that of dark matter candidates, but ERs are far more common. Thus, small changes in the ER model can have large effects on the dark matter sensitivity.

As discussed in section 2.3.3.3, the electronic recoil background of XENON1T is dominated by ^{222}Rn beta decays, which have an effectively flat energy spectrum below 10 keV. The internal ^{220}Rn isotope used for calibration (see section

2.3.3) has the same property. In principle, given enough events, we could simply histogram the calibration data in (cS1, cS2_b) and use it as our ER background model.

Unfortunately, the ²²⁰Rn also produces many high-energy events, as two alpha decays precede the emission of any betas from a decaying atom (see the decay chain in equations 2.25 and 2.26). These cause long tails of PMT afterpulses and delayed extraction electrons, contaminating the low-energy events of interest. Thus, almost all the ²²⁰Rn data acquired for the first science run had to be disregarded; leaving only the data shortly after the ²²⁰Rn injection stopped, when we no longer see the initial pair of alpha decays, and the beta decay rate is slowly falling.¹ This offers only 691 events in the search space of the first science run (see section 5.2.5), enough to necessitate a fitted model. A fitted model also has other advantages, such as the ability to account for differing detector conditions (especially electron lifetime) between calibration and search data.

As discussed in section 2.3, the charge yield Q_y and light yield L_y (both in quanta per unit energy) of liquid xenon for electronic recoils sum to a constant of $1/W \approx 73$ quanta/keV. With this assumption, we only need to fit Q_y to our data. Here we will fit a polynomial:

$$Q_y = a\epsilon^2 + b\epsilon + c, \quad (6.5)$$

$$\epsilon = \log(E_R/E_0),$$

with a, b, c free parameters, E_R the recoil energy and $E_0 = 5$ keV, an arbitrary reference energy. The two-dimensional pdf in (cS1, cS2_b) is constructed by Monte Carlo simulation, accounting for detection efficiencies and their fluctuation due to electron lifetime and spatially varying S1 response. Field distortion and variations in the S2 response cause minor fluctuations inside the fiducial volume (assuming our correction is accurate) and are ignored. The energy of an interaction does not divide binomially over photons and electrons according to charge and light yield, but suffers an additional fluctuation, presumed to be due to track-by-track differences in the recombination fraction [188, 112]. We will model by drawing the probability that a created quantum ends its life as an electron from a Gaussian with standard deviation σ (the mean follows immediately from Q_y).

Table 6.1 shows the best fit parameters, and figure 6.2 shows how the resulting Q_y as a function of energy compares to the NEST model and results from other experiments. We see Q_y is approximately linear in $\log E$; indeed, the coefficient a of the second power in 6.5 is smallest and most uncertain. Our Q_y is a bit lower than the results from LUX's tritium calibration [112]. This could be explained by their slightly higher drift field (180 V/cm versus

¹Stating we took 3.0 days of ²²⁰Rn calibration data, as we did in section 5.2.4, obscures the fact that the rate is far from constant: about half the events come from the first twelve hours.

| Parameter | Best-fit value | Correlation with | | | |
|-----------|-------------------|------------------|---------|---------|----------|
| | | a | b | c | σ |
| a | 15 ± 4 | | $+0.06$ | -0.62 | -0.2 |
| b | -27.7 ± 1.1 | $+0.06$ | | -0.38 | -0.3 |
| c | 32.5 ± 0.3 | -0.62 | -0.38 | | $+0.12$ |
| σ | $(4.5 \pm 0.3)\%$ | -0.2 | -0.3 | $+0.12$ | |

Table 6.1: ER best-fit parameters for equation 6.5 and their correlation coefficients, for ^{220}Rn calibration during XENON1T’s first science run. Uncertainties are 68% credible intervals (assuming a uniform prior in parameter space) derived from a Markov chain Monte Carlo sampler.

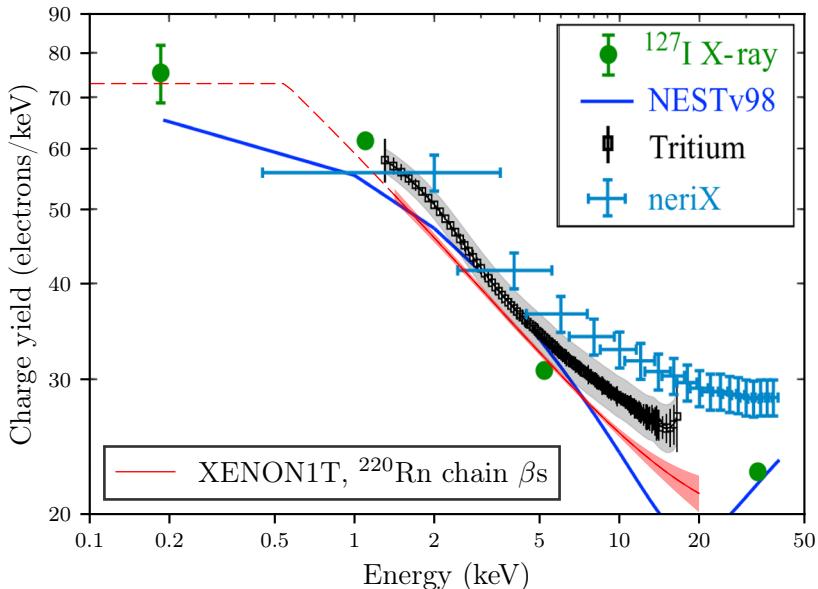


Figure 6.2: Charge yield (Q_y) that best fits XENON1T’s ER data (solid red), overplotted on results from NEST, LUX and neriX (from [111], see figure 2.7 for additional references and details). Shaded region represent the inner 68% of curves obtained from samples of the Markov chain sampler. The dashed line represents the extrapolation of the model to low-energy.

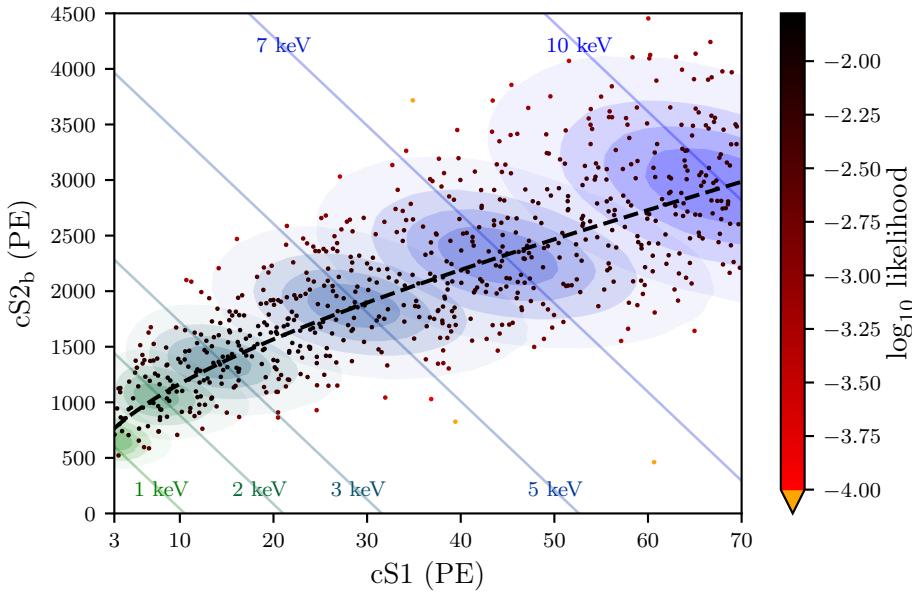


Figure 6.3: Dots: ^{220}Rn calibration data of XENON1T after all data quality cuts, with color shading indicating \log_{10} likelihood under the best-fit ER model. Solid lines indicate constant combined energy scale (equation 2.24) as labeled on the figure. Shaded regions show the response of a hypothetical mono-energetic ER source at the indicated energies: contours indicate the 25th, 50th, 75th and 99th percentile in density. The dashed line shows the nominal signal (with all fluctuations suppressed) of electronic recoils at different energies.

our 117 V/cm); in general, a higher drift field leads to a higher Q_y (see e.g. [113]), since the increased drift velocity suppresses recombination.

Figure 6.3 shows the median ER signal according to the fit in solid black, and the hypothetical responses of mono-energetic ER sources of various energies in green to blue. We see the major axes of the (approximately) ellipsoid responses do not align with contours of constant combined energy scale, but are somewhat more horizontal. This is because the statistical variation on the S1 signal is very significant at these small signal sizes. Figure 6.3 also shows the likelihood contribution (pdf times μ) of each observed ^{220}Rn event under the best-fit model. An outlier is visible at $\sim 60\text{PE}$ below the main band of events, to which no likelihood value can be assigned (based on 10^6 Monte Carlo events). The effect of outliers on the fit was limited by capping the most extreme likelihood contribution per event at 10^{-4} , which also affects two other events. Choosing 10^{-8} as a threshold (which affects only the extreme outlier) leads to the same result, i.e. the best-fit parameters change much less than the fitter’s statistical error. Section 6.3.1 discusses outlier events in likelihood analyses in more detail.

Table 6.2 shows the results of applying various one-dimensional goodness-of-fit tests to the data, in $cS1$, $cS2_b$, the combined energy scale (CES) of equation 2.24, and a ‘CES conjugate’ which is simply the orthogonal coordinate α to CES:

$$\alpha = W \left(\frac{cS1}{g_1} - \frac{cS2}{g_2} \right). \quad (6.6)$$

The results are consistent with chance fluctuations from the best-fit model, though it should be noted the binning for the χ^2 test-in the CES conjugate dimension does not extend far enough to include the outlier event; if it did, the p-value would of course be 0. Figure 6.4 illustrates the χ^2 -test in the CES dimension.

For the XENON1T first results presented in chapter 5, a different ER model was used, based on a fit with many more free parameters, including a more complicated parametrization of Q_y and the recombination fluctuation, as well as small uncertainties on several other variables (such as the processor’s efficiency, g_1 and g_2 , the double photoelectron emission probability, etc.) [189]. This could not practically be fitted by the method of section 6.1.2; instead, the event simulation was re-evaluated at each parameter space point during minimization. We argued against this approach in section 6.1, as statistical errors due to the finite Monte Carlo sample size will be significant near the best fit point, and likely broaden the estimated posterior. The computational requirements are also much larger, taking up to several days on a dedicated GPU cluster rather than a few minutes on an ordinary multi-core CPU machine. On the other hand, the method avoids PDF interpolation, and is free from any errors resulting from an unfortunate choice of anchor models.

Using many free parameters, especially on a limited dataset, increases the potential for overfitting. In fact, the XENON1T ER model had to be adjusted a few times for this reason, after goodness-of-fit tests showed it was performing poorly originally. The model as used in chapter 5 shows similar goodness of fit to the ^{220}Rn data as the model presented here.

A final issue is how to propagate the remaining uncertainty on the ER model to the WIMP search likelihood. Propagating the correlated uncertainties on all fit parameters shown in table 6.1 would take considerable computational expense, due to the four required shape uncertainties. In this case, there is no advantage to performing the ER fit separately; a simultaneous fit would be almost equally expensive (since the main computational expense is simulating events for PDF modelling). Choosing a subset of parameters to propagate is not ideal either: if an ignored parameter is correlated with a propagated parameter, varying only the propagated parameter results in models that may not fit the calibration data at all, leading to excessive uncertainty on the background shape.

Hence, we apply normalization and principal component analysis (PCA) on the posterior distributions of the polynomial coefficients $\mathbf{v} = (c, b, a)$ and transform them to three uncorrelated components $\mathbf{p} = (p_0, p_1, p_2)$ whose

| Statistic | cS1 | cS2 _b | CES | CES conjugate |
|-------------------------------|------|------------------|------|---------------|
| Kolmogorov-Smirnov | 0.74 | 0.34 | 0.57 | 0.41 |
| Anderson-Darling | 0.80 | 0.19 | 0.52 | 0.80 |
| χ^2 , 20 bins | 0.95 | 0.97 | 0.82 | 0.90 |
| χ^2 , 20 bins, 15 d.o.f. | 0.80 | 0.89 | 0.57 | 0.84 |

Table 6.2: ER fit goodness-of-fit test p-values for different tests (rows) in different dimensions (columns). For the Kolmogorov-Smirnov and Anderson-Darling tests, p-values are computed by comparing the goodness-of-fit statistic to that of simulated datasets drawn from the best-fit model. For the χ^2 test, p-values are computed in the same way (third row) and analytically from the χ^2 distribution with 15 degrees of freedom (fourth row). The latter accounts for the fact that the model has parameters which are fitted to the data (the four parameters of table 6.1 and the total event rate); this reduces the p-values, but not enough to alter the conclusions.

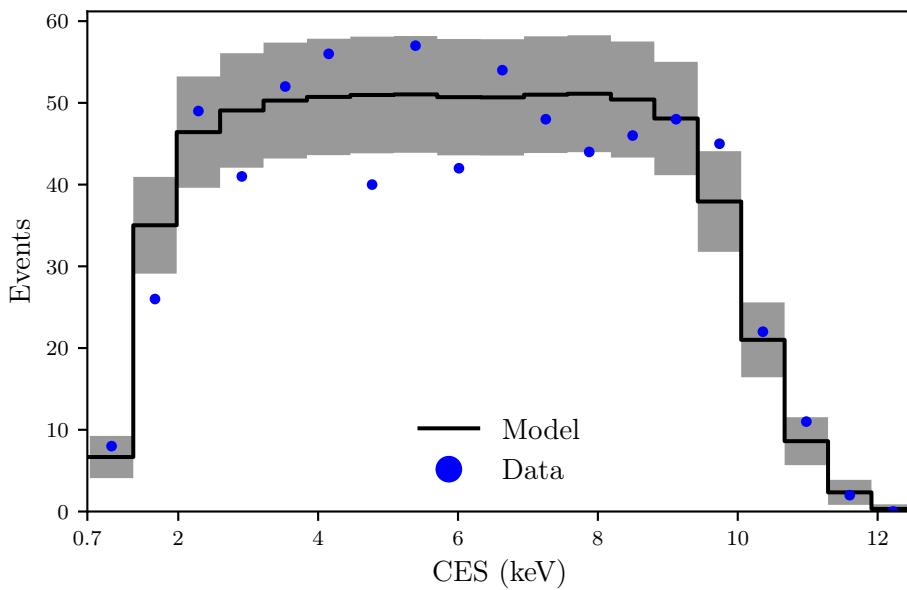


Figure 6.4: ER data in CES (blue dots) overlaid on best-fit model prediction (black). Grey bands show expected statistical uncertainty according to the best-fit model.

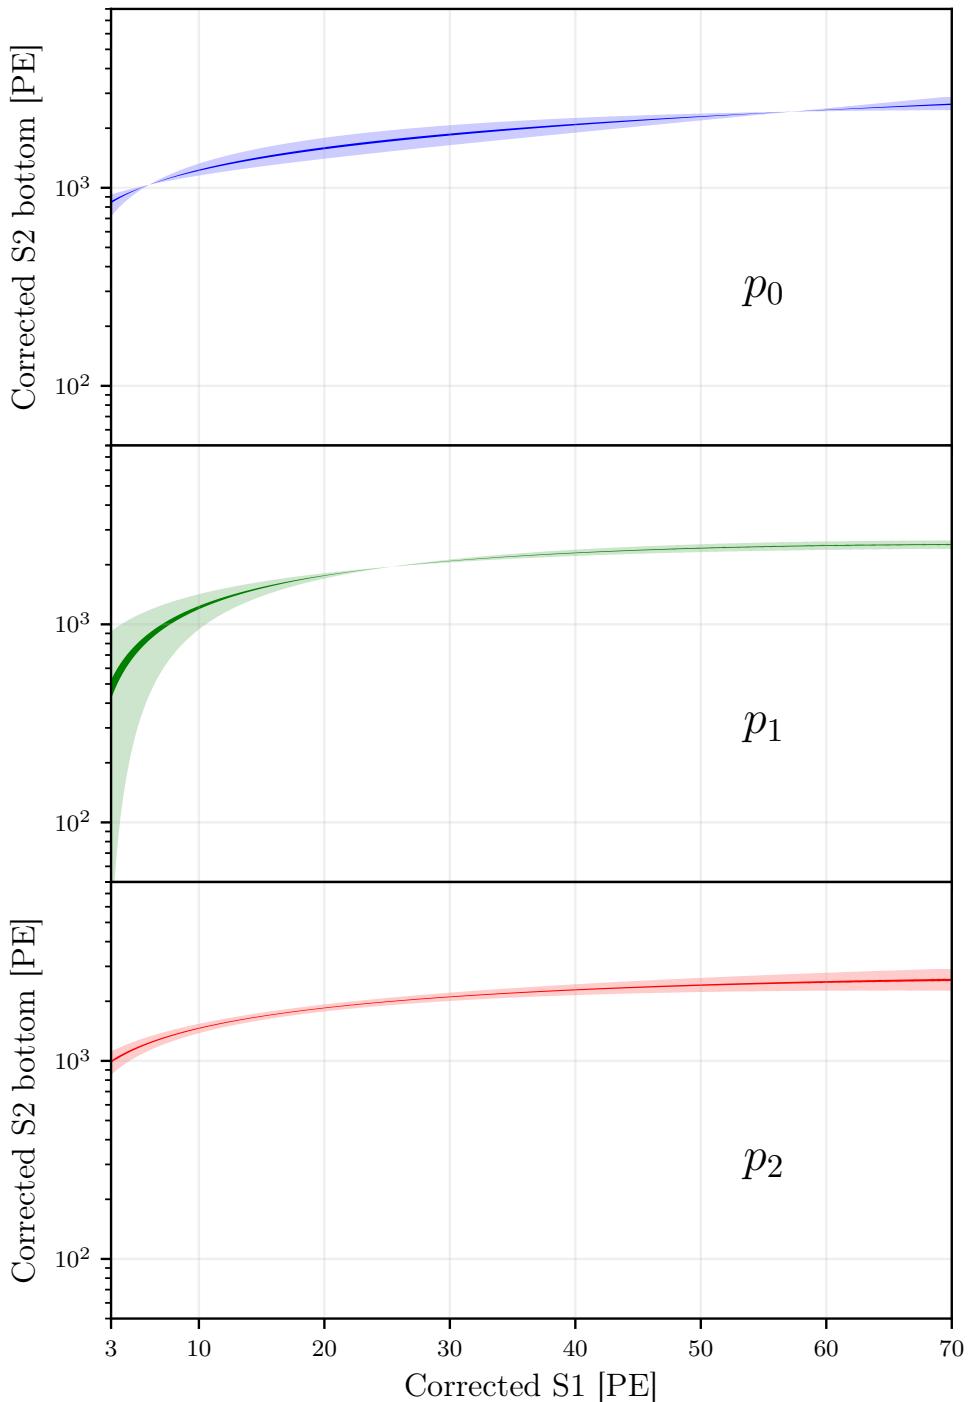


Figure 6.5: Mean ER signal under 1σ (solid colours) and, for illustration, 10σ (light colours) variations of the principal components of the Q_y fit.

posteriors have mean 0 and standard deviation 1:

$$\mathbf{p} = \left(\mathbf{P} \times \frac{\mathbf{v} - \boldsymbol{\mu}_{\text{pre}}}{\sigma_{\text{pre}}} \right) / \sigma_{\text{post}} \quad (6.7)$$

with $\boldsymbol{\mu}_{\text{pre}}$ and σ_{pre} the mean and standard deviation of the components of \mathbf{v} , respectively, \mathbf{P} the PCA transformation matrix, and σ_{post} the standard deviation of the posteriors after PCA (the mean of the \mathbf{p} is zero identically). The three principal components capture 61%, 29% and 10%, respectively, of the variance of \mathbf{v} . We treat the recombination fluctuation parameter σ in isolation, since its correlation with the Q_y polynomial coefficients is weak, and it modifies the ER response in a very different way (broadening the ER band rather than moving its mean).

Figure 6.5 shows the effect of the uncertainty on each of the principal components of the Q_y fit on the mean ER signal. Clearly it is not obvious which principal components (if any) we should propagate to the main analysis. Although p_0 captures more of the uncertainty on the (normalized) original polynomial coefficients than p_1 , the effect of p_1 on searches for especially low-mass WIMPs (which produce signals with small cS1 and cS2_b) will be greater. We will study the impact of each component on a 50 GeV/c² spin-independent WIMP search in section 6.3.3.

6.3 Dark matter limit setting

6.3.1 Impact of outlier events

All but about one in a million events seen by XENON1T are background events that can, for the most part, be easily removed. As described in section 2.4.3, these can consist of many possible event topologies, such as misidentified double scatter events, gas events, and events consisting of pile-up of single electrons. The chance of such an event surviving the gauntlet of data quality cuts is minute, but not zero.

For likelihood ratio testing, outlier events pose a particular challenge: observing an event in a location without signal or background models makes the likelihood (equation 6.1) zero and likelihood ratios not computable. Events that are possible, yet very unlikely under both the signal and background model, also cause problems, especially if $p_{\text{signal}} \gtrsim p_{\text{background}}$. Such an outlier would make the limit significantly higher (i.e. weaker), and in extreme cases could cause a false discovery. The problem is more manageable if $p_{\text{signal}} \ll p_{\text{background}}$: the event will not help fit more signal in, and its effect on the likelihood ratio mostly cancels out. This also requires that the background model is reasonably fixed: if it has nuisance parameters, they have strong enough constraints not to be unduly affected by an outlier.

There are several ways experiments have dealt with this problem:

- Construct a *full background model*. For every conceivable anomalous event source, we would need a model in each dimension involved in the data quality cuts, accurate to deep in the tails of the distribution. This is often impractical or impossible.
- If anomalous events are observed in calibration data, we can assume the rate of these events is *proportional to a regular source of events*, e.g. the rate of low-energy ERs for outliers observed in the ^{220}Rn calibration. We must also assume a distribution of the outliers: if many outliers are observed, this could be estimated from data. With fewer outliers, [183] advocated a conservative approach: assume the proportional anomalous background has the same shape as the signal of interest, and include the background rate as a nuisance parameter in a simultaneous fit of the calibration and science data.
- Unmodelled outliers in the science data can be removed by *post-unblinding cuts*. The danger with these is rejecting real WIMP candidates: looking at enough parameters, most events could be made to look extreme in some parameter.
- *Give up the possibility of a discovery*. If an experiment only quotes upper limits, unmodelled backgrounds always make the limit higher, i.e. more conservative. Often this is combined with a method more robust to unmodelled backgrounds occurring preferentially in a specific region of parameter space (e.g. Yellin's optimum interval method [190]).

None of these methods of dealing with rare outliers is completely satisfactory, which makes it unlikely that a discovery based on a handful of events in a single experiment will be widely accepted even if it is statistically significant based on the experiment's own background model. There is also the converse danger of overcompensating for outliers, inflating the background expectation, causing illusory under-fluctuations of the background and ultimately too strong exclusion limits.

For the analysis described in chapter 5, XENON1T used a mixture of the first two methods. For some ‘anomalous’ event sources, we included a full model: accidental coincidences, and misreconstructed events near the detector walls (‘wall events’). For accidental coincidences, the event distribution in all data quality parameters was inferred from isolated S1s and S2s; for wall events, the background rate was extrapolated into the fiducial volume from the measured rates outside (using a variety of functional fits, which give consistent results). We also included a proportional background model based on the single outlier observed in ^{220}Rn calibration data (see section 5.3.5). For more discussion of these backgrounds, see section 5.3.

Figures 6.6 shows the total resulting background model of XENON1T, with the observed science run events and signal density contours of a 50 GeV/c^2 WIMP overlaid; figure 6.7 shows which background is dominant

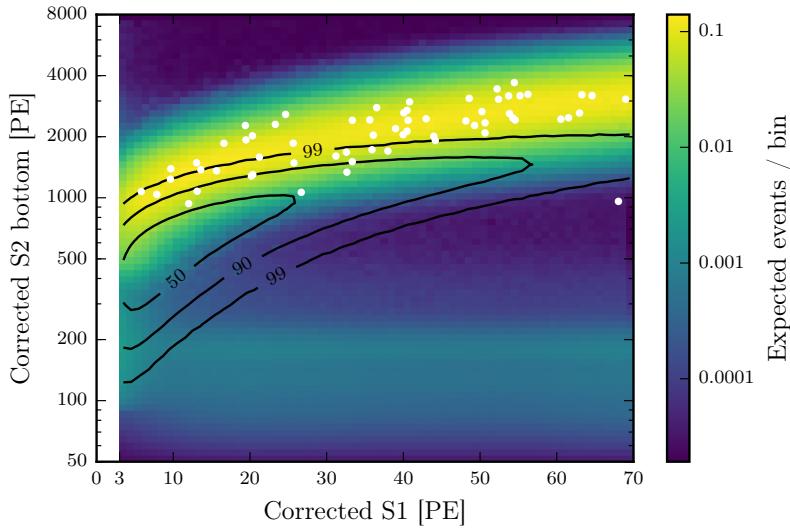


Figure 6.6: Total expected background in XENON1T’s first science run, in events per bin (1 PE bins in cS1, 70 logarithmically spaced bins in cS2_b). Note the logarithmic color scale. Contours show events/bin percentiles for spin-independent elastic nuclear recoil signals from a 50 GeV/c² WIMP. Dots show the observed events in the dark matter search data.

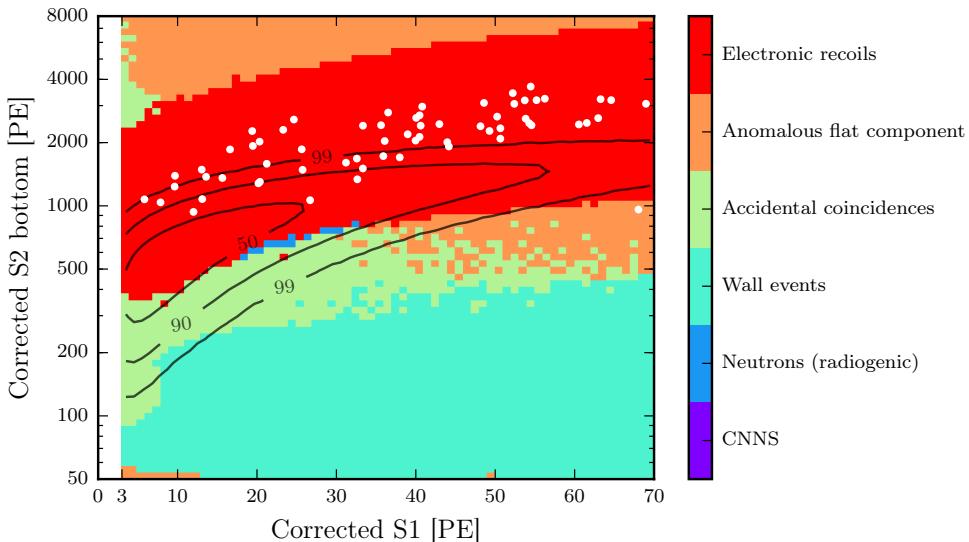


Figure 6.7: Dominant background component in each bin (bins as in figure 6.6). Contours show events/bin percentiles for spin-independent elastic nuclear recoil signals from a 50 GeV/c² WIMP. Dots show the observed events in the search data.

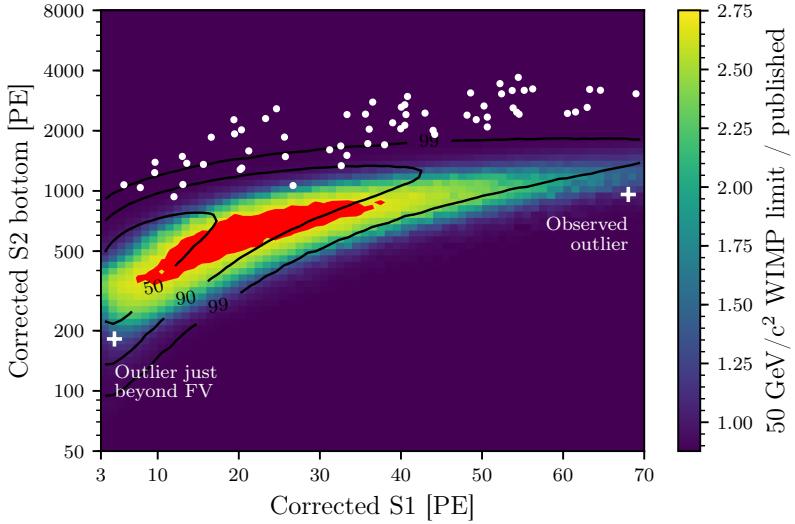


Figure 6.8: Colours: XENON1T first science run spin-independent WIMP-nucleon scattering cross-section limit (divided by the published limit) at 50 GeV/c^2 , if the outlier event at 68 PE (indicated by the cross on the right) would have appeared elsewhere. Red region: if the outlier appeared here, it would cause a WIMP ‘excess’ at $\geq 1\sigma$ significance. Contours: events/bin percentiles for spin-independent elastic nuclear recoil signals from a 50 GeV/c^2 WIMP. Dots: other observed events in the search data. The cross on the left of the figure shows the location of an event that appeared just outside the fiducial volume.

at each point in the search space. Quickly alternating colors indicate two backgrounds are so similar in size small statistical uncertainties dominate. We see there is a region between 500 and 3000 PE $cS2_b$ for $cS1 \gtrsim 30$ PE, which is both outside the main signal region and has an extremely low background expectation, dominated by the poorly understood ‘anomalous background’ (and to a lesser extent accidental coincidences). This is where outlier events would be most dangerous.

To illustrate this point, figure 6.8 shows the impact if the observed outlier event at 68 PE would have appeared elsewhere in the search region. Limits are computed according to the asymptotic likelihood distribution approximation (not with the method of section 6.3.2), and with all nuisance parameters fixed to the best-fit values to the original dataset (with the 68 PE outlier), to limit the computational burden.

We indeed see that in the aforementioned region, an event would have a considerable impact, although WIMPs are very unlikely to appear there. Specifically, at its current location, the outlier causes a 22% increase (i.e. deterioration) of the limit. Without the anomalous background model, this would be much larger at $\sim 70\%$, even though the probability for a WIMP event to

appear so far away from the bulk of its distribution is $\lesssim 0.0013$ for all WIMP masses. For such a clear outlier to have a significant impact on the limit is unsatisfactory. In the future, it would be better to remove regions with very small WIMP signal expectation (and/or where poorly understood backgrounds are dominant) from the search region before unblinding. We already do this for $cS1 > 70$ PE, < 3 PE, and several regions in other parameter spaces (our data quality cuts). This results in a small loss of sensitivity, but a more robust result.

We also indicated an event that appeared a mere 0.6 mm outside the fiducial volume (which was fixed before unblinding, as discussed in section 5.2.5). This event would have had a similar impact on the limit as the observed 68 PE outlier; although it is orders of magnitude more likely to be a WIMP, it has an even more straightforward interpretation as a wall leakage background – which, after all, had an expectation of 0.5 events in the run (see table 5.1). Finally, note that a single event in the red region of figure 6.8 can lead to a 1σ ‘excess’ over background; the maximum detection significance thus obtained is 91%.

6.3.2 Correction for non-asymptotic test statistic distribution

In section 5.4.3 we assumed the profile likelihood ratio test statistic for upper limits (equation 6.3) is distributed according to equation 6.4, which has its 90th percentile at 1.64. This is valid only in the limit of an infinitely large sample [182]. While we saw 63 events in XENON1T’s first science run, most of these are far from the region of interest; the WIMP models we tested predict only a handful expected number of events (see table 5.1). Thus, it is not clear that the asymptotic distribution assumption is accurate; our limit might not correspond to a 90% confidence level, but to a somewhat lower level of confidence.

Further evidence for this is shown in figure 6.9: we exclude some WIMP models that would predict less than one observed event in the run. For comparison, a background-free counting experiment that saw no events would report a 90% confidence level limit of 2.3 events, as indicated in purple. Underfluctuation (or overestimation) of background could also explain why a limit is unusually low (as discussed in section 6.1.1), but does not explain why the sensitivity (median expected limit in case WIMPs do not exist) is also much lower than a background-free experiment.

To compute an adjusted limit, we must determine the distribution (and in particular the 90th percentile) of $t(\sigma)$, for each σ , by repeatedly simulating datasets with WIMP cross-section σ then determining $t(\sigma)$ for these. For simple statistics this is generally known as a *Neyman construction* [191]. The resulting estimated 90% quantile for t is shown in figure 6.10 in solid black for different σ , with the asymptotic quantile in dashed gray. Our published limit corresponds to where the magenta line (the observed values of $t(\sigma)$) crosses

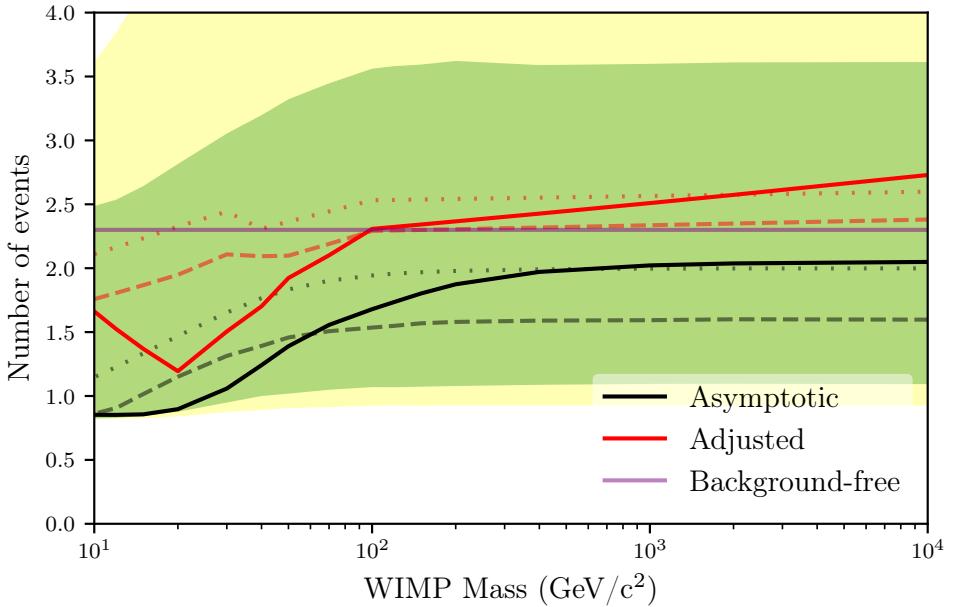


Figure 6.9: XENON1T first science run limit and sensitivity, expressed in expected number of WIMP events in the run. Solid lines show the limit, dashed lines the median sensitivity, and dotted lines the mean sensitivity (in \log_{10} cross-section). Line colour indicates whether the computation uses the asymptotic distribution of the test statistic (black) or the distribution inferred from repeated Monte Carlo calculation (red). The solid purple line shows the limit of a background-free counting experiment. Yellow and green bands show the 2 and 1σ sensitivity percentiles, respectively, according to the asymptotic assumption (as in figure 5.4).

the dashed line; a true 90% confidence limit would be at the intersection with the solid black line instead, at just under 2 expected events.

For this computation, the complete official XENON1T model was fitted to each Monte Carlo trial, using all nuisance parameters. For generating the events, a single model was used, with all nuisance parameters fixed to their nominal values. The procedure could also be done while varying the nuisance parameters (at least those for which a reasonable prior is available) in the generation step. This arguably yields a limit setting method that works reasonably well if the nuisance parameters are close to the estimated values, whereas the current procedure leads to a method that is exactly correct (in the limit of infinite simulation trials) if the nuisance parameters are exactly at the nominal values (and thus probably still reasonable if the true values are near).

We can repeat the construction for all WIMP masses, though at low WIMP masses ($m_\chi \lesssim 30 \text{ GeV}/c^2$) the critical 90% curve (solid black) and observed

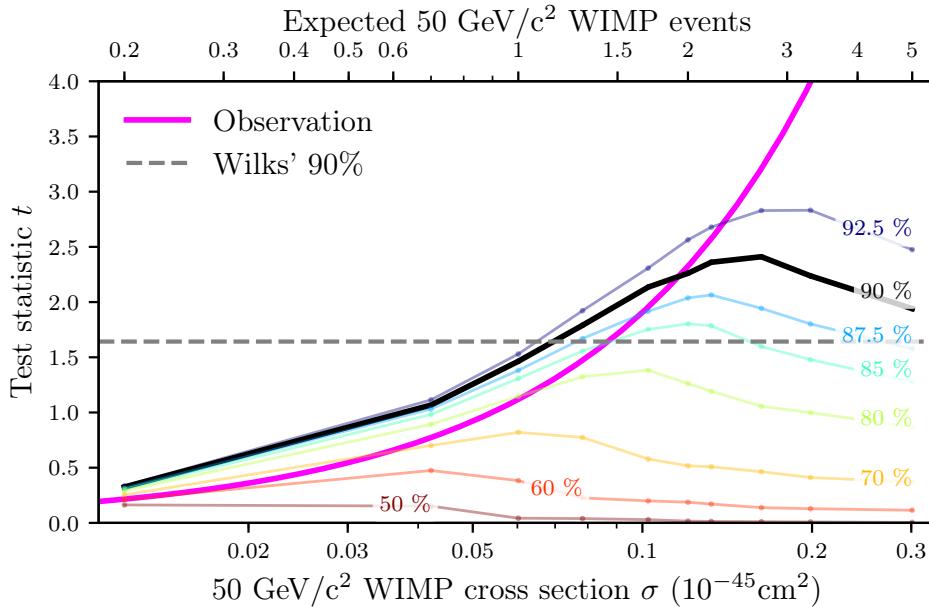


Figure 6.10: Coloured lines: Estimated quantiles of the distribution of $t(\sigma)$ (equation 6.3) for various cross-sections for a $50 \text{ GeV}/c^2$ WIMP. The estimated 90% quantile is indicated in thick black; the dashed line shows the 90% quantile from the asymptotic distribution. Thick magenta line shows $t(\sigma)$ for the observed XENON1T data (chapter 5). Upper x-axis shows the expected number of events for each σ .

curve (magenta) become almost parallel. This is because the region of interest shrinks to a small region just above our experimental thresholds; the likelihood then effectively becomes that of a counting experiment, for which $t(\sigma)$ has only a discrete set of possible outcomes corresponding to the observed number of events in the region. In these cases, computing the intersection of the observed $t(\sigma)$ (magenta in figure 6.10) and the critical 90% curve is very vulnerable to small errors due to limited statistics and interpolation between different σ points; hence, we add a small offset (0.2) to the 90% curve for $m_\chi \leq 30 \text{ GeV}/c^2$, chosen by hand based on the observed non-smoothness in the critical curves of several WIMP masses. This can result in slightly more conservative limits.

The adjusted limit for all m_χ is shown in red in figure 6.9, and the adjusted median sensitivity in dashed red. This limit is about 33% higher (weaker) than the published limit at high masses, and is up to twice as weak at the lowest WIMP masses. The sensitivity is now much closer to the background-free sensitivity of 2.3; the limit is still somewhat lower at low masses, indicating some background underfluctuation.

One might still wonder how our sensitivity could improve on a background-

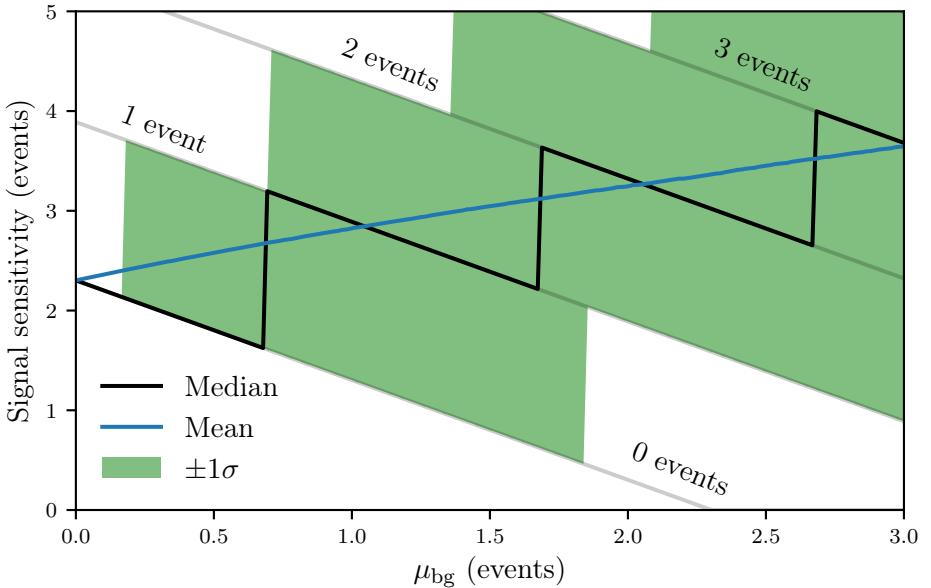


Figure 6.11: Mean (blue) and median (black) sensitivity at 90% confidence level for a counting experiment as a function of its expected background μ_{bg} . Thin grey lines show the possible limits the experiment can obtain: the lowest corresponds to 0 observed events, the next lowest to 1 event, etc. The green band shows the $\pm 1\sigma$ percentile of the expected limits.

free detector, while XENON1T certainly has relevant backgrounds. To this end, examine figure 6.11, which shows the median and mean expected limits (as well as the $\pm 1\sigma$ sensitivity percentiles) for a counting experiment with different number of expected background events μ_{bg} . The median sensitivity of a counting experiment actually *improves* (i.e. decreases) with adding a small background, from 2.3 at $\mu_{\text{bg}} = 0$ to about 1.6 near $\mu_{\text{bg}} = 0.69$. This counterintuitive result is caused by the discreteness of the test statistic, here simply the observed number of events. The limit for 0 observed events (lowest gray line) decreases with μ_{bg} , but for $\mu_{\text{bg}} < 0.69$ the experiment observes 0 events more than half of the time, so the median experimental result is an underfluctuation of the background. Similarly, the median sensitivity for $\mu_{\text{bg}} = 1.67$ is 2.2, again slightly lower (better) than a background-free experiment.

The *mean* sensitivity (blue in figure 6.11) of a counting experiment shows the expected behaviour: a monotonic deterioration (i.e. increase) with increasing μ_{bg} . However, note the mean sensitivity in number of events (or σ) is not equal to the mean sensitivity in $\log \sigma$. Figure 6.9 also shows the mean sensitivity for the adjusted limit setting method in number of events in dotted red. This is clearly worse (i.e. larger) than a background-free detector, except

(slightly) at low WIMP masses. This could be because the ad-hoc adjustment to the critical curve (see above) was insufficient; increasing the number of Monte Carlo trials and evaluating the distribution at more cross-sections would eventually resolve this (at the cost of much more CPU time).

Since the adjustment to the limit (black to red in figure 6.9) is well within the 1σ sensitivity band of the experiment (as stated in chapter 5), simply repeating the experiment would likely result in a larger change in the limit. Still, this shows the asymptotic approximation is inappropriate for short ‘first results’ runs of experiments like XENON1T; for future XENON1T results, a method similar to that described above will be followed (although the longer the run, the better the asymptotic approximation should perform).

6.3.3 Impact of uncertainties and sensitivity projection

XENON1T’s result relies on accurate signal and background models, which require assumptions on detector response parameters. Uncertainties in these are sometimes propagated to the final likelihood (e.g. the rate uncertainty on several backgrounds), sometimes not (e.g. differences in g_2 between calibration and search data), and sometimes it is unclear to what extent our methodology actually accounts for an uncertainty (e.g. various small uncertainties in the ER fitting which are not propagated to the likelihood, as discussed in section 6.2). It would be very useful to know which uncertainties actually matter and to what extent, that is, to know the impact of small parameter variations on our dark matter sensitivity.

We must first choose a sensitivity metric. Although XENON1T currently uses the median expected background-only limit to quantify its sensitivity, we saw in section 6.3.2 this can, counter-intuitively, improve when additional backgrounds are added. Hence, we will look at changes in the mean (in \log_{10} cross-section) expected background-only limit instead. It is useful to compare sensitivity changes to the intrinsic spread of the background-only results (due to stochastic variations in the observed dataset). We will use the quantity:

$$\text{impact} = 100 \times \frac{\mu_{\text{before}} - \mu_{\text{after}}}{s}, \quad (6.8)$$

where μ_{before} and μ_{after} are the mean \log_{10} of background-only limits before and after a change in settings, respectively, and s is the standard deviation of the \log_{10} ’s of background-only limits before the change.

To estimate the impact of changes in detector parameters such as g_2 or the electron lifetime, we need to regenerate the signal and background models. This is straightforward using `blueice` if we use the ER model from section 6.2 and the NR model from the global fit of [110]. The official XENON1T models were produced with a different framework, primarily intended to run on a GPU cluster; there is no evidence the official ER model describes the calibration data better (both fit the limited calibration data well), while the official NR model differs only in slight adjustments to the global fit’s parameters well

within its priors (which, as we'll see below, implies it is not significantly different). We cannot regenerate the small data-derived backgrounds (wall events, accidental coincidences, and anomalous events); though it is likely these would also change if parameters such as g_1 and g_2 are varied.

We will measure the impact of a parameter variation by setting limits on 10^4 toy Monte Carlo datasets created under different conditions. For most parameters, large variations are needed to get changes from the base scenario that are significantly different from the base model at this sample size, and the effect of small variations is estimated by interpolation. The sensitivity as a function of the parameter change is visually inspected to ensure the interpolation is reasonable. We do *not* include any nuisance parameters in the likelihood, except that we do leave the ER background rate free (as in the real likelihood). Including nuisance parameters would greatly increase the computational expense, and also interfere with an important motivation for this study: finding out which nuisance parameters are relevant enough to include in future science runs. Repeating this study with varying nuisance parameters would of course still be valuable, e.g. for probing the robustness of a given likelihood.

Table 6.3 shows the impact of various errors on and changes of the parameters of a hypothetical future 1-ton year science run under the same conditions as the first science run of XENON1T. The first column shows the impact of an overestimate of the reference value, i.e. the true value is actually at the low endpoint of the quoted uncertainty interval. The second column likewise shows the impact of an underestimate (true value at the high endpoint). In both cases, we change the underlying true value (for generating events) without updating the assumed value (for the likelihood inference), i.e. we are considering what would happen if the assumed parameter values are incorrect.

If small variations in the true value cause no change in the limit, it does not imply a parameter is irrelevant: perhaps the sensitivity should have changed, if the true value were used in the inference (as opposed to the original value). Thus, we must compare the effect of changing the true value *and* correctly updating our estimate, as shown in the third and fourth columns of table 6.3.

Below we discuss each parameter or set of parameters in turn: the source of the assumed reference value and uncertainty, the observed impact and its possible consequences.

- At the end of a run, the *exposure* is generally known to a few percent accuracy (see e.g. chapter 5), driven by the uncertainty on the fiducial mass. We use a variation of 1 ton-month here, to have a better reference for comparing other variations. As expected, the sensitivity increases monotonously with exposure, as shown in figure 6.12, though increases yield diminishing returns due to the presence of backgrounds. The sensitivity appears better than predicted at design time, though the difference is approximately the magnitude of the correction described in

| Parameter (unit) | Reference | Impact on sensitivity of | | | |
|------------------------------------|---------------------------------|--------------------------|---------------|----------|----------|
| | | Overestimate | Underestimate | Decrease | Increase |
| Exposure (tonne days) | 365.25 ± 30.00 | -9 | 7 | 10 | -7 |
| Electron lifetime (μ s) | 452 ± 10 | -20 | 40 | 4 | -2 |
| g_1 (PE/keV) | 0.144 ± 0.006 | -36 | 38 | 3 | -1 |
| g_2 (PE/keV) | 11.52 ± 0.46 | 112 | -77 | 3 | -2 |
| Single-PE response width (PE) | 0.42 ± 0.10 | -4 | 7 | -2 | 3 |
| ER p_0 of Q_y | 0 ± 1 | 9 | -7 | 1 | 0 |
| ER p_1 of Q_y | 0 ± 1 | 37 | -32 | 11 | -10 |
| ER p_2 of Q_y | 0 ± 1 | 21 | -21 | 3 | -4 |
| ER σ : fluctuation | 0.045 ± 0.003 | -19 | 31 | -5 | 6 |
| NR α | $1.24^{+0.079}_{-0.073}$ | -1 | 2 | 9 | -8 |
| NR ζ | $0.0472^{+0.0088}_{-0.0073}$ | 0 | -1 | -4 | 7 |
| NR β | 239 ± 19.2 | 0 | 0 | 2 | -1 |
| NR γ | $0.01385^{+0.00058}_{-0.00073}$ | -1 | 1 | 7 | -5 |
| NR δ | $0.062^{+0.0056}_{-0.0064}$ | 0 | -1 | -5 | 3 |
| NR Lindhard k | $0.1394^{+0.0032}_{-0.0026}$ | 0 | 1 | -1 | -1 |
| NR η | $3.3^{+5.3}_{-0.7}$ | 0 | -1 | -1 | 7 |
| NR λ | $1.14^{+0.45}_{-0.09}$ | -1 | 0 | 3 | -1 |
| NR σ : fluctuation | 0.02 ± 0.02 | 0 | -1 | 0 | 3 |
| ER background (events/day) | 1.8 ± 0.23 | -4 | 4 | -6 | 3 |
| Neutron background (events/day) | 0.0017 ± 0.00034 | -4 | 4 | -2 | 2 |
| CNNS background (events/day) | $2.5e-05^{+0.00056}_{-2.3e-05}$ | 0 | 6 | 0 | 0 |
| Anomalous background (events/day) | $0.0026^{+0.00024}_{-0.00078}$ | -3 | 1 | 0 | 0 |
| Accidental background (events/day) | 0.0064 ± 0.00029 | -2 | 1 | -1 | 0 |
| Wall background (events/day) | 0.015 ± 0.0091 | 0 | 1 | 0 | 0 |

Table 6.3: (facing page) Impact on the mean sensitivity (expected background-only limit), as quantified by equation 6.8, of under- and overestimations of parameter values (first two columns) and changes in the parameter values that are properly accounted for (third and fourth columns). The quoted uncertainty on the reference value sets the amplitude of the variation. Omitted values imply no significant variation was found even with amplitudes far exceeding the quoted uncertainty. Impact values with absolute value greater than 30 are highlighted in bold font.

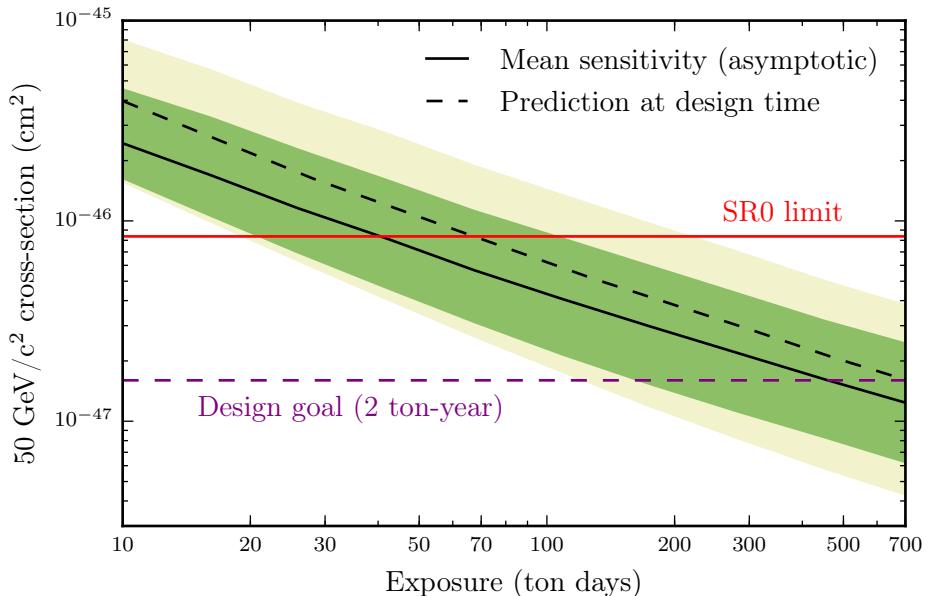


Figure 6.12: XENON1T sensitivity to $50 \text{ GeV}/c^2$ WIMPs as a function of exposure, assuming no change in the conditions of the first science run (SR0). Coloured bands show the 1 and 2σ sensitivity percentiles, as usual. The assumption of an asymptotic test statistic distribution is made throughout; see section 6.3.2 for more details. The limit of the first science run is indicated in red. Dashed lines indicate a prediction of the sensitivity made before XENON1T construction based on the parameters in [68].

section 6.3.2, which has not been applied here (computing the critical 90% curve for all parameter variations considered would incur tremendous computational expense).

- The *electron lifetime* uncertainty has been set at 2.4%, accounting for statistical uncertainties as well as systematic differences between estimation methods [192]. The electron lifetime varied considerably more during the science run, but this variation was accounted for in our models. For this study the electron lifetime is taken as constant. Increasing the lifetime leads to very slightly better limits, since this reduces the fluctuation of especially small S2s. However, an under- or overestimate of the electron lifetime affects the sensitivity significantly: it causes us to use a wrong spatial correction for the S2s, causing larger fluctuations and shifts of the mean $cS2_b$. If the estimation error is consistent over the run, the background model fits would correct some of this (at the cost of incorrect ER model parameters), but it cannot do so for systematic differences between ^{220}Rn calibration data and science data.
- We assume 4% uncertainties on the *signal scales* g_1 and g_2 , since this is approximately the variations in single electron gain we saw in section 3.4.1. Unlike electron lifetime, signal scale variations are currently not considered during model production. Changes in the signal scales have little impact, but under- and overestimates of especially g_2 have a very large impact, much larger than any other variation we considered. Such errors shift the models in $cS1$ and $cS2_b$; the effect of the latter is greatest, as signals are much less spread out in $cS2_b$ than in $cS1$ (since the S1 signal is made up of far fewer photons). In section 3.4.1 we saw variations in single electron area that coincided with the start and end of ^{220}Rn calibration, so a relevant systematic difference between calibration and real data is plausible. Clearly this deserves careful attention in the analysis of future XENON1T runs.
- The *single-photoelectron response* is assumed to be a Gaussian with a mean of 1 PE and a spread equal to the median of the assumed gain spreads of XENON1T’s PMTs divided by their median measured gain. Changes in the single-PE response mean will have a perfectly correlated impact on g_1 and g_2 . To test the impact of changes in the single-PE response width, we vary the gain spread by 25%, which should be easily visible. The impact is extremely small.
- For the *ER response*, we vary the principal components (p_0, p_1, p_2) of the Q_y fit derived in section 6.2 (which have a standard Gaussian uncertainty by construction), as well as the recombination fluctuation parameter σ within its uncertainty from the ER fit. We see that under- and overestimates of p_1 and to a lesser extent p_2 cause significant effects;

changes in the true value have a relatively small impact (though larger than any other change of value considered). Although p_0 was the first principal component (indicating it changed the underlying polynomial coefficients most), its effect on the dark matter search is small. Since the ER fit was statistically limited, taking more ^{220}Rn calibration data should reduce the impact of uncertainties on the ER model. As of December 2017, XENON1T has already taken ~ 7 times as much ^{220}Rn calibration data than was used for its first results (chapter 5).

- The *NR fit parameters* are taken from the global fit of [110] (discussed in section 2.3.2), and are varied individually within their published confidence intervals. No NR response is assumed below 1 keV. Again, the model parameters likely have some correlation, but this information is not available in [110] and not used by XENON1T. Notice first that over- and underestimates of any NR parameters leave the expected result essentially unchanged, since background-only datasets contain very few NR events. Of course the NR model is still relevant, as it affects the signal model, as we can see from the third and fourth columns. Still, the impact of changes in the NR parameters is small. This implies the global fit’s constraints are tight enough that additional constraints (e.g. from XENON1T’s $^{241}\text{AmBe}$ calibration discussed in chapter 5) will have a minor impact, unless they cause major deviations from the global fit’s priors. Requesting the parameter correlation information from [110] would seem to have a higher priority than taking more $^{241}\text{AmBe}$ calibration data. In [110] a sub-percent Fano factor on the recombination probability was added to the model; here we tested a simpler model, where the probability for a detected quantum to be an ion is fluctuating by σ (as described for the ER model, see section 6.2). The impact of even a several-percent fluctuation is found to be negligible.
- The *background rates* were varied within the intervals reported in table 5.1, with the exception of the CNNS background, whose rate we vary over a much larger interval that covers the factor ~ 20 difference between the rate expected under the NR model used here (see above) and the rate quoted in table 5.1. We see the impact on the $50 \text{ GeV}/c^2$ WIMP sensitivity of all variations considered is very small.

6.4 Recommendations for future runs

Based on the studies in this chapter, we can make a few recommendations for the long science run of XENON1T, and (to some extent) for future dark matter detectors:

- The assumption of an asymptotically distributed likelihood statistic should be abandoned in favour of deriving the distribution by repeated

Monte Carlo simulation, especially for short ‘first results’ runs. This is challenging for low WIMP masses with tiny regions of interest; for these, approximate methods could be investigated.

- Median sensitivity is a poor measure for comparing rare-event search experiments, as experiments with substantial backgrounds can have a better median sensitivity than a similar background-free experiment. The mean sensitivity (in log cross-section) can be an alternative.
- Regions of parameter space where WIMP signals are very unlikely should not be included in the limit setting, especially if the backgrounds in these regions is poorly understood (e.g. if it is dominated by ‘anomalous’ backgrounds assumed proportional to the ER rate). The risk of mistaken results induced by outlier events in these regions is substantial, while the sensitivity gain of including them is minimal. Naturally the decision which regions to exclude must be taken before unblinding, and events in such regions are still useful for other purposes (e.g. estimating the reliability of the background model).
- Unaccounted variations of g_1 and g_2 (and to a lesser extent electron lifetime) during the science run, or between the science run and calibration periods, can be of great consequence. We should make every effort to quantify these, and/or include appropriate uncertainties in the background model (as in, for example, [193]).
- ER calibration data (^{220}Rn) statistics should be substantially increased. For NR, the global fit of [110] is sufficiently accurate for our purposes; verifying this with independent small-scale detectors may be more practical than additional in-situ measurements in XENON1T.
- If uncertainties on the ER fit are substantial, it may not suffice to propagate only the component of the Q_y uncertainty that accounts for most variance in the coefficients of a polynomial Q_y model. Different principal components may be more consequential for the final result, which could even depend on which WIMP model is tested. We should examine the possibility of a simultaneous ER calibration and science data fit. This is likely feasible only if, instead of fitting a many-parameter model (with the added risk of overfitting) we study which specific uncertainties matter for the final result, then ignore the rest.

Chapter 7

Expanding the dark matter search

XENON1T is currently extending the hunt for dark matter by taking an unprecedented ~ 1 ton year of exposure at a record-low background level. We discussed the sensitivity prospects of this run in section 6.3.3, and made recommendations for its analysis in section 6.4. This chapter describes a parallel track of improvements to the XENON1T results: using more general dark matter models and search techniques to extract complementary dark matter constraints from the same data. We will use the data from XENON1T’s first science run (see chapter 5). Our aim is to demonstrate the viability and potential of these techniques; the results here should be considered preliminary and/or proofs of concept.

7.1 Introduction

The analysis of chapter 5 considered elastic spin-independent nuclear recoils detected by S1 and S2 signals. Below we consider the following alternative searches:

- Section 7.2 considers *spin-dependent scattering* (see section 2.1.3). This is relevant for models such as Majorana fermions coupling to the Z (see equation 1.1 and figure 1.5).
- Section 7.3 searches for WIMPs through the *bremsstrahlung* signature advanced by [96] and discussed in section 2.2.3. This extends the analysis to light WIMPs ($0.5 - 2$ GeV/c 2), which are kinematically unable to produce detectable elastic nuclear recoils.
- Section 7.4 performs a spin-independent WIMP search in the *S2-only channel*, disregarding the scintillation-light signal S1 (except perhaps as a veto). This lowers the effective detection threshold, since LXe’s

low-energy charge yield far exceeds its light yield (see section 2.3.2). While the lower threshold increases the expected rate of all WIMPs, this is relevant only for a narrow range of WIMP masses (3-7 GeV/c^2).

- Finally, section 7.5 searches for *bremsstrahlung signals in the S2-only channel*. The S1 threshold is far more limiting for an ER search than an NR search (since ERs produce much smaller S1s at the same S2), so releasing it yields correspondingly greater benefits. We extend XENON1T’s spin-independent WIMP-nucleon sensitivity to 100 MeV/c^2 , where no other DM-scattering limit currently exists.

The first three extensions listed above are independent choices that can be combined – for example, one could search for spin-dependent scattering in the S2-only channel. The analysis of section 7.5 is but one of these combinations.

For the S1S2 spin-dependent and bremsstrahlung searches, we could repeat the analysis of chapter 5 (hereafter called the ‘official analysis’) with new dark matter signal models. However, reproducing the exact software setup used by the collaboration to produce the signal and background models in the official analysis is nontrivial. Moreover, as we saw in section 6.3.2, the official analysis’ assumption of an asymptotically distributed test statistic leads to too strong limits. We also made other suggestions for improvement in section 6.4 which we now have the opportunity to try out.

Therefore, we make several changes to the official analysis, discussed in detail in section 7.1.1 below. Our results should consequently be considered preliminary. An even stronger caveat applies to the S2-only analysis: this must be constructed from the ground up, as the S1 signal is used for vertical fiducialization and several data quality cuts in the official analysis. This analysis should be considered a proof of concept.

7.1.1 Differences with the official analysis

For the spin-dependent and bremsstrahlung search, we make the following changes to the official analysis:

- **Signal yield models.** We use the ER model of section 6.2 and the NR model of [110]. Our ER model shows a good fit to the ^{220}Rn calibration (see table 6.2), as does the official model. The official NR model adjusted the parameters of [110] within their uncertainties based on the XENON1T $^{241}\text{AmBe}$ calibration. However, as shown in figure 2.6, these uncertainties are small, and the study of section 6.3.3 confirmed they have no significant influence on the WIMP search sensitivity (at least for spin-independent 50 GeV/c^2 WIMPs). As in the main analysis, we assume no detectable signals from WIMP-nucleus scattering with less than 1 keV deposited energy. For bremsstrahlung photons we make the same assumption; although charge yield has been demonstrated to lower energies than this

(especially for ERs), light yield is required to make S1s, and sub-keV light yield has not been demonstrated for either recoil type in LXe.

- **Nuisance parameters.** We include a different set of systematic uncertainties on the signal and background models as nuisance parameters. From the study of section 6.3.3, we know the most consequential parameters for the $50 \text{ GeV}/c^2$ WIMP search are the 4% systematic uncertainty on g_2 (with a uniform prior between -4% and $+4\%$) and the uncertainty on the ER Q_y principal component p_1 . Ideally, one would instead repeat the study of section 6.3.3 for several WIMP masses for each WIMP detection mechanism, and if necessary adjust the choice of nuisance parameters accordingly. We leave the ER background rate unconstrained (as in the official analysis).
- **Asymptotic approximation.** Section 6.3.2 showed the official analysis’ results for low WIMP masses are too strong due to an inaccurate assumption on the likelihood ratio test statistic distribution. Repeating this study for every analysis considered here would take considerable computational expense. Instead, we apply an approximate correction. If the expected number of WIMP events N_0 at the cross-section corresponding to the mean (in \log_{10} cross-section) background-only 90% C.L. limit (under the nominal signal model) is less than 2.3 events (the result for a background-free detector), we scale the result (and sensitivity band) by $2.3/N_0$. This is close to the effect of the full correction for the case shown in figure 6.10.

A more difficult choice concerns the anomalous flat background of section 5.3.5 and the outlier event at $cS1 \approx 68 \text{ PE}$ in the search data. Our two assumptions about the anomalous component – flat in our search space and proportional to the ER rate – are both tenuous: the first has no physical basis, and the second is challenged by ^{220}Rn calibration data from XENON1T taken after the first science run was completed, in which no anomalous population is observed [194]. This suggests we remove the background: although this is a post-unblinding modification, it results in more conservative limits.

However, there clearly is some source of events in the ^{220}Rn calibration data (at least one event, see figure 6.3) beyond what our electronic recoil model can accommodate. Attempts to explain this using accidental coincidences (section 5.3.3) or wall leakage (section 5.3.4) were unsuccessful (at least when the data was first analyzed, a reanalysis is ongoing as of this writing [194]).

Removing the flat background would also diminish the robustness of the analysis. As discussed in section 6.3.1, likelihood-ratio tests are vulnerable to events in regions where the signal and background expectation are both small (compared to unquantified uncertainties). The flat background ensures such events will be seen as part of a small background and have little effect; without it, one outlier can greatly skew the result. The outlier at $cS1 \approx 68 \text{ PE}$ in the

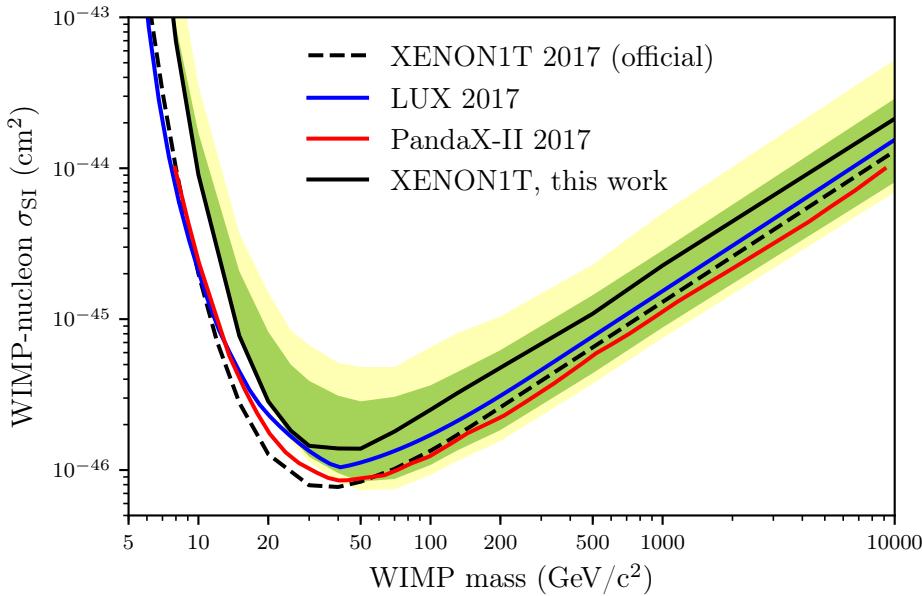


Figure 7.1: Solid black: Limit on spin-independent WIMP-nucleon scattering from XENON1T’s first science run according to the analysis described in this chapter. Green and yellow bands show the 1 and 2σ sensitivity quantiles, respectively. The limit from chapter 5 is shown in dotted black. Limits from PandaX-II [195] (red) and LUX [196] (blue) are shown for comparison.

science data would have worsened the limit by up to $\sim 70\%$ without the flat background, though this event is unlikely to be a WIMP ($p \lesssim 10^{-3}$). It seems more plausible that this event has the same origin as the anomalous event(s) in the ^{220}Rn data – exactly what the likelihood *with* the flat background does.

Hence, we will *keep* the anomalous background for the extended analyses in this chapter. The only other option would be to remove it while taking additional measures against outliers (e.g. clipping the signal model beyond its 99th percentile). This would arguably have been a better choice before unblinding; now, it effectively cuts the observed outlier at a minute sensitivity penalty. This seems a dubious post-unblinding intervention, which would result in a more aggressive result.

7.1.2 Comparison of spin-independent results

Figure 7.1 shows the upper limits on elastic spin-independent WIMP-nucleon scattering from the first science run of XENON1T which result from our revised analysis. The limit is higher (worse) than the official analysis at all masses due to the correction of section 6.3.2. The changes to the signal and background models have no substantial effect on the limit; even removing all

nuisance parameters from the likelihood does not considerably change the result. However, the sensitivity band at low masses is narrower than in the official analysis. In general, searches for signals near detector thresholds are very sensitive to assumptions on systematic uncertainties. As we based our choice of nuisance parameters on studying 50 GeV/c² WIMP searches in a hypothetical 1-year run, it is possible we omitted a parameter that is relevant for the lowest-mass WIMP searches.

7.2 Spin-dependent search

7.2.1 Signal model

As discussed in section 2.1.3, spin-dependent WIMP-nucleon scattering uses a different form factor (eq. 2.16) and a different conversion between WIMP-nucleus and WIMP-nucleon cross-section (eq. 2.19). We must also sum the contribution of the different isotopes in the target; for xenon, the two relevant isotopes are listed in table 7.1. The differential WIMP-nucleus cross-section becomes:

$$\frac{d\sigma_{\text{SD}}}{dE_R} = \sum_A f_A \frac{4\pi\mu_N^2}{3\mu_n^2(2J+1)} S_A(E_R) \sigma_n / E_{\text{max}}, \quad (7.1)$$

with the sum running over each isotope in table 7.1, σ_n the WIMP-nucleon cross-section, E_{max} the kinematic recoil energy maximum (eq. 2.11, note this depends on the WIMP velocity), f_A the fractional abundance of each isotope, J the nuclear spin, S_A the structure function, μ_N and μ_n the WIMP-nucleus and WIMP-nucleon reduced mass, respectively. Note f_A , J , S_A , μ_N , and E_{max} are different for each isotope, though the latter two only because each isotope has a slightly different mass.

The structure functions S_A can be computed with chiral effective field theory models of the nucleus. We use the results from [84], whose model currently has world-leading accuracy in reproducing the measured nuclear excitation spectra. For perspective, this means the ordering of states is approximately reproduced, but factor 2 discrepancies between predicted and observed excitation energies remain (see figure 1 of [84]). The same structure functions were used for the recent LUX results [85, 86].

While experiments usually quote one result for spin-independent scattering (corresponding to equal WIMP-proton and WIMP-neutron coupling), they

| Isotope | Natural abundance (f_A) | Nuclear spin (J_A) |
|-------------------|-----------------------------|------------------------|
| ¹²⁹ Xe | 26.4 % | 1/2 |
| ¹³¹ Xe | 21.2 % | 3/2 |

Table 7.1: Stable xenon isotopes that carry nuclear spin in their ground state (due to their odd number of neutrons).

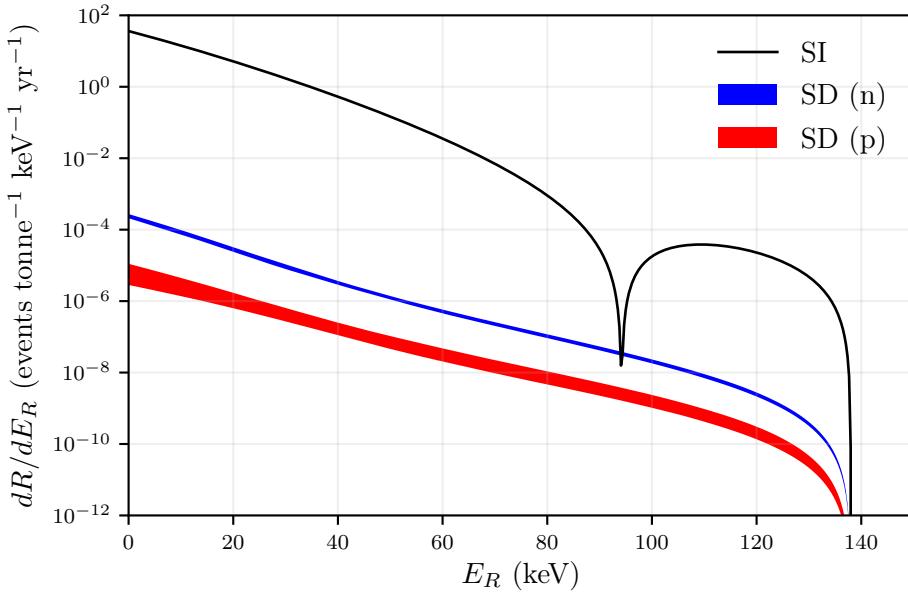


Figure 7.2: Differential rate of elastic WIMP-nucleus scattering for a $m_\chi = 50 \text{ GeV}/c^2$ WIMP with WIMP-nucleon scattering cross-section $\sigma_0 = 10^{-45} \text{ cm}^2$ versus recoil energy E_R , under different WIMP-nucleon coupling assumptions: spin-independent (SI, black), spin-dependent neutron-primary (SD,n) and spin-dependent proton-primary (SD,p). Bands on the spin-dependent rates indicate the uncertainty from theoretical uncertainties on the structure function as quoted in [84]. Nuclear physics uncertainties on the SI scattering rate have negligible impact on WIMP searches [81] and not shown.

quote two results for spin-dependent scattering, corresponding (at first order) to neutron-only ($\sigma_{\text{SD},n}$) and proton-only coupling ($\sigma_{\text{SD},p}$), calculated using different structure functions. This is because the sensitivity to $\sigma_{\text{SD},n}$ and $\sigma_{\text{SD},p}$ varies greatly between target elements. Xenon has an even number (54) of protons, which thus have paired spins in their ground state in every xenon isotope, and participate very weakly in spin-dependent scattering. Xenon experiments are therefore far less sensitive to $\sigma_{\text{SD},p}$ than $\sigma_{\text{SD},n}$; its sensitivity to $\sigma_{\text{SD},p}$ comes primarily from higher-order interactions (in which the WIMP interacts with a pair of nuclei). The case is reversed for experiments using elements with an odd proton count such as iodine. Very few elements have stable isotopes with odd numbers of both protons and neutrons, so only experiments using a combination of detector materials have good sensitivity to $\sigma_{\text{SD},n}$ and $\sigma_{\text{SD},p}$ both.

Figure 7.2 compares the differential rate of elastic WIMP-nucleus scattering (equation 2.2) in xenon for spin-independent and spin-dependent scattering. Note the rate for spin-independent scattering is highest (due to the A^2 coherent

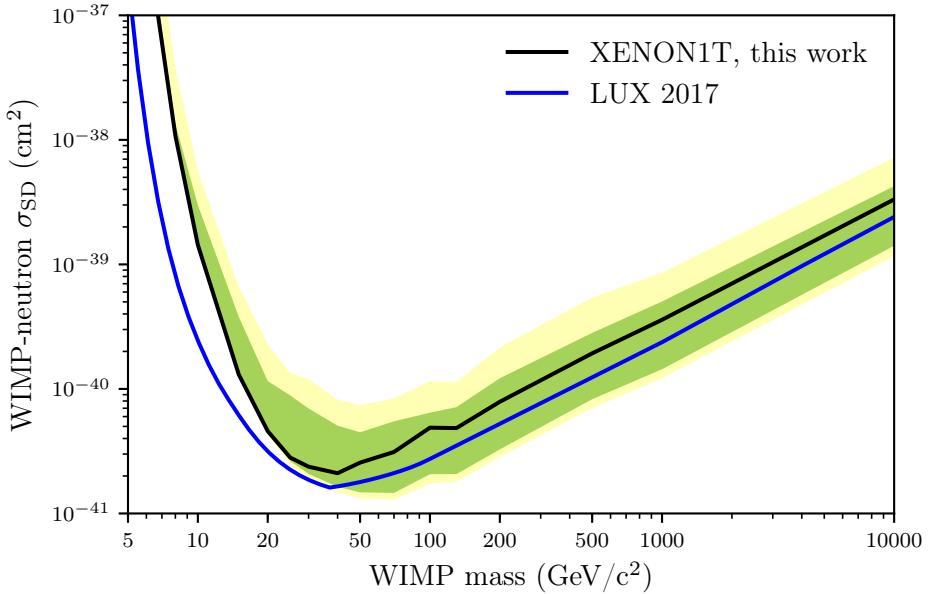


Figure 7.3: Solid black: Limit on spin-dependent WIMP-neutron scattering from XENON1T’s first science run (black). The result from LUX [86] is shown for comparison in blue. Coloured bands are as in figure 7.1.

enhancement), followed by spin-dependent neutron scattering and finally spin-dependent proton scattering. All rates drop off at the same maximum energy, set by kinematics and the maximum observed WIMP velocity in the standard halo model (see section 2.1.2). Coherent scattering off all nucleons causes a diffraction minimum in the Helm form factor for SI-scattering, which is absent in SD-scattering. Except for the reduction in total rate, the recoil energy spectra of all three scattering modes in figure 7.2 are very similar in the region of interest (< 40 keV), so our SD results should be approximately upward translations of the SI limit in figure 7.1.

7.2.2 Results

We proceed with the analysis as described in section 7.1, adding one additional nuisance parameter parametrizing the theoretical uncertainty on the differential scattering rate (due to the nuclear models that provide the structure functions) as shown in figure 7.2 (with a uniform prior). Figures 7.3 and 7.4 show the resulting 90% C.L. upper limits on $\sigma_{SD,n}$ and $\sigma_{SD,p}$, respectively. As expected, the limit curves have the same shape as in the spin-independent case (figure 7.1), and are positioned similarly to (i.e. slightly weaker than) the results from LUX. PICO-60 provides the leading WIMP-proton scattering result (due to using target nuclei with odd proton count).

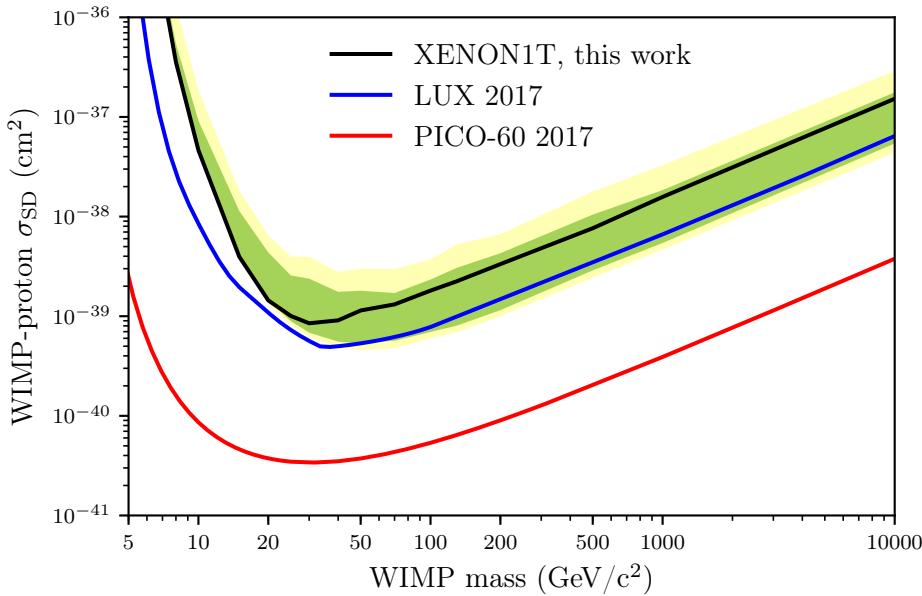


Figure 7.4: Solid black: Limit on spin-dependent WIMP-proton scattering from XENON1T’s first science run. Limits from LUX [86] and PICO-60 [197] (red) are shown for comparison. Coloured bands are as in figure 7.1.

7.3 Bremsstrahlung search

7.3.1 Signal model

As discussed in section 2.2.3, [96] pointed out that any WIMP-nucleus collision can cause emission of a Bremsstrahlung photon, which is detectable through an electronic recoil signal. This gives xenon detectors sensitivity to signals from sub-GeV WIMPs (see figure 2.3).

Figure 7.5 shows the differential rate of bremsstrahlung photon production from WIMP-nucleus scattering (equation 2.22) for spin-independent and spin-dependent WIMP-nucleon scattering. The valleys in the rate below 1 keV are due to the X-ray form factor in equation 2.23, which parametrizes the effects of atomic electrons on the bremsstrahlung emission. We follow [97] in using the form factor from [198].

As mentioned in section 7.1.1, we assume no response for bremsstrahlung photons with $\omega < 1 \text{ keV}$. This is consistent with the convention of LXe experiments (discussed in section 2.3.2) to not extrapolate signals yields below energies probed by calibrations, even if such extrapolations are better motivated theoretically than a hard cutoff. While charge yield for electronic recoils has been demonstrated down to 180 eV [111], light yield has only been measured down to $\sim 1 \text{ keV}$ [112, 113] (it is difficult to quote a precise value,

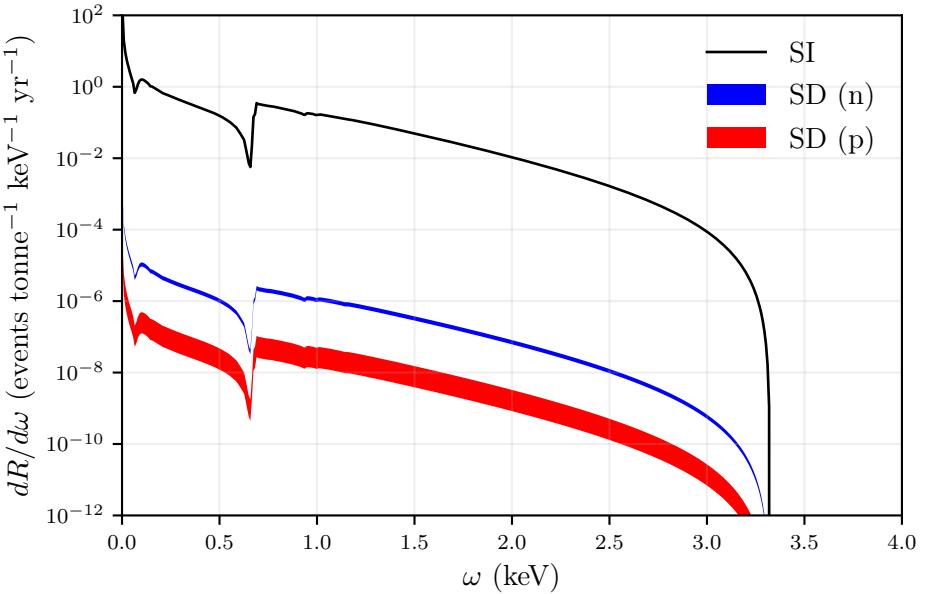


Figure 7.5: Differential rate of Bremsstrahlung emission from WIMP-nucleus scattering for a $m_\chi = 1 \text{ GeV}/c^2$ WIMP with WIMP-nucleon scattering cross-section $\sigma_0 = 10^{-38} \text{ cm}^2$ versus the Bremsstrahlung photon energy ω , under different WIMP-nucleon coupling assumptions (as in figure 7.2). The spectrum agrees with that shown in figure 1 of [96] (after multiplying by the ratio of the shown cross sections).

since these are fits to continuous spectra). The 1 keV cutoff restricts our sensitivity to WIMPs with $m_\chi > 0.3 \text{ GeV}/c^2$ (see figure 2.3).

Figure 7.6 compares the distribution of bremsstrahlung events for a $m_\chi = 1 \text{ GeV}/c^2$ WIMP in the XENON1T first science run in the (cS1, cS2_b) search space with the background model of the search (see also figure 6.6) and the observed events in the search data (see also figure 5.2c). Notice the signal region overlaps with the electronic recoil background, since the signals are themselves electronic recoils (caused by bremsstrahlung photons).

7.3.2 Results

The resulting limits on spin-independent WIMP-nucleon scattering from Bremsstrahlung emission are shown in figure 7.7. We see the limit is not competitive with existing constraints; in particular, CRESST-III [199] and a prototype gram-scale calorimeter for CRESST [201] have set much stronger constraints on low-mass WIMPs. A stronger limit has also been derived by [97] in a bremsstrahlung search using data from LUX; however, [97] cuts off the WIMP signal below 0.19 keV instead of below 1 keV, and thus relies

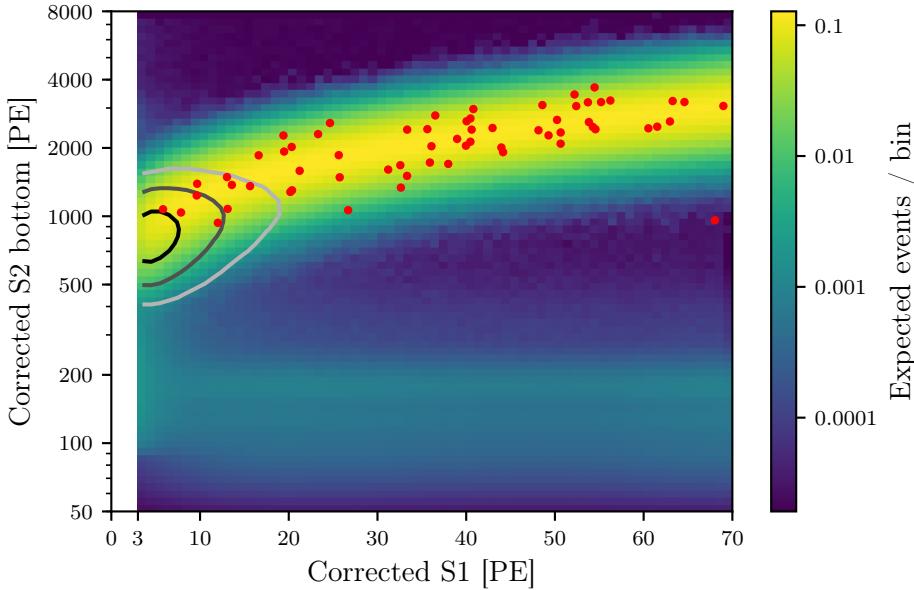


Figure 7.6: Contours: signal model of bremsstrahlung photons from a $m_\chi = 1 \text{ GeV}/c^2$ WIMP, from dark to light: 50, 90, and 99th percentiles of density. Coloured background: total expected background in XENON1T’s first science run, similar to figure 6.6. Red dots: observed events in the dark matter search data.

on extrapolated light yields, resulting in stronger limits. LXe experiments conventionally avoid such extrapolations (see section 2.3.2), although [97] shows the result is robust to using one of a variety of physically motivated extrapolations. This shows the sensitivity of xenon detectors to WIMPs below $0.5 \text{ GeV}/c^2$ via bremsstrahlung could be much greater after a demonstration of light yield from ERs in LXe at sub-keV energies (e.g. using dedicated setups).

The electronic recoil background rate is completely unconstrained in our analysis – as remarked in section 5.3.1, the dark matter search data itself provides the strongest and most reliable constraint on this. For the $m_\chi = 1 \text{ GeV}/c^2$ WIMP considered in figures 7.5 and 7.6, the signals are of sufficiently low energy that the ER background can still be constrained from the higher-energy part of the search space. For higher-mass WIMPs, bremsstrahlung and ER background will be more similar in our search space, so including additional constraints on our ER background (e.g. from higher-energy search data) might increase the sensitivity. As a test, we fixed the ER background arbitrarily at the best-fit level from section 5.4.2, but found no appreciable improvement in the limit at any mass considered in figure 7.7. At even higher masses, Bremsstrahlung constraints are irrelevant compared to constraints from elastic nuclear recoil searches (which reach $\mathcal{O}(10^{-46} \text{cm}^2)$, see figure 7.1).

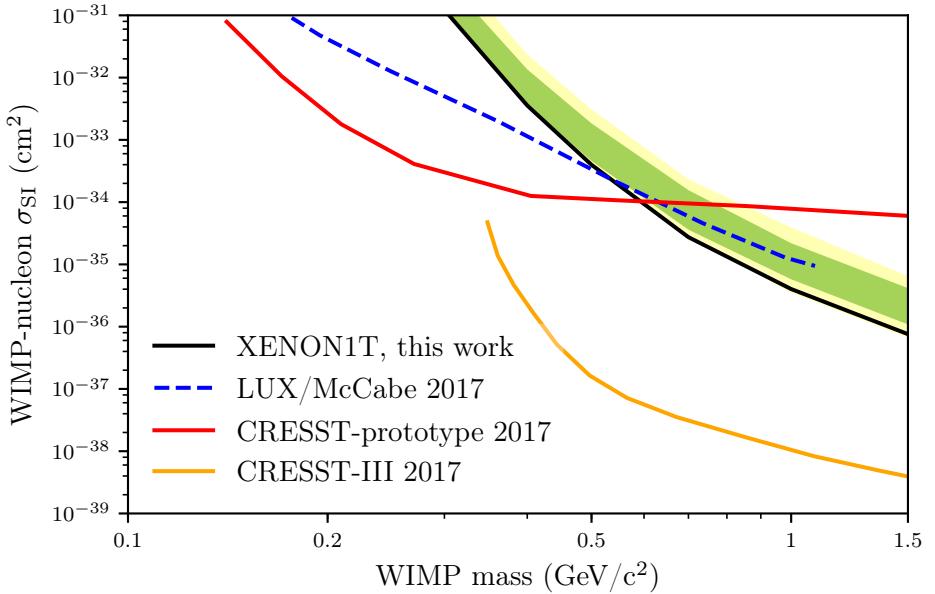


Figure 7.7: Solid black: Limits on the spin-independent WIMP-nucleon scattering cross-section from a search for bremsstrahlung photon emission in XENON1T’s first science run. For comparison, we show results from a bremsstrahlung search [97] using data from LUX [196] without the 1 keV cutoff, and results from elastic NR searches in CRESST-III [199] (orange) and a CRESST prototype detector [200] (red).

7.4 S2-only search

An S2-only analysis considers events without a detectable S1 signal, which are removed in analyses that use both S1 and S2 (termed ‘S1S2 analyses’ hereafter) [202, 203]. Although the name suggests differently, events with an S1 may still be included, though perhaps not analyzed in the same way as an S1S2 analysis. S2-only analyses have a lower detection threshold, resulting in sensitivity to lower-mass WIMPs and a higher expected event count for all WIMP masses. However, S2-only analyses face larger and less well-understood backgrounds, lack S1-based data quality cuts, drift-time based vertical fiducialization, and S1/S2-based ER/NR discrimination.

We can see the S1 threshold limits low-mass WIMP searches from figure 5.2b. NRs with the lowest detectable S1s have $cS2_b \approx 400$ PE, so $cS2 \approx 1000$ PE, compared to the analysis threshold of $S2 = 200$ PE and the 50% trigger efficiency point at $S2 \approx 110$ PE (see figure 4.8). With the $\approx 7 e^-/\text{keV}$ charge yield of low-energy NRs (see figure 2.6) and the ≈ 23 photons/e⁻ single electron gain (see section 3.4.1), an S2-only search should have appreciable sensitivity down to 1 keV, as opposed to the quick drop below ~ 5 keV in the

S1S2 analysis (see figure 5.1). The rate of low-mass WIMP scatters will be correspondingly increased (see figure 2.4).

XENON1T is well-suited to S2-only analyses for several reasons. Its ER background is at a record-low level (see section 5.4.2), so the loss of ER/NR discrimination is not as impactful. However, the main reason is that XENON1T is long and has a low field, which implies that S2 widths vary greatly throughout the TPC (due to diffusion, see section 3.5.3). Thus, the S2 signal width allows a limited (especially at very low S2) but effective vertical fiducialization. Hence, we can remove most backgrounds from interactions near the surface and (to some extent) the cathode with acceptable losses. This can reduce the remaining number of events remaining after cuts by several orders of magnitude.

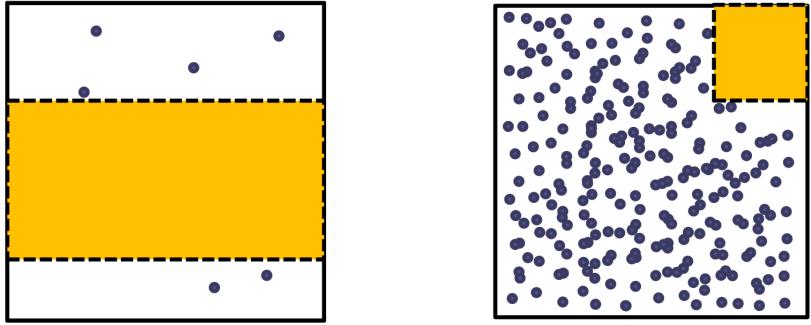
7.4.1 Limit setting among unknown backgrounds

An analysis that could make a dark matter discovery must deal with every background, either by modelling it, or by including a cut to reduce it to negligible levels. The S1S2 analysis uses both approaches; for example, it uses models for ER and accidental coincidences (see section 5.3) but cuts for gas events and unresolved double scatters (see section 2.4.3).

This is difficult for most backgrounds relevant in an S2-only analysis. Without the S1 signal, we have no accurate estimate of the depth of an interaction. Thus, backgrounds from near the surface (e.g. Compton scatters, see section 2.3.3.2), gas events (see section 2.4.3), and events originating from the cathode electrodes are difficult to identify. At low S2, inward-reconstructed events from the TPC wall also become more numerous, as position reconstruction worsens with lower S2 (fewer photons). Finally, backgrounds due to poorly-understood detector effects (e.g. pileup of spontaneous electron emission) cannot be ruled out.

Hence, S2-only analyses can currently only set limits on dark matter; even if it exists, they could not discover it. Since likelihood methods require a full background model, we also have to use a cut-and-count method for limit setting. Unfortunately, choosing effective cuts is difficult as well. While we know qualitatively what backgrounds to expect and what cuts would attenuate them (see table 7.2), choosing a good strength (or equivalently parameter boundaries) of each cut without a quantitative background model is nontrivial. In principle, this is a problem for S1S2 analyses too, but less severe, as unmodelled backgrounds are smaller and easier to remove.

The desired information about shapes and relative strengths of each background is in theory contained in the search data itself. One might wish to *tune cuts* on the search data, as illustrated in figure 7.8a, by selecting a region in parameter space (cS2 and data quality dimensions) where the signal model predicts relatively many WIMPs compared to the observed events. Clearly we cannot set ordinary counting limits with such regions: we expect less events



(a) Tuning to statistical fluctuations (b) Tuning to signal model errors

Figure 7.8: Problems with cut tuning illustrated in a toy model, with a 2d parameter space in which the assumed signal distribution is homogeneous. Circles represent observed events, yellow box the cut tuning result, i.e. the counting box for limit setting. **a)** For low event counts, we can select a region where the number of events is low by chance. This can be mitigated by simulating and correcting for the incurred bias (as in [190]) or by using different datasets for tuning and limit setting. **b)** For high event counts, the box is attracted to a region with an overestimated detection efficiency (here the upper right). This can only be avoided by using conservative signal models.

in a cut box tuned on data than in a pre-chosen cut box. If the cut tuning procedure is completely specified in advance, however, we can correct for this bias using our knowledge of the signal model and the procedure, as is done in the so-called maximum gap and optimum interval methods [190].

A different approach is to tune cuts on a different dataset than the data used in the dark matter search. In the S1S2 analysis, several cuts are tuned (non-algorithmically) this way on data outside the main region of interest (e.g. below the S2 threshold or beyond the fiducial volume). Calibration data is suboptimal for cut tuning, as background distributions are likely quite different, but this still suffices for some straightforward cuts in the S1S2 analysis (e.g. for removing gas events).

Here, we will split the original exposure in two; then use part of it as search data and part of it for cut tuning, as illustrated in figure 7.9. We split the exposure evenly between data for cut tuning ('training data') and data for limit setting ('test data'), according to the parity of the DAQ event numbers. For a sufficiently high number of cuts on uncorrelated parameters, it can be shown that such an even split is close to optimal [204].

Both methods of cut tuning (bias-correction as in optimum interval, and using separate datasets) are vulnerable to signal model uncertainties. If the signal expectation is overestimated in some region of parameter space, cut tuning will preferentially select it, as illustrated in figure 7.8b. As a result, the resulting limit can be overstated. This problem is not avoided by tuning

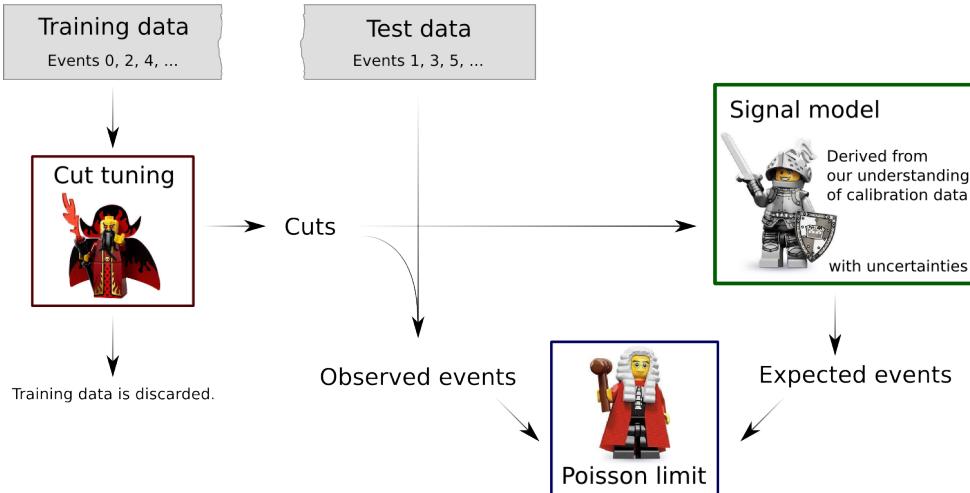


Figure 7.9: Overview of the S2-only analysis method presented in this section. Data of XENON1T’s first science run is split in two; one half (‘training data’) is used for tuning cuts, while the other (‘test data’) is kept blind until its use in the final limit setting. A signal model then conservatively estimates the number of WIMP events remaining after cuts, which we use to set a Poisson (cut-and-count) limit on the test data.

cuts on calibration data, since it is equally exposed to the most likely sources of unquantified biases (event processing, cuts, and assumptions on detector parameters). Instead, we have to parametrize our signal model uncertainties, then make a conservative estimate of the expected signal events in any box considered during cut tuning. This is discussed in section 7.4.4 below.

7.4.2 Signal model

We compute the dark matter event rate $r(S2, z)$ for a given WIMP mass m_χ and cross-section σ in a finely binned 2d-grid in $\log_{10} S2$ and z (depth below the gate mesh). Although z cannot be derived from an observed drift time without an S1, the signal model varies substantially with z (due to electron lifetime) and some cuts have a z -dependent acceptance (especially the width cut). Cut efficiencies are computed in the same $(S2, z)$ space.

Unlike for the S1S2 analysis, the signal model is straightforwardly computed without Monte Carlo simulations. Starting from the differential WIMP rate $f(E) = dR/dE$ (equation 2.2), the rate e_p of events with n produced electrons is:

$$e_p(n) = \int dE f(E) \text{Poisson}(n | EQ_y(E)), \quad (7.2)$$

where $Q_y(E)$ is the charge yield function in electrons per keV and E the recoil energy in keV. We assume the charge yield of [110] (figure 2.6), cut off below 0.7 keV, the lowest energy at which NR charge yield has been observed [109].

The rate e_d of events with k *detected* electrons depends on z through the electron lifetime τ :

$$e_d(k, z) = \sum_n e_p(n) \text{Binom}(k | n, p = \eta e^{-\frac{z}{v_d \tau}}). \quad (7.3)$$

Here η is the electron extraction efficiency and v_d the electron drift velocity. Since τ varies throughout the run, we compute eq. 7.3 for several lifetimes and average the result proportionally to the occurrence of each lifetime during the run. Finally, the signal model $r(S2, z)$ is found by convolving the approximately Gaussian single electron spectrum:

$$r(S2, z) = \sum_k c(k) \text{Normal}(S2 | \mu = \mu_{\text{SE}} k, \sigma = \sigma_{\text{SE}} \sqrt{k}), \quad (7.4)$$

with μ_{SE} and σ_{SE} the mean and standard deviation of the single electron size distribution in PE, respectively. The spatial variation in XENON1T’s S2 response (due to varying extraction efficiency and PMT response) is only $\sim 5\%$ RMS within the radii considered in this analysis, and will be ignored.

We use the depth-dependent trigger efficiency from figure 4.7. As we will see below, the cut tuning chooses to remove data with $S2 \approx 140$ PE; thus, we only rely on signal response from regions with $> 80\%$ trigger efficiency.

The resulting signal model is shown in figure 7.10. Background models for ER and CNNS are computed similarly. For the ER background we assume a $0.193 \text{ events}/(\text{kg} \times \text{day} \times \text{keV}_{\text{ee}})$ flat energy spectrum (see section 5.3.1 and the charge yield from the fit of section 6.2); for CNNS we take the energy spectrum of a $6 \text{ GeV}/c^2$ WIMP with a WIMP-nucleon cross section of $4 \times 10^{-45} \text{ cm}^2$ (an excellent approximation, see figure 3 of [168]).

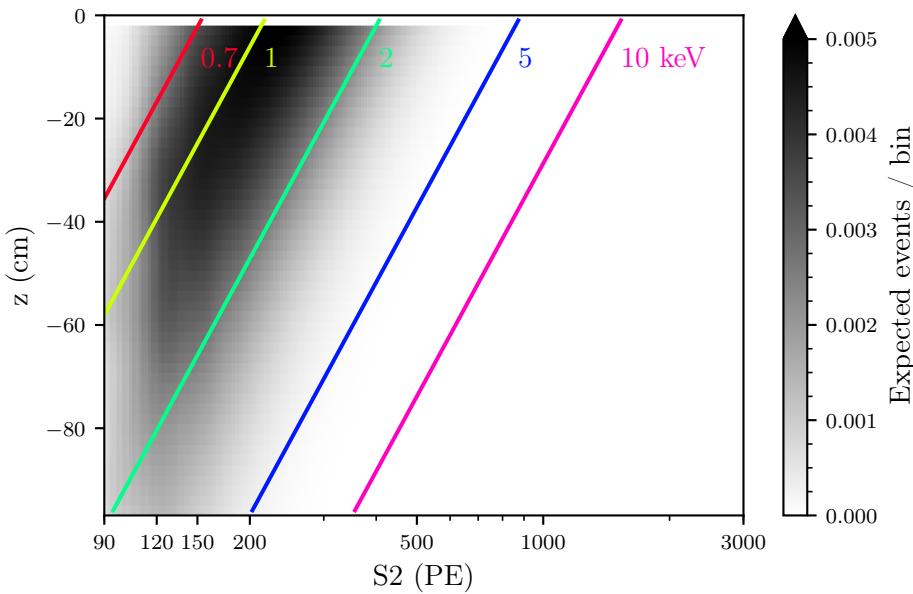
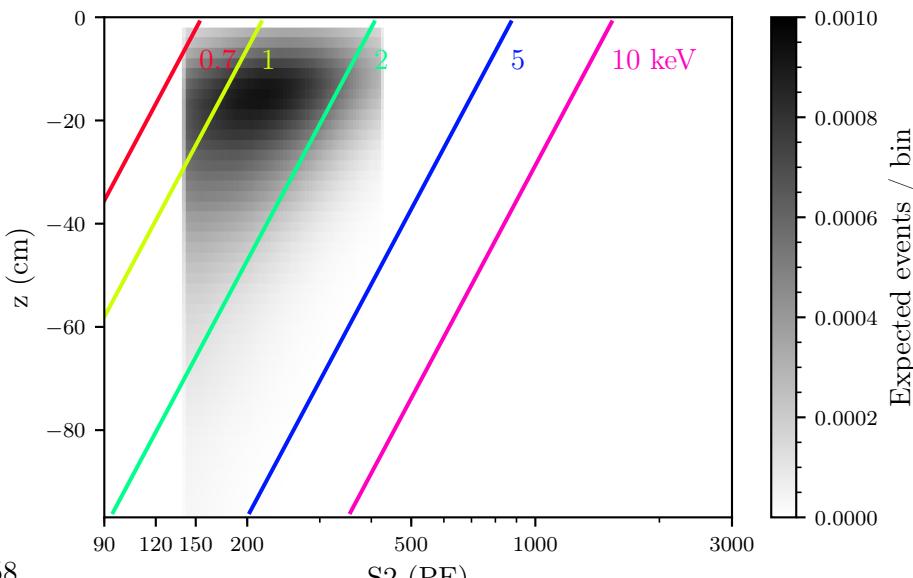


Figure 7.10: Shading: signal model for elastic spin-independent WIMP-nucleon scattering for $m_\chi = 6 \text{ GeV}/c^2$ and $\sigma = 4 \times 10^{-45} \text{ cm}^2$, and equivalently background model for CNNS events, for XENON1T’s first science run, before applying cuts, as a function of S2 size and interaction depth. Coloured lines show the nominal signal of a mono-energetic nuclear recoil source of the indicated energy in keV. The total expected number of WIMP events is 8.0.

Figure 7.11: As figure 7.10, after cuts optimized for $6 \text{ GeV}/c^2$ WIMPs (see text). Note the different shading scale. The total expected number of WIMP events is 0.66.



7.4.3 Event selection

Table 7.2 lists the seven major event selections we use in the analysis, along with which backgrounds they are primarily targeted against. These cuts are tuned on training data, as described in section 7.4.4. This section discusses details of the cut’s definitions and signal acceptance determination.

For the signal width and shape cuts, we will use two variables related to the S2 shape illustrated in figure 7.12: the 50% area width (or interquartile range) $\Delta t_{50\%}$ discussed in section 3.5.3, and the ‘full rise time’ Δt_r , which is the time between the start of the signal and its center (where 50% of its charge has been deposited). For ordinary S2s, these are both strongly correlated with the signal depth, and $\Delta t_r \approx 2 \Delta t_{50\%}$. Waveforms other than S2s do not generally obey this. Thus, we define transformed variables:

$$\text{width} = (\Delta t_{50\%} + \Delta t_r)/2 \quad (7.5)$$

$$\text{shape} = (\Delta t_{50\%} - \Delta t_r)/2 \quad (7.6)$$

which we will perform separate cuts on.

In the order of table 7.2, the cuts we use are:

- *Area cut:* Which S2 areas a WIMP-nucleon scattering produces depends strongly on the WIMP mass, so a different range in S2 is optimal for each mass. The signal loss is determined trivially, as S2 is a primary signal model dimension.

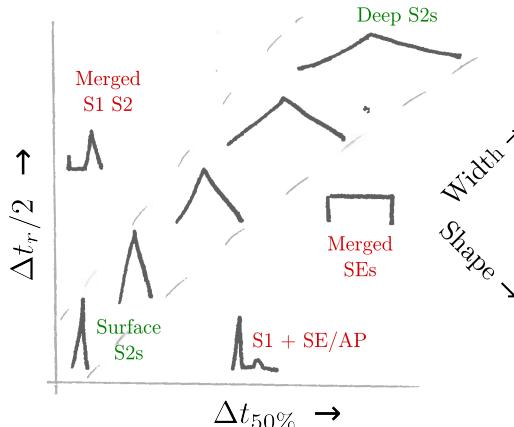


Figure 7.12: Illustration of waveform shapes encountered in the S2-only analysis and their width and shape parameters as defined in the text. Real S2s lie in the central band, with S2s from interactions deeper in the TPC further towards the top right. Outside the band are various artifacts: S2s merged with a preceding S1 (top left), merged single electrons (center right) and S1s merged with single electrons and afterpulses, then misclassified as S2s (bottom right).

| | Area | Width | Shape | Pattern | AFT | Tail | R |
|----------------------|------|-------|-------|---------|-----|------|-----|
| Events removed | 31 | 117 | 3 | 6 | 145 | 19 | 890 |
| Signal loss (%) | 30 | 62 | 6 | 10 | 19 | 15 | 45 |
| Target background | | | | | | | |
| ER | | X | | | | | |
| CNNS | | | | | | | |
| Wall | X | | | | | | X |
| Surface, Cathode | X | X | | | | | |
| Gas | X | X | X | X | X | | |
| Processing artefacts | X | X | X | X | X | | |
| SE tails | X | X | X | X | | | X |
| Spontaneous SE | X | X | X | X | | | |

Table 7.2: Columns: major event selections (‘cuts’) used in the S2-only analysis (‘R’ denotes the cut on the observed radial position), described in the text. Three events remain after all cuts, which together cause a signal loss of 92%. The top two rows list the number of events in the test data and expected signal fraction removed in the 4 GeV/c² WIMP analysis after applying all other cuts (see figures 7.13-7.16, second panel from bottom). The rest of the table indicates which background populations each cut is targeted against. The S2 area cut is effective unless a background has a very similar S2 spectrum as the signal, as is the case for CNNS (especially for 6 GeV/c²WIMPs). Processing artefacts are primarily S1s merged with photoionization and afterpulse signals, then misclassified as S2s. ‘SE tails’ refer to pileup of single electrons from delayed extraction after high-energy events; ‘spontaneous SE’ means pileup of single electrons from spontaneous electron emission (e.g. the ‘impulsive emitters’ from [205]).

- *Width cut*: As discussed in section 3.5.3, the S2 signal width increases with interaction depth due to diffusion. This allows an S2 width cut to select events by depth (although far less accurately than the drift time cut in the S1S2 analysis), suppressing large backgrounds from near the cathode and particularly the surface. We use the width parameter defined in equation 7.5. The signal loss is estimated from simulated waveforms using the waveform simulator discussed in chapter 3. We simulate S2s homogeneously in the $(S2, z)$ search space, then estimate the distributions in $(S2, z, width)$ and $(S2, z, shape)$ using kernel density estimation (KDE).
- *Shape cut*: As mentioned above, the conjugate of the width parameter (equation 7.6) is useful for removing events with processor artifacts misidentified as S2s. The signal loss is estimated as for the width cut, using a separate 3-dimensional KDE. Estimating a single distribution

in $(S2, z, width, shape)$ with a 4D KDE would be more accurate, but is computationally much more complex.

- *Pattern cut*: Remove events whose S2 pattern across the top PMTs is very different from the expected pattern at their reconstructed (x, y) position. This primarily targets unresolved double scatters and single-electron pileup. The cut bound is set to ensure a fixed signal loss of 10% at all S2 energies, as estimated from downsampling hitpatterns of large S2s found in data. We cannot use simulated S2s generated from the expected hitpattern maps (which would give an underestimated signal loss unless the maps are perfect).
- *AFT cut*: The fraction of the S2 signal observed by the top array ('area fraction top' or simply AFT) discriminates well between S2s from events in the liquid and scintillation signals in the gas phase higher in the detector. We correct the observed AFT for a slight dependence on the (x, y) position (due to varying light collection efficiency of especially the top array). The resulting distribution is straightforwardly modeled using binomial photon statistics and a small intrinsic variation visible primarily at high energy (where it is well-described by a student-t distribution), possibly due to imperfections in the correction map and inaccuracy of the position reconstruction resolution. The model reproduces the AFT distribution of $^{241}\text{AmBe}$ events remaining after the S1S2-analysis cuts, and we thus use it for estimating the signal loss of the cut.
- *Tail cut*. Some events with large ($\gtrsim 10^6$ PE) S2s are followed by a tail of events consisting mostly of single electrons, possibly from delayed charge extraction [140]. Pile-up of such single electrons can cause fake S2 signals. Though these are usually wider than allowed by our width cut, we further reduce this background by a cut on the parameter:

$$T = \max_{1 \leq i \leq 100} S2_i / \Delta t_i, \quad (7.7)$$

where $S2_i$ denotes the largest S2 of the event occurring i events before the current event, and Δt_i the time between that event and the current event. The same parameter is used in the main XENON1T analysis. The distribution of T is easily computed from the observed S2s in the full (unblind, but completely uncut) data. Since T does not depend on properties of the dark matter candidate event itself, this cut has a constant efficiency in $(S2, z)$.

- *Radial cut*: We remove events with reconstructed radial position $r^2 < 1200 \text{ cm}^2$, suppressing the wall leakage events discussed in section 5.3.4. Since we have no accurate estimate of the event depth, we cannot correct the radial position for field distortion effects (see section 2.4.2). Fortunately, our analysis will focus mostly on the top of the detector,

where field distortion effects are less severe. Regardless, we still compute the proper z -dependent signal loss of this cut using the inverse of the field distortion correction map.

Besides these major cuts, we apply three smaller selections (not tuned on training data):

- Remove events with DAQ BUSY signals and events near regions with known acquisition problems, exactly as in the S1S2 analysis. The signal loss is accounted for in the DAQ deadtime, and therefore the exposure calculation of the run.
- Remove events with $S1 > 70\text{ PE}$, which are also removed in the S1 analysis. The signal loss is negligible for the low-mass WIMPs of interest here. More stringent S1 cuts could be beneficial, but would require an additional signal model dimension to fully account for their signal loss.
- Remove events with an S1 that have a drift time $> 650\text{ }\mu\text{s}$, close to the drift time of the cathode (see figure 3.14). This removes events from regions near the cathode with suppressed charge yield (and therefore a low S2). The signal loss is negligible, considering that few WIMP signals come from deep in the detector (see figure 7.10) and deep events will be removed anyway by the S2 width cut (see below).

7.4.4 Cut tuning and limit setting

The seven major cuts of the S2-only analysis are tuned on the training data. A few of cut boundaries were chosen by hand (for the r cut and the signal shape cut) for simplicity. The other cut bounds are chosen algorithmically: we vary the boundaries of the cuts until we find a sparse region in the training data, as quantified by the limit setting method described below. This cannot be done with ordinary derivative-based optimizers, as the objective function jumps discontinuously at training data events, so we used a global optimization based on differential evolution [206].

After this optimization, we make each cut bound slightly stricter than what would appear optimal on training data, incurring 3% extra signal loss for each tuned bound (a total of 24% signal loss). This avoids using regions with very small signal expectation where (by chance) no background events are seen in the training data. We verified this is likely to improve the end result by repeatedly splitting the training dataset randomly in even halves, tuning cuts on one half, and examining the result on the other half.

As explained in section 7.4.1, cut tuning methods can yield overstated results if the signal expectation is overestimated in any region of parameter space (see figure 7.8b). To prevent this, we consider several uncertainties on the signal model:

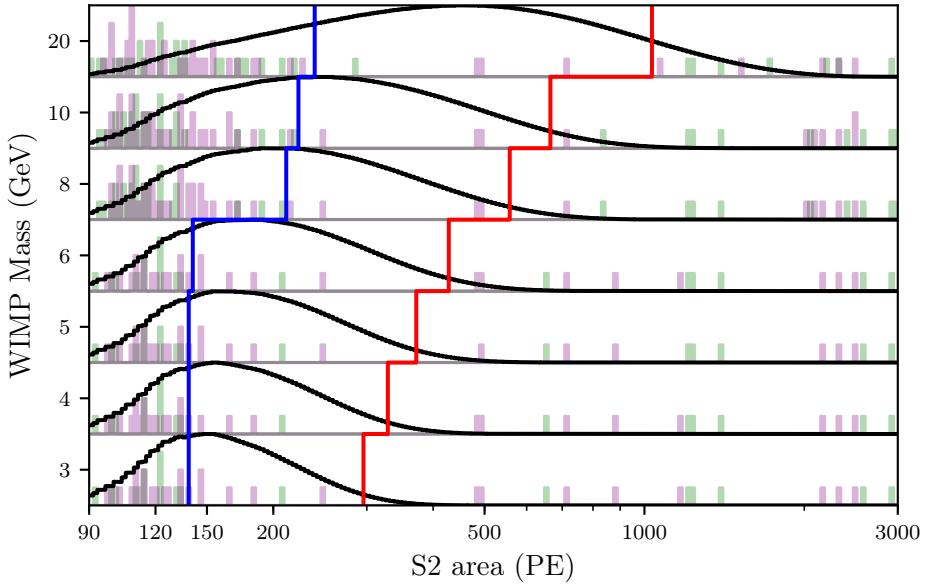


Figure 7.13: Histograms: training (green) and test (purple) events after all cuts except the S2 area selection, for cuts optimized for different WIMP masses. Note cuts other than S2 area also differ for each WIMP mass. Histogram values range from 0 to 4 events, with higher counts off-scale (i.e. shown as 4 events). Blue and red lines show lower and upper cut bounds, respectively. Black lines show WIMP signal models in arbitrary units, also after all cuts except the S2 area selection.

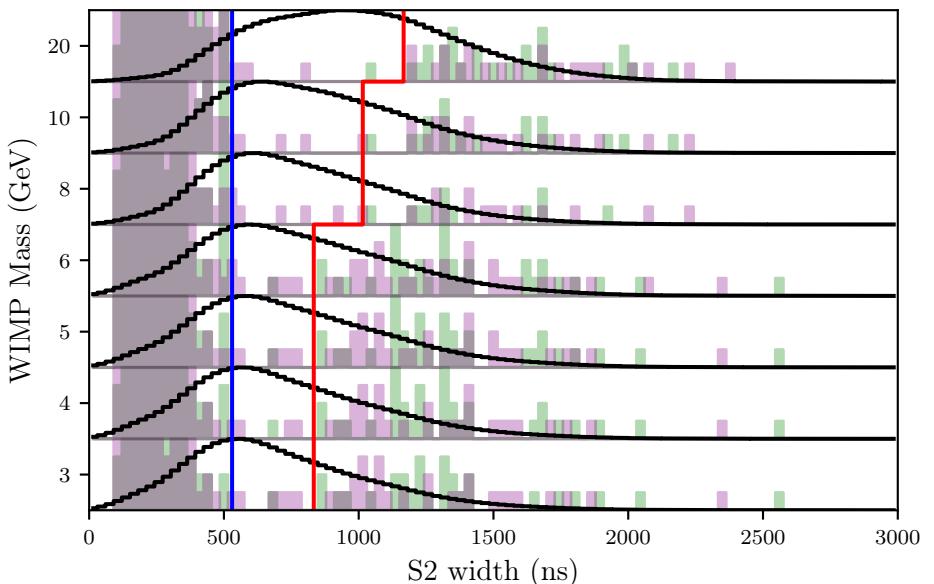


Figure 7.14: As figure 7.13, with the signal width cut left out instead. Surface events are clearly visible at low widths.

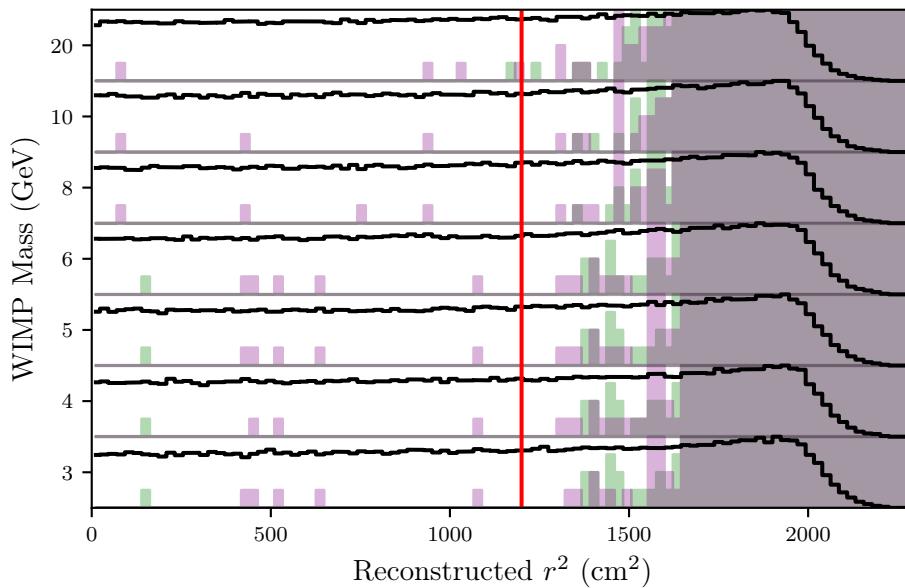


Figure 7.15: As figure 7.13, with the radial cut left out instead. Events from near the TPC wall are visible at high radii.

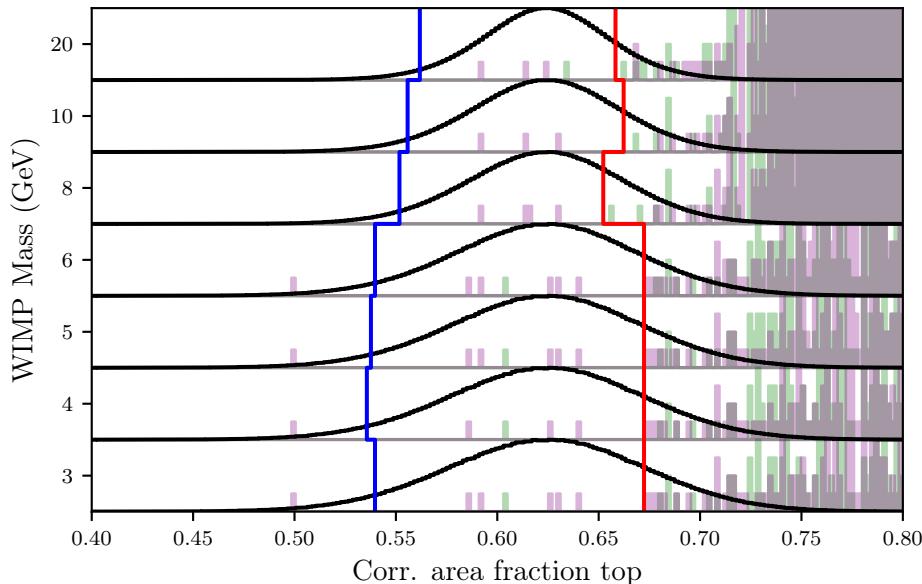


Figure 7.16: As figure 7.13, with the cut on the fraction of S2 observed in the top array (corrected for a small position dependence) left out instead. Gas events are visible at high values.

- *Single electron gain* of (29.7 ± 2) PE. The central value is the average over SR0 (using the method of section 3.4.1), the uncertainty is motivated by nonlinearities of the event processing (estimated at 1 – 3%) and variations during the run (see figure 3.11, which is expressed in number of hits rather than PE).
- *Extraction efficiency* of $96.1 \pm 4.6\%$, as discussed in section 3.4.1. The upper bound of the extraction efficiency is taken at 100% rather than the non-physical 100.7%.
- *Electron lifetime* uncertainty of 5%. Although the lifetime is known to 2.4% accuracy, we use a larger uncertainty to account for any systematics introduced by calculating the signal model as the proportional average of models computed with several lifetimes (see section 7.3.1).
- *Surface cutoff* of (2.5 ± 1) cm. At depths below this, no signal response will be assumed. It was observed the data processor sometimes merges S1 and S2s together for events in the first ~ 0.7 mm [207]; to be conservative we cut off a considerably larger region.

When computing the expected number of signal events in a box, we compute its uncertainty by adding the effects of the above uncertainties in quadrature. Then, we use the 90% worst-case, i.e. the nominal value minus 1.28 times the uncertainty, in the limit setting. Unless a large signal model uncertainty is left unquantified, this should be a conservative estimate of the expected number of signal events.

7.4.5 Results

Figures 7.13-7.16 show the cut bounds that result from cut tuning on the training data for the four most important cuts, for several WIMP masses. Notice we always select regions where the signal model shows a high density, rather than relying on a (perhaps poorly understood) tail. Figure 7.14 clearly shows that the width removes a large background of narrower signals (likely from events at the surface) and wider signals (e.g. from higher-energy events deeper in the detector, SE pileup, etc). Other clear populations include gas events (figure 7.16 at high AFT) and events near the wall (figure 7.15 at high radius). Figure 7.11 shows the signal model of a $6 \text{ GeV}/c^2$ WIMP after cuts.

The purple histograms in figures 7.13-7.16 show the test data events that remain after applying the tuned cuts. For $m_\chi < 10 \text{ GeV}/c^2$, three to five events remain in the selected region, while the background expectation from CNNS and ER together is about 0.7 events (dominated by CNNS). We stress this should not be regarded as evidence for dark matter, but evidence for the relevance of the many remaining unmodelled backgrounds (see table 7.2). Our treatment of signal model uncertainties causes a $\sim 20\%$ reduction in the

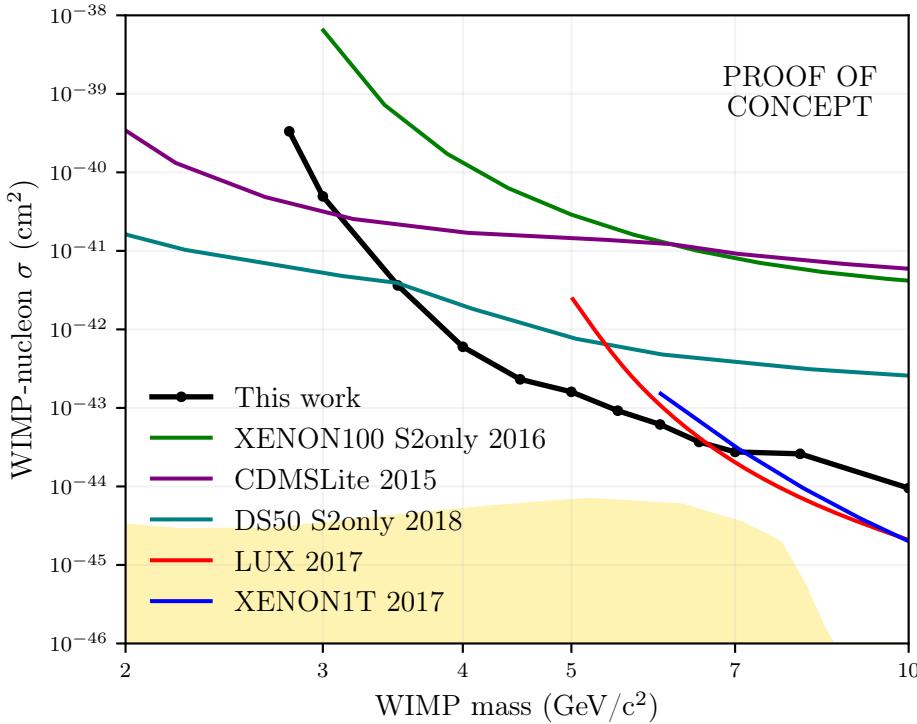


Figure 7.17: 90% C.L. limit on spin-independent WIMP-nucleon scattering set by an S2-only analysis of XENON1T’s first science run (black line, proof of concept). Coloured lines show results from S2-only analyses from XENON100 [203] (green) and DarkSide-50 (DS-50, [208]), and main results from XENON1T (chapter 5, blue), LUX [196] (red), and CDMSLite ([209]). The yellow-shaded region shows the neutrino-floor sensitivity [4].

expected number of WIMP events for most masses – a small reduction, as we would expect from selecting a relatively central region in every cut.

Figure 7.17 shows the limits on spin-independent WIMP-nucleon scattering resulting from this analysis. No sensitivity estimate can be shown as we have only a partial background model. Our constraints improve upon the current world-leading limits by up to a factor 5, though only in the narrow range of WIMP masses ($3.5\text{--}6.5\text{ GeV}/c^2$) where S2-only searches are relevant. S2-only searches extend the power of LXe TPCs – caused by their large mass, high purity, and high A^2 coherent enhancement (see section 2.1.4) – a bit further out into the parameter space of light WIMPs.

A comparison with XENON100’s S2-only analysis [203] (green in figure 7.17) is instructive. The most significant difference is that we used S2 width to effectively cut backgrounds from the surface and from deeper regions of the detector. Without it, we would have seen 117 rather than 3 events in the 4

GeV/c² search box, causing an approximately twenty times higher (i.e. worse) limit. The XENON100 analysis used a very loose width cut – but even if it had used a tighter cut, that would not have worked as well, since XENON100 is shorter and has a lower diffusion constant than XENON1T (due to its higher drift field). XENON100 also had a far larger low-S2 background event rate: inside the fiducial volume used for its S2-only analysis, the rate per kilogram before data quality cuts is more than twice as high as in XENON1T *before* fiducialization [204]. Finally, the XENON100 analysis assumed a more conservative Q_y (~ 6 e⁻/keV rather than ~ 7.5 e⁻/keV at low energy), and only tuned its S2 area cut to background measurements (using optimum-interval limit setting).

7.4.6 Discussion

Our proof-of-concept result (figure 7.17) clearly shows the power of S2-only analyses with modern detectors such as XENON1T. Besides its competitive dark matter limit, it provides a starting point for more detailed studies of the low-S2 background of LXe TPCs. These could improve S2-only searches further: currently, we lose half our exposure to cut tuning, but the cuts could still be quite suboptimal due to limited statistics for rare backgrounds in the training data. If some remaining detector backgrounds could be characterized even approximately, that could be used to define better cuts, leading to better separation of the remaining backgrounds. The ultimate goal would be to cut away or model all relevant backgrounds, and make S2-only analyses capable of making discovery claims.

As we have emphasized, this analysis is a proof of concept, and therefore still has remaining issues. Most importantly, we need additional validation of the cut acceptances used. Ideally, for each cut, we would cross-check the acceptance by a suitable control sample from calibration data. Unfortunately, these are challenging to obtain, as backgrounds that affect the S2-only search are also present (and perhaps even larger) in calibration data. The main S1S2 analysis cuts can provide cleaner samples, but these do not extend below the main analysis threshold of S2 = 200 PE, whereas 150-200 PE is most important for $m_\chi < 6$ GeV/c² (see figure 7.13).

Validating the S2 width cut on calibration data is particularly problematic. The S2 width depends crucially on depth (which is the reason we cut on it), but depth can only be measured reliably using drift time, i.e. on events with S1s. However, S2-only searches are by definition interested in the (S2, depth) regions where even nuclear recoils almost never make S1s. We would have to obtain control samples from special cases where the S2 is small but the S1 nonetheless available – e.g. double scatters, events affected by charge loss at the detector walls, or early XENON1T data with very low detector purity – but contaminations and/or differences with the main event population would make obtaining reliable measurements with such samples a daunting

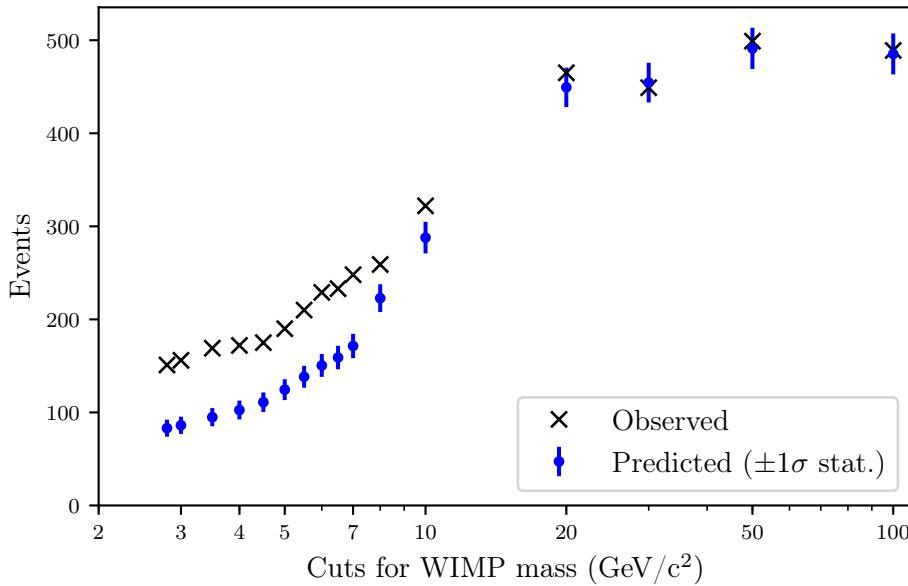


Figure 7.18: Blue bars: predicted number of events in XENON1T $^{241}\text{AmBe}$ calibration data (see section 5.2.4) after S2-only analysis cuts for different WIMP masses (recall our analysis uses a different set of cuts for each WIMP mass). Crosses: observed events after these cuts. An additional cut against double scatters (the same as used in the main S1S2 analysis) is applied to ensure the Monte Carlo simulation results can be compared with data; its signal acceptance is $> 95\%$ at all S2 in the main analysis, and approaching 100% for low S2s [210]. The distribution of the tail parameter (equation 7.7) has been recalculated, since $^{241}\text{AmBe}$ has a different rate of high-energy events than background data.

prospect. An alternative would be a dedicated calibration campaign with an external source positioned at different heights. A consolation is that for the relatively central selections shown in figures 7.13-7.16, large underestimates of acceptance losses could only come from very substantial mismodelling.

A possibly useful test for the signal model as a whole, including cut acceptances, is predicting how many events survive after applying the S2-only cuts to $^{241}\text{AmBe}$ calibration data. Since we have a different set of cuts for each WIMP mass, this gives several different (but not completely independent) tests. Figure 7.18 shows the results of these tests for this analysis (without the 0.7 keV cutoff and signal model uncertainties). We see as many events as predicted after cuts optimized for high WIMP masses, but see *more* events than predicted after the cuts of lower-mass WIMPs. Perhaps $^{241}\text{AmBe}$ has a significant extra background source (e.g. from delayed extraction tails of AmBe's high-energy gamma rays), or the low-energy charge yield is higher

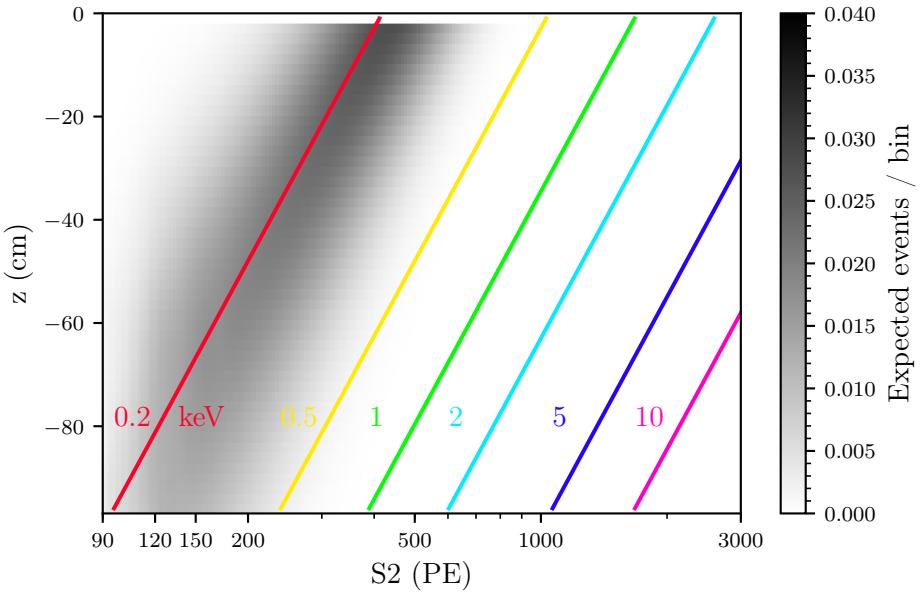


Figure 7.19: Signal model for spin-independent WIMP-nucleon scattering detected through bremsstrahlung, for $m_\chi = 0.1 \text{ GeV}/c^2$, $\sigma = 2 \times 10^{-31} \text{ cm}^2$, for XENON1T’s first science run, before applying cuts. Coloured lines show the nominal signal of a mono-energetic nuclear recoil source of the indicated energy in keV. The expected number of events is 45 (4.2 after cuts).

than assumed (which measurements from LUX [109] also seem to indicate), or our analysis has a problem which leads to overly conservative predictions. While neither of these invalidates our results, a better matching would clearly bolster confidence in the accuracy of our models. Further research is needed on this point.

7.5 S2-only bremsstrahlung search

As a final extension of the dark matter analysis, we repeat the S2-only analysis of section 7.4 with a signal model based on bremsstrahlung (as in section 7.3) rather than elastic nuclear recoils. Since the S2-only analysis only relies on charge yield, we can now use the signal response cutoff corresponding to the lowest measured ER charge yield at 0.19 keV, rather than the 1 keV cutoff at to the lowest measured ER light yield (see section 7.3 for discussion and references).

Figure 7.19 shows the S2-only signal model for an $m_\chi = 0.1 \text{ GeV}/c^2$ WIMP. Notice that the signal cutoff is still relevant for the analysis; signals from ERs below 0.19 keV would be detectable, at least in the top of the detector, if we

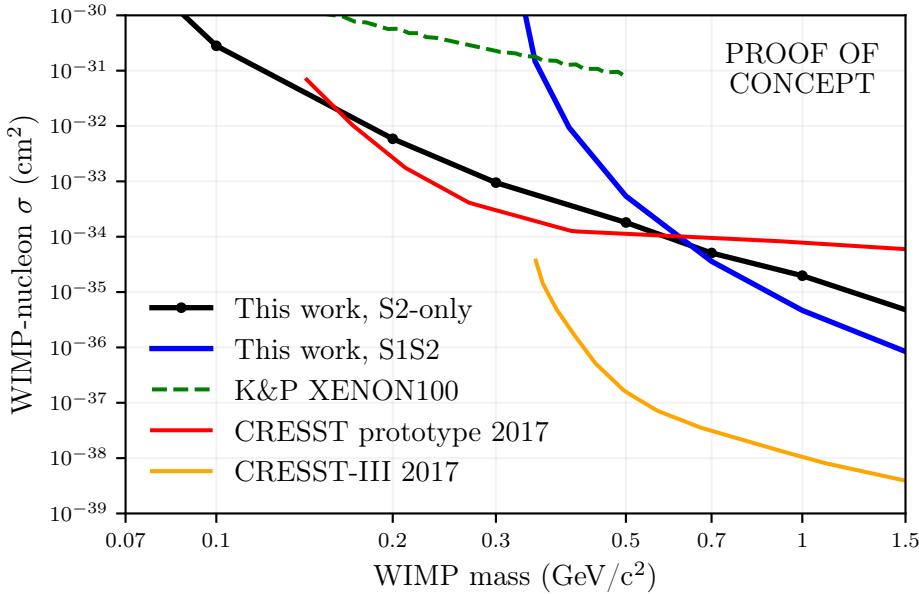


Figure 7.20: Black line: 90% C.L. limit on bremsstrahlung photons from spin-independent WIMP-nucleon scattering in an S2-only analysis of XENON1T’s first science run (proof of concept). The dashed green line shows the result of the original bremsstrahlung paper [96] using the XENON100 S2-only analysis [203] (dashed green), and the blue line shows the result from the elastic nuclear recoil bremsstrahlung analysis (section 7.3). Other coloured lines are as in figure 7.7.

could be sure they exist.

We follow the same cut optimization as described in section 7.4. For $m_\chi < 0.3 \text{ GeV}/c^2$, the optimized cuts are similar to each other (most signal is expected near the cutoff) and similar to the optimal cuts for the $8 \text{ GeV}/c^2$ nuclear recoil search (see figures 7.13 - 7.16).

After cuts, 2 to 3 events remain in the test data for $m_\chi < 1 \text{ GeV}/c^2$. The total expected background from ER and CNNS for these analyses is less than 0.5 events, dominated by ER – again, this is not evidence for dark matter, but evidence the remaining unmodelled backgrounds are significant. Figure 7.20 shows the resulting limits on the spin-independent WIMP-nucleon scattering cross-section. Our constraints are competitive with the current world-leading low-mass limits from [200], even extending these to slightly lower masses. The limits are stronger than our S1S2 bremsstrahlung analysis (in blue, section 7.3) below $m_\chi = 0.6 \text{ GeV}/c^2$ due to the lower signal response cutoff. Even if the S1S2 analysis could use the lower 0.19 keV cutoff, it would still be weaker than the S2-only analysis at the lowest masses, since almost all quanta in ER tracks at low energy are converted to electrons rather than photons (see

section 2.3.2).

Our limits are stronger than those of the original bremsstrahlung paper [96], which were based on the XENON100 S2-only analysis [203]. The difference in limits is similar to the difference between the XENON100 and XENON1T S2-only searches for elastic NR (see figure 7.17) at $0.5 \text{ GeV}/c^2$, but less at lower WIMP masses. This is likely because [96] did not use the 0.19 keV cutoff, and instead assumed charge yield down to the ionization threshold of 12 eV (or at least down to a point where XENON100 would no longer be sensitive).

We do not show limits higher than 10^{-30} cm^2 : at such high cross-sections, dark matter particles would scatter significantly inside the earth [96], reducing the energy and flux as a result (see also section 1.2.4). At slightly lower cross-sections ($\sim 10^{-31} \text{ cm}^2$), constraints from earth heating due to annihilation of dark matter particles and antiparticles becomes relevant (see section 1.2.5) if such dark matter annihilation can indeed occur.

7.6 Outlook

The search for dark matter is in full swing, and DM-scattering experiments such as XENON1T are leading the charge. XENON1T will soon unblind a $\sim 1 \text{ ton}$ year run under similar conditions as the data analysed in this work: an exposure three times larger than that of the three currently world-leading results (from XENON1T, LUX, and PandaX) put together. This represents a giant leap in sensitivity and an exciting opportunity for discovery.

After a short phase dedicated to R&D and specialized searches, XENON1T will be decommissioned and replaced with its upgrade, XENONnT. This will feature a 6-ton active mass, a dedicated cryogenic distillation column to reduce the ^{222}Rn background dominant in XENON1T, and a neutron veto subdetector to attenuate backgrounds from radio- and cosmogenic neutrons [2]. Plans for dark matter detectors after XENONnT that reach neutrino-floor sensitivity are already drawn up [3]. In parallel, collider and annihilation searches for dark matter similarly plan strides in scale to discover complementary signatures of dark matter. Within a decade, we will either rule out most currently popular dark matter models, or make a sensational discovery, and finally understand what much of the universe is made of.

Bibliography

- [1] Planck Collaboration et al. “Planck 2015 Results. XIII. Cosmological Parameters”. *Astronomy & Astrophysics* 594 (Oct. 2016), A13. doi: 10.1051/0004-6361/201525830. arXiv: 1502.01589 (cit. on pp. 6, 14, 15, 21, 107, 186).
- [2] XENON Collaboration. *XENONnT Technical Design Report (in Preparation)*. 2018 (cit. on pp. 10, 171, 189).
- [3] J. Aalbers et al. “DARWIN: Towards the Ultimate Dark Matter Detector”. *Journal of Cosmology and Astroparticle Physics* 2016.11 (Nov. 8, 2016), pp. 017–017. doi: 10.1088/1475-7516/2016/11/017. arXiv: 1606.07001 (cit. on pp. 10, 171, 189).
- [4] 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* 89.2 (Jan. 27, 2014). doi: 10.1103/PhysRevD.89.023524. arXiv: 1307.5458 (cit. on pp. 10, 46, 166, 191).
- [5] J. Aalbers and C. Tunnell. *The Pax Data Processor v6.4.2*. Zenodo, Feb. 19, 2017. doi: 10.5281/zenodo.546239. URL: <https://zenodo.org/record/546239#.WZWq1CdLfbo> (visited on 08/17/2017) (cit. on pp. 11, 54, 101).
- [6] E. Aprile et al. “First Dark Matter Search Results from the XENON1T Experiment”. *Physical Review Letters* 119.18 (Oct. 30, 2017). doi: 10.1103/PhysRevLett.119.181301 (cit. on pp. 11, 107).
- [7] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric Dark Matter”. *Physics Reports* 267.5 (Mar. 1, 1996), pp. 195–373. doi: 10.1016/0370-1573(95)00058-5 (cit. on pp. 12, 22, 28, 107).
- [8] Gianfranco Bertone, Dan Hooper, and Joseph Silk. “Particle Dark Matter: Evidence, Candidates and Constraints”. *Physics Reports* 405 (5–6 Jan. 2005), pp. 279–390. doi: 10.1016/j.physrep.2004.08.031. arXiv: hep-ph/0404175 (cit. on pp. 12, 17, 33).
- [9] M. Milgrom. “A Modification of the Newtonian Dynamics as a Possible Alternative to the Hidden Mass Hypothesis”. *The Astrophysical Journal* 270 (July 1, 1983), pp. 365–370. doi: 10.1086/161130 (cit. on p. 12).
- [10] Jacob D. Bekenstein. “Relativistic Gravitation Theory for the MOND Paradigm”. *Physical Review D* 71.6 (Mar. 14, 2005). doi: 10.1103/PhysRevD.71.069901. arXiv: astro-ph/0403694 (cit. on p. 12).
- [11] Scott Dodelson. “The Real Problem with MOND”. *International Journal of Modern Physics D* 20.14 (Dec. 31, 2011), pp. 2749–2753. doi: 10.1142/S0218271811020561. arXiv: 1112.1320 (cit. on pp. 12, 15).
- [12] The LIGO Scientific Collaboration and the Virgo Collaboration. “Observation of Gravitational Waves from a Binary Black Hole Merger”. *Physical Review Letters* 116.6 (Feb. 11, 2016). doi: 10.1103/PhysRevLett.116.061102. arXiv: 1602.03837 (cit. on pp. 12, 17).

[13] Paul M. Chesler and Abraham Loeb. “Constraining Relativistic Generalizations of Modified Newtonian Dynamics with Gravitational Waves”. *Physical Review Letters* 119.3 (July 18, 2017). DOI: 10.1103/PhysRevLett.119.031102. arXiv: 1704.05116 (cit. on p. 12).

[14] The LIGO Scientific Collaboration and the Virgo Collaboration. “Tests of General Relativity with GW150914”. *Physical Review Letters* 116.22 (May 31, 2016). DOI: 10.1103/PhysRevLett.116.221101. arXiv: 1602.03841 (cit. on p. 12).

[15] J. C. Kapteyn. “First Attempt at a Theory of the Arrangement and Motion of the Sidereal System”. *The Astrophysical Journal* 55 (May 1, 1922), p. 302. DOI: 10.1086/142670 (cit. on p. 13).

[16] J. H. Oort. “The Force Exerted by the Stellar System in the Direction Perpendicular to the Galactic Plane and Some Related Problems”. *Bulletin of the Astronomical Institutes of the Netherlands* 6 (Aug. 1, 1932), p. 249 (cit. on p. 13).

[17] F. Zwicky. “On the Masses of Nebulae and of Clusters of Nebulae”. *The Astrophysical Journal* 86 (Oct. 1, 1937), p. 217. DOI: 10.1086/143864 (cit. on p. 13).

[18] Guido D’Amico, Marc Kamionkowski, and Kris Sigurdson. “Dark Matter Astrophysics” (July 10, 2009). arXiv: 0907.1912 [astro-ph] (cit. on p. 13).

[19] Matts Roos. “Dark Matter: The Evidence from Astronomy, Astrophysics and Cosmology” (Jan. 4, 2010). arXiv: 1001.0316 [astro-ph] (cit. on p. 13).

[20] Chandra Collaboration. *1E 0657-56: NASA Finds Direct Proof of Dark Matter*. Aug. 21, 2006. URL: <http://chandra.harvard.edu/photo/2006/1e0657/> (visited on 01/02/2018) (cit. on p. 13).

[21] Douglas Clowe et al. “A Direct Empirical Proof of the Existence of Dark Matter”. *The Astrophysical Journal* 648.2 (Sept. 10, 2006), pp. L109–L113. DOI: 10.1086/508162. arXiv: astro-ph/0608407 (cit. on pp. 13, 15).

[22] G. B. Airy. “Account of Some Circumstances Historically Connected with the Discovery of the Planet Exterior to Uranus.” *Astronomische Nachrichten* 25.10 (Jan. 1, 1847), pp. 131–148. DOI: 10.1002/asna.18470251002 (cit. on p. 14).

[23] Albert Einstein. “Explanation of the Perihelion Motion of Mercury from the General Theory of Relativity”. *Sitzungsber.Preuss.Akad.Wiss.Berlin (Math.Phys.)* 1915 (1915), pp. 831–839 (cit. on p. 14).

[24] Simon Foreman and Douglas Scott. “What Do Gas-Rich Galaxies Actually Tell Us about Modified Newtonian Dynamics?” *Physical Review Letters* 108.14 (Apr. 5, 2012). DOI: 10.1103/PhysRevLett.108.141302. arXiv: 1108.5734 (cit. on p. 14).

[25] J. G. de Swart, G. Bertone, and J. van Dongen. “How Dark Matter Came to Matter”. *Nature Astronomy* 1.3 (Mar. 2, 2017), pages. DOI: 10.1038/s41550-017-0059 (cit. on p. 14).

[26] Wayne Hu. *CMB Intermediate*. 2001. URL: <http://background.uchicago.edu/~whu/intermediate/intermediate.html> (visited on 10/05/2017) (cit. on p. 15).

[27] Constantinos Skordis. “The Tensor-Vector-Scalar Theory and Its Cosmology”. *Class.Quant.Grav.* 26 (2009), p. 143001. DOI: 10.1088/0264-9381/26/14/143001 (cit. on p. 15).

[28] Wayne Hu and Scott Dodelson. “Cosmic Microwave Background Anisotropies”. *Annual Review of Astronomy and Astrophysics* 40.1 (Sept. 2002), pp. 171–216. DOI: 10.1146/annurev.astro.40.060401.093926. arXiv: astro-ph/0110414 (cit. on p. 14).

[29] Joel R. Primack. “Cosmological Structure Formation” (May 11, 2015). arXiv: 1505.02821 [astro-ph] (cit. on pp. 15, 16).

- [30] Scott Dodelson and Michele Liguori. “Can Cosmic Structure Form without Dark Matter?” *Physical Review Letters* 97.23 (Dec. 7, 2006). DOI: 10.1103/PhysRevLett.97.231301. arXiv: astro-ph/0608602 (cit. on p. 15).
- [31] Bernard Carr et al. “Primordial Black Hole Constraints for Extended Mass Functions”. *Physical Review D* 96.2 (July 14, 2017). DOI: 10.1103/PhysRevD.96.023514. arXiv: 1705.05567 (cit. on p. 17).
- [32] Simeon Bird et al. “Did LIGO Detect Dark Matter?” *Physical Review Letters* 116.20 (May 19, 2016). DOI: 10.1103/PhysRevLett.116.201301. arXiv: 1603.00464 (cit. on p. 17).
- [33] Katherine J. Mack, Jeremiah P. Ostriker, and Massimo Ricotti. “Growth of Structure Seeded by Primordial Black Holes”. *The Astrophysical Journal* 665.2 (2007), p. 1277. DOI: 10.1086/518998 (cit. on p. 17).
- [34] Sofiane M. Boucenna et al. “Novel Constraints on Mixed Dark-Matter Scenarios of Primordial Black Holes and WIMPs” (Dec. 18, 2017). arXiv: 1712.06383 [astro-ph] (cit. on p. 17).
- [35] Felix Kahlhoefer. “Review of LHC Dark Matter Searches”. *International Journal of Modern Physics A* 32.13 (May 10, 2017), p. 1730006. DOI: 10.1142/S0217751X1730006X. arXiv: 1702.02430 (cit. on pp. 17, 20).
- [36] Jan Conrad and Olaf Reimer. “Indirect Dark Matter Searches in Gamma- and Cosmic Rays”. *Nature Physics* 13.3 (Mar. 2, 2017), pp. 224–231. DOI: 10.1038/nphys4049. arXiv: 1705.11165 (cit. on p. 18).
- [37] Jennifer M. Gaskins. “A Review of Indirect Searches for Particle Dark Matter”. *Contemporary Physics* 57.4 (Oct. 2016), pp. 496–525. DOI: 10.1080/00107514.2016.1175160. arXiv: 1604.00014 (cit. on p. 18).
- [38] F. Mayet et al. “A Review of the Discovery Reach of Directional Dark Matter Detection”. *Physics Reports* 627 (Apr. 2016), pp. 1–49. DOI: 10.1016/j.physrep.2016.02.007. arXiv: 1602.03781 (cit. on p. 19).
- [39] C. E. Aalseth et al. “CoGeNT: A Search for Low-Mass Dark Matter Using p-Type Point Contact Germanium Detectors”. *Physical Review D* 88.1 (July 2013). DOI: 10.1103/PhysRevD.88.012002. arXiv: 1208.5737 [astro-ph, physics:hep-ex, physics:hep-ph, physics:nucl-ex, physics:physics] (cit. on p. 19).
- [40] G. Angloher et al. “Results from 730 Kg Days of the CRESST-II Dark Matter Search”. *The European Physical Journal C* 72.4 (Apr. 2012). DOI: 10.1140/epjc/s10052-012-1971-8. arXiv: 1109.0702 (cit. on p. 19).
- [41] CDMS Collaboration et al. “Silicon Detector Dark Matter Results from the Final Exposure of CDMS II” (Apr. 15, 2013). arXiv: 1304.4279 [astro-ph, physics:hep-ex, physics:physics] (cit. on p. 19).
- [42] R. Bernabei et al. “Final Model Independent Result of DAMA/LIBRA-Phase1”. *The European Physical Journal C* 73.12 (Dec. 2013). DOI: 10.1140/epjc/s10052-013-2648-7. arXiv: 1308.5109 (cit. on p. 19).
- [43] XENON100 Collaboration et al. “Dark Matter Results from 225 Live Days of XENON100 Data”. *Physical Review Letters* 109.18 (Nov. 2, 2012). DOI: 10.1103/PhysRevLett.109.181301. arXiv: 1207.5988 (cit. on p. 19).
- [44] LUX Collaboration et al. “First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility”. *Physical Review Letters* 112.9 (Mar. 4, 2014). DOI: 10.1103/PhysRevLett.112.091303. arXiv: 1310.8214 (cit. on pp. 19, 80).

[45] Jonathan H. Davis. “The Past and Future of Light Dark Matter Direct Detection”. *International Journal of Modern Physics A* 30.15 (May 30, 2015), p. 1530038. DOI: 10.1142/S0217751X15300380. arXiv: 1506.03924 (cit. on p. 19).

[46] The XENON Collaboration et al. “Exclusion of Leptophilic Dark Matter Models Using XENON100 Electronic Recoil Data”. *Science* 349.6250 (Aug. 21, 2015), pp. 851–854. DOI: 10.1126/science.aab2069. arXiv: 1507.07747 (cit. on pp. 19, 35).

[47] Benjamin D. Wandelt et al. “Self-Interacting Dark Matter” (June 23, 2000). arXiv: astro-ph/0006344 (cit. on p. 19).

[48] Jianglai Liu, Xun Chen, and Xiangdong Ji. “Current Status of Direct Dark Matter Detection Experiments”. *Nature Physics* 13.3 (Mar. 2017), pp. 212–216. DOI: 10.1038/nphys4039 (cit. on p. 19).

[49] Teresa Marrodan Undagoitia and Ludwig Rauch. “Dark Matter Direct-Detection Experiments”. *Journal of Physics G: Nuclear and Particle Physics* 43.1 (Jan. 1, 2016), p. 013001. DOI: 10.1088/0954-3899/43/1/013001. arXiv: 1509.08767 (cit. on pp. 19, 107).

[50] Guillaume Plante. “The XENON100 Dark Matter Experiment: Design, Construction, Calibration and 2010 Search Results with Improved Measurement of the Scintillation Response of Liquid Xenon to Low-Energy Nuclear Recoils”. New York: Columbia University, 2012. URL: <http://xenon.astro.columbia.edu/thesis/plante-thesis-201201.pdf> (cit. on p. 19).

[51] Richard J. Gaitskell. “Direct Detection of Dark Matter”. *Annual Review of Nuclear and Particle Science* 54.1 (2004), pp. 315–359. DOI: 10.1146/annurev.nucl.54.070103.181244 (cit. on p. 19).

[52] Gregory D. Mack, John F. Beacom, and Gianfranco Bertone. “Towards Closing the Window on Strongly Interacting Dark Matter: Far-Reaching Constraints from Earth’s Heat Flow”. *Physical Review D* 76.4 (Aug. 28, 2007). DOI: 10.1103/PhysRevD.76.043523. arXiv: 0705.4298 (cit. on p. 20).

[53] Jalal Abdallah et al. “Simplified Models for Dark Matter Searches at the LHC”. *Physics of the Dark Universe* 9-10 (Sept. 2015), pp. 8–23. DOI: 10.1016/j.dark.2015.08.001. arXiv: 1506.03116 (cit. on p. 20).

[54] Julien Lesgourgues and Sergio Pastor. “Cosmological Implications of a Relic Neutrino Asymmetry”. *Physical Review D* 60.10 (Oct. 27, 1999). DOI: 10.1103/PhysRevD.60.103521. arXiv: hep-ph/9904411 (cit. on p. 21).

[55] Lawrence J. Hall et al. “Freeze-In Production of FIMP Dark Matter”. *Journal of High Energy Physics* 2010.3 (Mar. 2010). DOI: 10.1007/JHEP03(2010)080. arXiv: 0911.1120 (cit. on p. 21).

[56] Nicolás Bernal et al. “The Dawn of FIMP Dark Matter: A Review of Models and Constraints” (June 22, 2017). arXiv: 1706.07442 [astro-ph] (cit. on p. 21).

[57] Kim Griest. “Unitarity Limits on the Mass and Radius of Dark-Matter Particles”. *Physical Review Letters* 64.6 (1990), pp. 615–618. DOI: 10.1103/PhysRevLett.64.615 (cit. on p. 21).

[58] Vid Iršič et al. “New Constraints on the Free-Streaming of Warm Dark Matter from Intermediate and Small Scale Lyman-\$\alpha\$ Forest Data”. *Physical Review D* 96.2 (July 19, 2017). DOI: 10.1103/PhysRevD.96.023522. arXiv: 1702.01764 (cit. on p. 21).

[59] Chi Man Ho and Robert J. Scherrer. “Limits on MeV Dark Matter from the Effective Number of Neutrinos”. *Physical Review D* 87.2 (Jan. 8, 2013). DOI: 10.1103/PhysRevD.87.023505. arXiv: 1208.4347 (cit. on p. 21).

- [60] Asher Berlin and Nikita Blinov. “Thermal Dark Matter Below an MeV” (June 21, 2017). arXiv: 1706.07046 [astro-ph, physics:hep-ph] (cit. on p. 21).
- [61] Tracy R. Slatyer. “Dark Matter Theory Overview”. July 24, 2017. URL: <https://indico.cern.ch/event/606690/sessions/222457/> (cit. on p. 21).
- [62] Tracy R. Slatyer. “Indirect Dark Matter Signatures in the Cosmic Dark Ages I. Generalizing the Bound on s-Wave Dark Matter Annihilation from Planck”. *Physical Review D* 93.2 (Jan. 25, 2016). doi: 10.1103/PhysRevD.93.023527. arXiv: 1506.03811 (cit. on p. 21).
- [63] Fermi-LAT Collaboration. “Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi-LAT Data”. *Physical Review Letters* 115.23 (Nov. 30, 2015). doi: 10.1103/PhysRevLett.115.231301. arXiv: 1503.02641 (cit. on pp. 22, 23).
- [64] Gary Steigman, Basudeb Dasgupta, and John F. Beacom. “Precise Relic WIMP Abundance and Its Impact on Searches for Dark Matter Annihilation”. *Physical Review D* 86.2 (July 3, 2012). doi: 10.1103/PhysRevD.86.023506. arXiv: 1204.3622 (cit. on p. 22).
- [65] Kim Griest. “Three Exceptions in the Calculation of Relic Abundances”. *Physical Review D* 43.10 (1991), pp. 3191–3203. doi: 10.1103/PhysRevD.43.3191 (cit. on p. 22).
- [66] Jonathan L. Feng. “Dark Matter at the Fermi Scale”. *Journal of Physics G: Nuclear and Particle Physics* 32.1 (Jan. 1, 2006), R1–R24. doi: 10.1088/0954-3899/32/1/R01. arXiv: astro-ph/0511043 (cit. on p. 22).
- [67] Miguel Escudero et al. “Toward (Finally!) Ruling Out Z and Higgs Mediated Dark Matter Models”. *Journal of Cosmology and Astroparticle Physics* 2016.12 (Dec. 15, 2016), pp. 029–029. doi: 10.1088/1475-7516/2016/12/029. arXiv: 1609.09079 (cit. on pp. 23–25).
- [68] The XENON collaboration et al. “Physics Reach of the XENON1T Dark Matter Experiment”. *Journal of Cosmology and Astroparticle Physics* 2016 (04 Apr. 14, 2016), pp. 027–027. doi: 10.1088/1475-7516/2016/04/027. arXiv: 1512.07501 (cit. on pp. 23, 44, 107, 112, 115, 139).
- [69] Giorgio Arcadi et al. “The Waning of the WIMP? A Review of Models, Searches, and Constraints” (Mar. 21, 2017). arXiv: 1703.07364 [astro-ph] (cit. on p. 24).
- [70] The GAMBIT Collaboration et al. “A Global Fit of the MSSM with GAMBIT” (May 22, 2017). arXiv: 1705.07917 [astro-ph] (cit. on p. 24).
- [71] The GAMBIT Collaboration et al. “Global Fits of GUT-Scale SUSY Models with GAMBIT” (May 22, 2017). arXiv: 1705.07935 [astro-ph] (cit. on p. 24).
- [72] J.D. Lewin and P.F. Smith. “Review of Mathematics, Numerical Factors, and Corrections for Dark Matter Experiments Based on Elastic Nuclear Recoil”. *Astroparticle Physics* 6.1 (Dec. 1996), pp. 87–112. doi: 10.1016/S0927-6505(96)00047-3 (cit. on pp. 28, 30, 33, 116).
- [73] Anne M. Green. “Astrophysical Uncertainties on the Local Dark Matter Distribution and Direct Detection Experiments”. *Journal of Physics G: Nuclear and Particle Physics* 44.8 (Aug. 1, 2017), p. 084001. doi: 10.1088/1361-6471/aa7819. arXiv: 1703.10102 (cit. on pp. 29, 30).
- [74] Jonathan D. Sloane et al. “Assessing Astrophysical Uncertainties in Direct Detection with Galaxy Simulations”. *The Astrophysical Journal* 831.1 (Oct. 28, 2016), p. 93. doi: 10.3847/0004-637X/831/1/93. arXiv: 1601.05402 (cit. on p. 30).

[75] JiJi Fan, Matthew Reece, and Lian-Tao Wang. “Non-Relativistic Effective Theory of Dark Matter Direct Detection”. *Journal of Cosmology and Astroparticle Physics* 2010.11 (Nov. 29, 2010), pp. 042–042. doi: 10.1088/1475-7516/2010/11/042. arXiv: 1008.1591 (cit. on p. 32).

[76] Spencer Chang, Aaron Pierce, and Neal Weiner. “Momentum Dependent Dark Matter Scattering”. *Journal of Cosmology and Astroparticle Physics* 2010 (01 Jan. 4, 2010), pp. 006–006. doi: 10.1088/1475-7516/2010/01/006. arXiv: 0908.3192 (cit. on p. 32).

[77] Tai Li, Sen Miao, and Yu-Feng Zhou. “Light Mediators in Dark Matter Direct Detections”. *Journal of Cosmology and Astroparticle Physics* 2015 (03 Mar. 17, 2015), pp. 032–032. doi: 10.1088/1475-7516/2015/03/032. arXiv: 1412.6220 (cit. on p. 32).

[78] Eugenio Del Nobile, Manoj Kaplinghat, and Hai-Bo Yu. “Direct Detection Signatures of Self-Interacting Dark Matter with a Light Mediator”. *Journal of Cosmology and Astroparticle Physics* 2015.10 (Oct. 27, 2015), pp. 055–055. doi: 10.1088/1475-7516/2015/10/055. arXiv: 1507.04007 (cit. on p. 32).

[79] A. Liam Fitzpatrick et al. “The Effective Field Theory of Dark Matter Direct Detection”. *Journal of Cosmology and Astroparticle Physics* 2013 (02 Feb. 5, 2013), pp. 004–004. doi: 10.1088/1475-7516/2013/02/004. arXiv: 1203.3542 (cit. on p. 32).

[80] E. Aprile et al. “Effective Field Theory Search for High-Energy Nuclear Recoils Using the XENON100 Dark Matter Detector”. *Physical Review D* 96.4 (Aug. 31, 2017). doi: 10.1103/PhysRevD.96.042004. arXiv: 1705.02614 (cit. on p. 32).

[81] David G. Cerdeno et al. “Nuclear Uncertainties in the Spin-Dependent Structure Functions for Direct Dark Matter Detection”. *Physical Review D* 87.2 (Jan. 18, 2013). doi: 10.1103/PhysRevD.87.023512. arXiv: 1208.6426 (cit. on pp. 33, 148).

[82] J. Menendez, D. Gazit, and A. Schwenk. “Spin-Dependent WIMP Scattering off Nuclei”. *Physical Review D* 86.10 (Nov. 7, 2012). doi: 10.1103/PhysRevD.86.103511. arXiv: 1208.1094 (cit. on p. 33).

[83] XENON100 collaboration et al. “Limits on Spin-Dependent WIMP-Nucleon Cross Sections from 225 Live Days of XENON100 Data”. *Physical Review Letters* 111.2 (July 2013). doi: 10.1103/PhysRevLett.111.021301. arXiv: 1301.6620 [astro-ph] (cit. on pp. 33, 34).

[84] P. Klos et al. “Large-Scale Nuclear Structure Calculations for Spin-Dependent WIMP Scattering with Chiral Effective Field Theory Currents”. *Physical Review D* 88.8 (Oct. 17, 2013), p. 083516. doi: 10.1103/PhysRevD.88.083516 (cit. on pp. 33, 147, 148).

[85] Christopher Savage et al. “LUX Likelihood and Limits on Spin-Independent and Spin-Dependent WIMP Couplings with LUXCalc” (Feb. 9, 2015). arXiv: 1502.02667 [astro-ph] (cit. on pp. 33, 147).

[86] LUX Collaboration et al. “Limits on Spin-Dependent WIMP-Nucleon Cross Section Obtained from the Complete LUX Exposure”. *Physical Review Letters* 118.25 (June 23, 2017). doi: 10.1103/PhysRevLett.118.251302. arXiv: 1705.03380 (cit. on pp. 33, 147, 149, 150).

[87] Joachim Kopp et al. “DAMA/LIBRA and Leptonically Interacting Dark Matter”. *Physical Review D* 80.8 (Oct. 2, 2009). doi: 10.1103/PhysRevD.80.083502. arXiv: 0907.3159 (cit. on pp. 34, 35).

[88] Rouven Essig et al. “First Direct Detection Limits on Sub-GeV Dark Matter from XENON10”. *Physical Review Letters* 109.2 (July 12, 2012). doi: 10.1103/PhysRevLett.109.021301. arXiv: 1206.2644 (cit. on p. 34).

[89] Rouven Essig, Tomer Volansky, and Tien-Tien Yu. “New Constraints and Prospects for Sub-GeV Dark Matter Scattering off Electrons in Xenon” (Mar. 2, 2017). arXiv: 1703.00910 [astro-ph] (cit. on p. 34).

[90] Joachim Kopp et al. “Leptophilic Dark Matter in Direct Detection Experiments and in the Sun” (Nov. 5, 2010). arXiv: 1011.1398 [hep-ph] (cit. on p. 35).

[91] John Ellis, R. A. Flores, and J. D. Lewin. “Rates for Inelastic Nuclear Excitation by Dark Matter Particles”. *Physics Letters B* 212.3 (Sept. 29, 1988), pp. 375–380. doi: 10.1016/0370-2693(88)91332-9 (cit. on p. 36).

[92] L. Baudis et al. “Signatures of Dark Matter Scattering Inelastically Off Nuclei”. *Physical Review D* 88.11 (Dec. 26, 2013). doi: 10.1103/PhysRevD.88.115014. arXiv: 1309.0825 (cit. on p. 36).

[93] Christopher McCabe. “Prospects for Dark Matter Detection with Inelastic Transitions of Xenon”. *Journal of Cosmology and Astroparticle Physics* 2016 (05 May 16, 2016), pp. 033–033. doi: 10.1088/1475-7516/2016/05/033. arXiv: 1512.00460 (cit. on p. 36).

[94] David Tucker-Smith. “Status of Inelastic Dark Matter”. *Physical Review D* 72.6 (2005). doi: 10.1103/PhysRevD.72.063509 (cit. on p. 36).

[95] E. Aprile et al. “Implications on Inelastic Dark Matter from 100 Live Days of XENON100 Data”. *Physical Review D* 84.6 (Sept. 6, 2011), p. 061101. doi: 10.1103/PhysRevD.84.061101 (cit. on p. 36).

[96] Chris Kouvaris and Josef Pradler. “Probing Sub-GeV Dark Matter with Conventional Detectors”. *Physical Review Letters* 118.3 (Jan. 20, 2017). doi: 10.1103/PhysRevLett.118.031803. arXiv: 1607.01789 (cit. on pp. 36, 37, 143, 150, 151, 170, 171).

[97] Christopher McCabe. “New Constraints and Discovery Potential of Sub-GeV Dark Matter with Xenon Detectors” (Feb. 15, 2017). arXiv: 1702.04730 [astro-ph] (cit. on pp. 36, 150–153).

[98] A. Kramida et al. *NIST: Atomic Spectra Database - Ionization Energies Form*. 2017. URL: <https://physics.nist.gov/PhysRefData/ASD/ionEnergy.html> (visited on 11/09/2017) (cit. on p. 37).

[99] Vitaly Chepel and Henrique Araújo. “Liquid Noble Gas Detectors for Low Energy Particle Physics”. *Journal of Instrumentation* 8 (04 Apr. 4, 2013), R04001–R04001. doi: 10.1088/1748-0221/8/04/R04001. arXiv: 1207.2292 [astro-ph, physics:hep-ex, physics:physics] (cit. on pp. 38, 40, 47).

[100] Jeremy Mock et al. “Modeling Pulse Characteristics in Xenon with NEST”. *Journal of Instrumentation* 9 (04 Apr. 3, 2014), T04002–T04002. doi: 10.1088/1748-0221/9/04/T04002. arXiv: 1310.1117 (cit. on pp. 38, 71, 76, 79, 83).

[101] Brian Lenardo. “Liquid Xenon Scintillation Measurements and Pulse Shape Discrimination in the LUX Dark Matter Detector”. June 15, 2017. URL: http://www.creteconf.org/presentations/Lenardo_Crete17.pdf (visited on 10/30/2017) (cit. on pp. 38, 84).

[102] Darryl Masson. *Convection Mapping in Xe1T with 220 RnPo Coincidence*. Jan. 11, 2017. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:darryl:convection_1701 (visited on 01/30/2018) (cit. on p. 38).

[103] Oliver Hilt and Werner F. Schmidt. “Positive Hole Mobility in Liquid Xenon”. *Chemical Physics* 183.1 (May 15, 1994), pp. 147–153. doi: 10.1016/0301-0104(94)00051-4 (cit. on p. 38).

[104] James Ziegler. *SRIM & TRIM*. URL: <http://srim.org/> (visited on 11/02/2017) (cit. on p. 38).

[105] A. Mozumder. “Free-Ion Yield and Electron-Ion Recombination Rate in Liquid Xenon”. *Chemical Physics Letters* 245 (Nov. 1, 1995), pp. 359–363. doi: 10.1016/0009-2614(95)01024-4 (cit. on p. 38).

[106] NIST. *ESTAR: Stopping-Power and Range Tables for Electrons*. URL: <http://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html> (visited on 01/26/2017) (cit. on p. 38).

[107] Akira Hitachi et al. “Effect of Ionization Density on the Time Dependence of Luminescence from Liquid Argon and Xenon”. *Physical Review B* 27.9 (May 1, 1983), pp. 5279–5285. doi: 10.1103/PhysRevB.27.5279 (cit. on p. 40).

[108] J. I. Collar and D. N. McKinsey. “Comments on ”First Dark Matter Results from the XENON100 Experiment”” (May 5, 2010). arXiv: 1005.0838 [astro-ph, physics:hep-ex, physics:hep-ph, physics:physics] (cit. on p. 40).

[109] LUX Collaboration et al. “Low-Energy (0.7–74 keV) Nuclear Recoil Calibration of the LUX Dark Matter Experiment Using D-D Neutron Scattering Kinematics” (Aug. 18, 2016). arXiv: 1608.05381 [astro-ph] (cit. on pp. 40, 45, 156, 169).

[110] Brian Lenardo et al. “A Global Analysis of Light and Charge Yields in Liquid Xenon”. *IEEE Transactions on Nuclear Science* 62.6 (Dec. 2015), pp. 3387–3396. doi: 10.1109/TNS.2015.2481322. arXiv: 1412.4417 (cit. on pp. 40, 41, 113, 136, 141, 142, 144, 156).

[111] LUX Collaboration et al. “Ultra-Low Energy Calibration of LUX Detector Using $\hat{\{127\}}\text{Xe}$ Electron Capture” (Sept. 3, 2017). arXiv: 1709.00800 [physics] (cit. on pp. 40, 42, 123, 150).

[112] LUX Collaboration et al. “Tritium Calibration of the LUX Dark Matter Experiment” (Dec. 9, 2015). arXiv: 1512.03133 [astro-ph, physics:hep-ex, physics:physics] (cit. on pp. 42, 44, 112, 122, 150).

[113] L. W. Goetzke et al. “Measurement of Light and Charge Yield of Low-Energy Electronic Recoils in Liquid Xenon” (Nov. 30, 2016). arXiv: 1611.10322 [astro-ph, physics:physics] (cit. on pp. 40–42, 124, 150).

[114] M. Szydagis et al. “NEST: A Comprehensive Model for Scintillation Yield in Liquid Xenon” (June 8, 2011). arXiv: 1106.1613 [physics] (cit. on p. 40).

[115] KamLAND-Zen Collaboration. “Measurement of the Double- β Decay Half-Life of $\hat{\{136\}}\text{Xe}$ with the KamLAND-Zen Experiment”. *Physical Review C* 85.4 (Apr. 19, 2012). doi: 10.1103/PhysRevC.85.045504. arXiv: 1201.4664 (cit. on p. 44).

[116] R. F. Lang et al. “Characterization of a Deuterium-Deuterium Plasma Fusion Neutron Generator” (May 12, 2017). arXiv: 1705.04741 [hep-ex, physics:nucl-ex, physics:physics] (cit. on p. 45).

[117] Curt Suplee. *XCOM: Photon Cross Sections Database*. 2009-09-17T17:33-04:00. URL: <https://www.nist.gov/pml/xcom-photon-cross-sections-database> (visited on 08/16/2017) (cit. on pp. 45, 47).

[118] Dave Brown. *Evaluated Nuclear Data File (ENDF) Retrieval & Plotting*. URL: <http://www.nndc.bnl.gov/sigma/index.jsp?as=132&lib=endfb7.1&nsub=10> (visited on 11/13/2017) (cit. on p. 45).

[119] D. Akimov et al. “Observation of Coherent Elastic Neutrino-Nucleus Scattering”. *Science* 357.6356 (Sept. 15, 2017), pp. 1123–1126. doi: 10.1126/science.aoa0990. arXiv: 1708.01294 (cit. on p. 46).

[120] M. Schumacher. “Photonuclear Reactions”. *Journal of Physics G: Nuclear Physics* 14 (S 1988), S235. doi: 10.1088/0305-4616/14/S/026 (cit. on p. 47).

- [121] Bernd Crasemann. *Atomic Inner-Shell Physics*. Springer Science & Business Media, Mar. 9, 2013. 760 pp. (cit. on p. 47).
- [122] R. E. Mikkelsen, A. H. Sørensen, and U. I. Uggerhøj. “Elastic Photonuclear Cross Sections for Bremsstrahlung from Relativistic Ions”. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 372 (Apr. 2016), pp. 58–66. doi: 10.1016/j.nimb.2016.02.003. arXiv: 1510.02600 (cit. on p. 47).
- [123] K. P. Schelhaas et al. “Nuclear Photon Scattering by ^{208}Pb ”. *Nuclear Physics A* 489 (Nov. 1, 1988), pp. 189–224. doi: 10.1016/0375-9474(88)90149-2 (cit. on p. 47).
- [124] *Thermal Neutron Capture Gamma Rays*. URL: <https://www-nds.iaea.org/capgam/indexbyn.htmlx> (visited on 08/16/2017) (cit. on p. 47).
- [125] Tenzing H.Y. Joshi. “Producing Nuclear Recoils”. Dec. 7, 2012. URL: https://neutrinos.llnl.gov/cns_workshop/Joshi_CNNWorkshop-LLNL-PRES-606757.pdf (cit. on p. 47).
- [126] E. Aprile et al. “The XENON1T Dark Matter Experiment”. *Eur.Phys.J.* C77 (Dec. 18, 2017), p. 881. doi: 10.1140/epjc/s10052-017-5326-3 (cit. on pp. 47, 48, 54, 97).
- [127] Elena Aprile et al. *Noble Gas Detectors*. Weinheim: WILEY-VCH Verlag GmbH & Co. KGaA, 2006. URL: <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-3527405976.html> (cit. on pp. 47, 76, 79).
- [128] P. Barrow et al. “Qualification Tests of the R11410-21 Photomultiplier Tubes for the XENON1T Detector”. *Journal of Instrumentation* 12 (01 Jan. 30, 2017), P01024–P01024. doi: 10.1088/1748-0221/12/01/P01024. arXiv: 1609.01654 (cit. on pp. 49, 65, 85, 108).
- [129] E. Aprile et al. “Lowering the Radioactivity of the Photomultiplier Tubes for the XENON1T Dark Matter Experiment”. *The European Physical Journal C* 75.11 (Nov. 2015). doi: 10.1140/epjc/s10052-015-3657-5. arXiv: 1503.07698 (cit. on pp. 49, 108).
- [130] XENON Collaboration et al. “Material Radioassay and Selection for the XENON1T Dark Matter Experiment” (May 4, 2017). arXiv: 1705.01828 [physics] (cit. on p. 49).
- [131] E. Aprile et al. “Conceptual Design and Simulation of a Water Cherenkov Muon Veto for the XENON1T Experiment” (June 9, 2014). arXiv: 1406.2374 [astro-ph, physics:physics] (cit. on pp. 49, 108).
- [132] XENON Collaboration et al. “Removing Krypton from Xenon by Cryogenic Distillation to the Ppq Level”. *The European Physical Journal C* 77.5 (May 2017). doi: 10.1140/epjc/s10052-017-4757-1. arXiv: 1612.04284 (cit. on pp. 49, 108).
- [133] XENON Collaboration et al. “Online $\hat{222}\text{Rn}$ Removal by Cryogenic Distillation in the XENON100 Experiment”. *The European Physical Journal C* 77.6 (June 2017). doi: 10.1140/epjc/s10052-017-4902-x. arXiv: 1702.06942 (cit. on pp. 49, 108).
- [134] Tadayoshi Doke et al. “Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles”. *Japanese Journal of Applied Physics* 41 (3R Mar. 2002), p. 1538. doi: 10.1143/JJAP.41.1538 (cit. on p. 49).
- [135] Ye Jingqiang et al. *Summary of the G1, G2 Investigation from Mono-Energetic Lines*. Mar. 14, 2017. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon-xenonit:analysis:subgroup:energyscale:g1g2firstresult_summarynote#electron_extraction_efficiency (visited on 01/29/2018) (cit. on pp. 50, 70).

[136] B. Edwards et al. “Measurement of Single Electron Emission in Two-Phase Xenon”. *Astroparticle Physics* 30.2 (Sept. 2008), pp. 54–57. DOI: 10.1016/j.astropartphys.2008.06.006. arXiv: 0708.0768 (cit. on pp. 54, 68, 108).

[137] E. Aprile et al. “Observation and Applications of Single-Electron Charge Signals in the XENON100 Experiment”. *Journal of Physics G: Nuclear and Particle Physics* 41.3 (Mar. 1, 2014), p. 035201. DOI: 10.1088/0954-3899/41/3/035201. arXiv: 1311.1088 [astro-ph, physics:hep-ex, physics:physics] (cit. on pp. 54, 56, 68, 108).

[138] Hamamatsu. *Photomultiplier Tubes: Basics and Applications (Edition 3a)*. 2007. URL: https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf (visited on 10/26/2017) (cit. on pp. 56, 65).

[139] D. Yu Akimov et al. “Observation of Delayed Electron Emission in a Two-Phase Liquid Xenon Detector”. *Journal of Instrumentation* 11 (03 2016), p. C03007. DOI: 10.1088/1748-0221/11/03/C03007 (cit. on p. 56).

[140] Peter Sorensen. “Electron Train Backgrounds in Liquid Xenon Dark Matter Search Detectors Are Indeed Due to Thermalization and Trapping” (Feb. 15, 2017). arXiv: 1702.04805 [astro-ph, physics:physics] (cit. on pp. 56, 68, 114, 161).

[141] XENON100 Collaboration et al. “The XENON100 Dark Matter Experiment”. *Astroparticle Physics* 35.9 (Apr. 2012), pp. 573–590. DOI: 10.1016/j.astropartphys.2012.01.003. arXiv: 1107.2155 [astro-ph] (cit. on pp. 57, 107).

[142] Dominic Cichon. *XENON1T Multi-Event Noise Analysis*. Mar. 2016. URL: https://xe1t-wiki.lngs.infn.it/lib/exe/fetch.php?media=xenon:xenon1t:analysis:dominick:xenon1t_multi-event_noise.html (visited on 01/29/2018) (cit. on p. 57).

[143] Richard Saldanha. *Filter Box for CAEN PMT HV Supply*. Sept. 2016. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:saldanha:xe1t:pmt_hv_filter_box (visited on 01/29/2018) (cit. on p. 57).

[144] Alex Kish and Michelle Galloway. *Private Communication*. Nov. 27, 2017 (cit. on p. 60).

[145] Daniel Mayani. “Photomultiplier Tubes for the XENON1T Dark Matter Experiment and Studies on the XENON100 Electromagnetic Background”. Zurich: University of Zurich, 2017. URL: http://www.physik.uzh.ch/groups/groupbaudis/darkmatter/theses/xenon/PhD_thesis_Mayani.pdf (cit. on p. 60).

[146] M. Anthony et al. “Characterization of Photomultiplier Tubes with a Realistic Model through GPU-Boosted Simulation” (Sept. 30, 2017). arXiv: 1710.00219 [astro-ph] (cit. on p. 60).

[147] R. Saldanha et al. “Model Independent Approach to the Single Photoelectron Calibration of Photomultiplier Tubes”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 863 (Aug. 2017), pp. 35–46. DOI: 10.1016/j.nima.2017.02.086. arXiv: 1602.03150 (cit. on pp. 60, 108).

[148] C. H. Faham et al. “Measurements of Wavelength-Dependent Double Photoelectron Emission from Single Photons in VUV-Sensitive Photomultiplier Tubes”. *Journal of Instrumentation* 10 (09 Sept. 9, 2015), P09010–P09010. DOI: 10.1088/1748-0221/2015/9/P09010. arXiv: 1506.08748 (cit. on pp. 62, 64).

[149] Hamamatsu. *R11410-21 Datasheet (XENOn1T, Rev. 3)*. Sept. 30, 2014. URL: https://xe1t-wiki.lngs.infn.it/lib/exe/fetch.php?media=xenon:xenon1t:dsg:sensors:r11410_21_datasheet_140930.pdf (cit. on pp. 65, 84).

- [150] H. L. Brooks et al. “Electron Drift Velocities in Xenon”. *Journal of Physics D: Applied Physics* 15.6 (June 14, 1982), p. L51. doi: 10.1088/0022-3727/15/6/002 (cit. on p. 73).
- [151] C. M. B. Monteiro et al. “Secondary Scintillation Yield in Pure Xenon”. *Journal of Instrumentation* 2 (05 2007), P05001. doi: 10.1088/1748-0221/2/05/P05001 (cit. on p. 73).
- [152] Christopher Geis. *Detector Leveling Using S2 Width and S2 Width vs. Liquid Level*. Oct. 10, 2016. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:analysis:chg:s2width_leveling (visited on 01/29/2018) (cit. on p. 74).
- [153] P. A. Breur. “Backgrounds in XENON1T”. Amsterdam: University of Amsterdam, 2018 (in preparation) (cit. on p. 78).
- [154] V. N. Lebedenko et al. “Results from the First Science Run of the ZEPLIN-III Dark Matter Search Experiment”. *Physical Review D* 80.5 (Sept. 25, 2009). doi: 10.1103/PhysRevD.80.052010. arXiv: 0812.1150 (cit. on p. 80).
- [155] PandaX-II Collaboration et al. “Dark Matter Results from First 98.7-Day Data of PandaX-II Experiment”. *Physical Review Letters* 117.12 (Sept. 16, 2016). doi: 10.1103/PhysRevLett.117.121303. arXiv: 1607.07400 (cit. on pp. 80, 115).
- [156] Andrea Tiseni. *Private Communication*. E-mail. May 10, 2016 (cit. on p. 80).
- [157] Peter Sorensen. “Anisotropic Diffusion of Electrons in Liquid Xenon with Application to Improving the Sensitivity of Direct Dark Matter Searches” (Feb. 14, 2011). *Nucl. Instr. Meth. A* 635 (2011) 41-43. arXiv: 1102.2865 [astro-ph] (cit. on pp. 79, 80).
- [158] S N. Anisimov, A Bolozdynya, and V Stekhanov. “Electron Localization and Drift under the Surface of Condensed Krypton”. *Soviet Journal of Experimental and Theoretical Physics Letters* (Aug. 1, 1984) (cit. on p. 79).
- [159] V. N. Solovov et al. “Measurement of the Refractive Index and Attenuation Length of Liquid Xenon for Its Scintillation Light”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 516 (2-3 Jan. 2004), pp. 462–474. doi: 10.1016/j.nima.2003.08.117. arXiv: physics/0307044 (cit. on p. 84).
- [160] M. Moongweluhan. “The Impact of Photon Flight Path on S1 Pulse Shape Analysis in Liquid Xenon Two-Phase Detectors” (Nov. 13, 2015). arXiv: 1511.04501 [hep-ex, physics:physics] (cit. on p. 84).
- [161] J. Kwong et al. “Scintillation Pulse Shape Discrimination in a Two-Phase Xenon Time Projection Chamber”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 612.2 (Jan. 2010), pp. 328–333. doi: 10.1016/j.nima.2009.10.106. arXiv: 0908.0790 (cit. on p. 84).
- [162] K. Ueshima et al. “Scintillation-Only Based Pulse Shape Discrimination for Nuclear and Electron Recoils in Liquid Xenon” (June 11, 2011). arXiv: 1106.2209 [hep-ex] (cit. on p. 84).
- [163] Daniel Coderre, Cyril Grignon, and Christopher Tunnell. *XENON1T/Kodiaq: Production Release*. Zenodo, Nov. 8, 2016. doi: 10.5281/zenodo.165486. URL: <https://zenodo.org/record/165486#.WZWplidLdhE> (visited on 08/17/2017) (cit. on p. 90).
- [164] MongoDB Inc. *MongoDB*. 2018. URL: <https://www.mongodb.com> (visited on 01/22/2018) (cit. on p. 90).
- [165] Pivotal Software. *RabbitMQ*. 2018. URL: <https://www.rabbitmq.com/> (visited on 01/22/2018) (cit. on pp. 91, 101).

[166] Stefan Shindler and Andrew Brown. *Software Trigger Efficiency for XENON1T*. Mar. 2017. URL: https://xe1t-wiki.lngs.infn.it/lib/exe/fetch.php?media=xenon:xenon1t:shindler:low_energy_event_loss_v2.html (visited on 01/25/2018) (cit. on pp. 99, 100).

[167] A. J. Anderson et al. “Coherent Neutrino Scattering in Dark Matter Detectors”. *Physical Review D* 84.1 (July 2011). doi: 10.1103/PhysRevD.84.013008. arXiv: 1103.4894 [hep-ex, physics:hep-ph] (cit. on p. 104).

[168] L. Baudis et al. “Neutrino Physics with Multi-Ton Scale Liquid Xenon Detectors”. *Journal of Cosmology and Astroparticle Physics* 2014 (01 Jan. 28, 2014), pp. 044–044. doi: 10.1088/1475-7516/2014/01/044. arXiv: 1309.7024 [astro-ph] (cit. on pp. 104, 157).

[169] Roland Diehl et al. “Radioactive 26Al from Massive Stars in the Galaxy”. *Nature* 439.7072 (Jan. 5, 2006), pp. 45–47. doi: 10.1038/nature04364 (cit. on p. 104).

[170] Rafael F. Lang et al. “Supernova Neutrino Physics with Xenon Dark Matter Detectors: A Timely Perspective”. *Physical Review D* 94.10 (Nov. 23, 2016). doi: 10.1103/PhysRevD.94.103009. arXiv: 1606.09243 (cit. on p. 104).

[171] P. Antonioli. “SNEWS: The SuperNova Early Warning System”. *New Journal of Physics* 6 (Sept. 2, 2004), pp. 114–114. doi: 10.1088/1367-2630/6/1/114. arXiv: astro-ph/0406214 (cit. on p. 105).

[172] Scott Dodelson. *Modern Cosmology*. Amsterdam: Academic Press, 2003 (cit. on p. 107).

[173] Edward Kolb and Michael Turner. *The Early Universe*. Frontiers in physics 69. 1990 (cit. on p. 107).

[174] Mark W. Goodman and Edward Witten. “Detectability of Certain Dark Matter Candidates”. *Phys.Rev.* D31 (1985), p. 3059. doi: 10.1103/PhysRevD.31.3059 (cit. on p. 107).

[175] E. Aprile et al. “Design and Performance of the XENON10 Dark Matter Experiment”. *Astroparticle Physics* 34.9 (Apr. 2011), pp. 679–698. doi: 10.1016/j.astropartphys.2011.01.006. arXiv: 1001.2834 (cit. on p. 107).

[176] Sebastian Lindemann and Hardy Simgen. “Krypton Assay in Xenon at the Ppq Level Using a Gas Chromatographic System and Mass Spectrometer”. *The European Physical Journal C* 74.2 (Feb. 2014). doi: 10.1140/epjc/s10052-014-2746-1. arXiv: 1308.4806 (cit. on p. 108).

[177] The XENON Collaboration et al. “Results from a Calibration of XENON100 Using a Source of Dissolved Radon-220”. *Physical Review D* 95.7 (Apr. 11, 2017). doi: 10.1103/PhysRevD.95.072008. arXiv: 1611.03585 (cit. on p. 110).

[178] Eric Dahl. “The Physics of Background Discrimination in Liquid Xenon, and First Results from Xenon10 in the Hunt for WIMP Dark Matter”. Princeton University, 2009. URL: <https://www.princeton.edu/physics/graduate-program/theses/theses-from-2009/E.Dahlthesis.pdf> (cit. on p. 110).

[179] Brian Lenardo et al. “A Global Analysis of Light and Charge Yields in Liquid Xenon” (Dec. 14, 2014). arXiv: 1412.4417 [astro-ph, physics:physics] (cit. on p. 113).

[180] XENON100 Collaboration et al. “XENON100 Dark Matter Results from a Combination of 477 Live Days” (Sept. 20, 2016). arXiv: 1609.06154 [astro-ph] (cit. on pp. 113–115).

[181] D. S. Akerib et al. “Results from a Search for Dark Matter in LUX with 332 Live Days of Exposure” (Aug. 26, 2016). arXiv: 1608.07648 [astro-ph] (cit. on p. 115).

[182] Glen Cowan et al. “Asymptotic Formulae for Likelihood-Based Tests of New Physics”. *The European Physical Journal C* 71.2 (Feb. 2011), pp. 1–19. doi: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [hep-ex]. Erratum-ibid. (Cit. on pp. 116, 119, 132).

[183] Nadav Priel et al. “A Model Independent Safeguard for Unbinned Likelihood”. *Journal of Cosmology and Astroparticle Physics* 2017 (05 May 4, 2017), pp. 013–013. doi: 10.1088/1475-7516/2017/05/013. arXiv: 1610.02643 (cit. on pp. 116, 129).

[184] Glen Cowan et al. “Power-Constrained Limits” (May 16, 2011). arXiv: 1105.3166 [hep-ex, physics:physics] (cit. on pp. 116, 120).

[185] Wouter Verkerke. “Nikhef Topical Lectures on Statistics”. 2016 (cit. on p. 118).

[186] Arjen Wildeboer. “Threshold Optimization and Bayesian Inference on the XENON1T Experiment”. Amsterdam: University of Amsterdam, Aug. 2017 (cit. on p. 119).

[187] Daniel Foreman-Mackey et al. “Emcee: The MCMC Hammer”. *Publications of the Astronomical Society of the Pacific* 125.925 (Mar. 2013), pp. 306–312. doi: 10.1086/670067. arXiv: 1202.3665 (cit. on p. 120).

[188] Atilla Dobi. “Measurement of the Electron Recoil Band of the LUX Dark Matter Detector with a Tritium Calibration Source”. College Park: University of Maryland, 2014. URL: <http://lux.brown.edu/papers/mainthesis-AD.pdf> (visited on 12/18/2017) (cit. on p. 122).

[189] Qing Lin. *Brief Summary of ER and Flat Background Modeling in Xe1T SR0*. Mar. 31, 2017. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:mcfate:summary_of_er_background_model_xe1t (visited on 01/27/2018) (cit. on p. 125).

[190] S. Yellin. “Finding an Upper Limit in the Presence of Unknown Background”. *Physical Review D* 66.3 (Aug. 2002). doi: 10.1103/PhysRevD.66.032005. arXiv: physics/0203002 (cit. on pp. 129, 155).

[191] J. Neyman. “Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability”. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 236.767 (1937), pp. 333–380. JSTOR: 91337 (cit. on p. 132).

[192] P.A. Breur et al. *Summary Note on E- Lifetime Correction for Sciencerun 0*. Feb. 12, 2017. URL: <https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:analysis:firstresults:elife> (visited on 01/27/2018) (cit. on p. 140).

[193] LUX Collaboration et al. “Improved WIMP Scattering Limits from the LUX Experiment” (Dec. 10, 2015). arXiv: 1512.03506 [astro-ph, physics:hep-ex, physics:physics] (cit. on p. 142).

[194] Shingo Kazama. *Anomalous Leakage BG in SR1 Rn220 Calibration Data*. Feb. 7, 2018. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:analysis:sciencerun1:anomalous_background_update (visited on 02/12/2018) (cit. on p. 145).

[195] PandaX-II Collaboration et al. “Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment”. *Physical Review Letters* 119.18 (Oct. 30, 2017). doi: 10.1103/PhysRevLett.119.181302. arXiv: 1708.06917 (cit. on p. 146).

[196] D. S. Akerib et al. “Results from a Search for Dark Matter in the Complete LUX Exposure”. *Physical Review Letters* 118.2 (Jan. 11, 2017). doi: 10.1103/PhysRevLett.118.021303. arXiv: 1608.07648 (cit. on pp. 146, 153, 166).

[197] C. Amole et al. “Dark Matter Search Results from the PICO-60 C\$-3\$F\$-8\$ Bubble Chamber”. *Physical Review Letters* 118.25 (June 23, 2017). doi: 10.1103/PhysRevLett.118.251301. arXiv: 1702.07666 (cit. on p. 150).

[198] C.T. Chantler et al. *X-Ray Form Factor, Attenuation, and Scattering Tables*. 2005. URL: <http://physics.nist.gov/ffast> (visited on 02/01/2018) (cit. on p. 150).

[199] CRESST collaboration et al. “First Results on Low-Mass Dark Matter from the CRESST-III Experiment” (Nov. 21, 2017). arXiv: 1711.07692 [astro-ph] (cit. on pp. 151, 153).

[200] G. Angloher et al. “Results on MeV-Scale Dark Matter from a Gram-Scale Cryogenic Calorimeter Operated above Ground”. *The European Physical Journal C* 77.9 (Sept. 2017). doi: 10.1140/epjc/s10052-017-5223-9. arXiv: 1707.06749 (cit. on pp. 153, 170).

[201] CRESST Collaboration et al. “Results on Low Mass WIMPs Using an Upgraded CRESST-II Detector”. *The European Physical Journal C* 74.12 (Dec. 2014). doi: 10.1140/epjc/s10052-014-3184-9. arXiv: 1407.3146 (cit. on p. 151).

[202] J. Angle et al. “A Search for Light Dark Matter in XENON10 Data”. *Physical Review Letters* 107.5 (July 2011). doi: 10.1103/PhysRevLett.107.051301. arXiv: 1104.3088 [astro-ph] (cit. on p. 153).

[203] E. Aprile et al. “Low-Mass Dark Matter Search Using Ionization Signals in XENON100”. *Physical Review D* 94.9 (Nov. 3, 2016). doi: 10.1103/PhysRevD.94.092001 (cit. on pp. 153, 166, 170, 171).

[204] J. Aalbers. *S2-Only Analysis for XENON1T SR0*. June 2017. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:aalbers:s2only_sr0 (visited on 02/12/2018) (cit. on pp. 155, 167).

[205] A. Tomás et al. “Study and Mitigation of Spurious Electron Emission from Cathodic Wires in Noble Liquid Time Projection Chambers” (Jan. 22, 2018). arXiv: 1801.07231 [physics] (cit. on p. 160).

[206] Rainer Storn and Kenneth Price. “Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”. *Journal of Global Optimization* 11 (Jan. 12, 1997), pp. 341–359. doi: 10.1023/A:1008202821328 (cit. on p. 162).

[207] J. Howlett, Z. Greene, and F. Gao. *Summary Note on Missing Top Population for SR1*. June 21, 2017. URL: https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:analysis:sciencerun1:headless_detector (visited on 02/13/2018) (cit. on p. 165).

[208] The DarkSide Collaboration et al. “Low-Mass Dark Matter Search with the DarkSide-50 Experiment” (Feb. 20, 2018). arXiv: 1802.06994 [astro-ph] (cit. on p. 166).

[209] SuperCDMS Collaboration et al. “WIMP-Search Results from the Second CDMSlite Run”. *Physical Review Letters* 116.7 (Feb. 17, 2016). doi: 10.1103/PhysRevLett.116.071301. arXiv: 1509.02448 (cit. on p. 166).

[210] T. Zhu. *Single Scatter (S2) Cut SR1 Summary*. Nov. 22, 2017. URL: <https://xe1t-wiki.lngs.infn.it/doku.php?id=xenon:xenon1t:sim:notes:tzhu:s2singlescattersr1summary> (visited on 02/13/2018) (cit. on p. 168).

Chapter Z

Donkere materie zoeken met XENON1T (samenvatting)

Natuurkunde streeft naar een fundamenteel en exact begrip van het universum. De afgelopen decennia zagen we spectaculaire vooruitgang naar dit doel, met ontdekkingen zoals zwaartekrachtgolven en het Higgs boson als hoogtepunt. Toch begrijpen we nog steeds niet waar het grootste deel van het universum van gemaakt is. Voor iedere kilo gewone materie bevat het heelal 5.4 kg *donkere materie* [1]: onbekende deeltjes die slechts zeer zwak reageren met atomen en licht. Ontdekken uit welke deeltjes donkere materie bestaat is één van de grootste uitdagingen in de hedendaagse deeltjesfysica.

Figuur Z.1 vat de eigenschappen samen van donkere en gewone materie in het vroege en huidige universum. Het jonge heelal was zo warm dat gewone materie op zichzelf niet zou kunnen samenklonteren: stralingsdruk bood voldoende tegenwicht aan de zwaartekracht. In het koelere heelal van vandaag kan materie *wel* samenklonteren tot sterren, door de energie die tijdens het instorten vrijkomt weg te stralen. Donkere materie reageert nauwelijks op straling, en kon dus in het vroege universum al samenklonteren. Het kon ook materie met zich meeslepen (door de zwaartekracht) en zo kosmische structuren zoals sterrenstelsels mogelijk maken. Donkere materie vormt echter geen sterren, omdat het de energie die bij het instorten vrijkomt niet kan wegstralen. Vandaag vinden we donkere materie daarom in dynamische (viriale) evenwichtsstructuren, zoals halo's rond sterrenstelsels.

Bewijs voor donkere materie komt van een groot aantal onafhankelijke waarnemingen. Halo's van donkere materie rond sterrenstelsels en clusters (rechtsonder in figuur Z.1) zijn zichtbaar door het effect van hun zwaartekracht op materie (door de beweging van sterren en sterrenstelsels) en licht (door lenswerking van de zwaartekracht). Het touwtrekken tussen zwaartekracht en stralingsdruk in het vroege universum (linksboven in figuur Z.1) heeft sporen achtergelaten in de kosmische achtergrondstraling (CMB) – straling die vrijkwam toen het heelal genoeg was afgekoeld om transparant te worden –

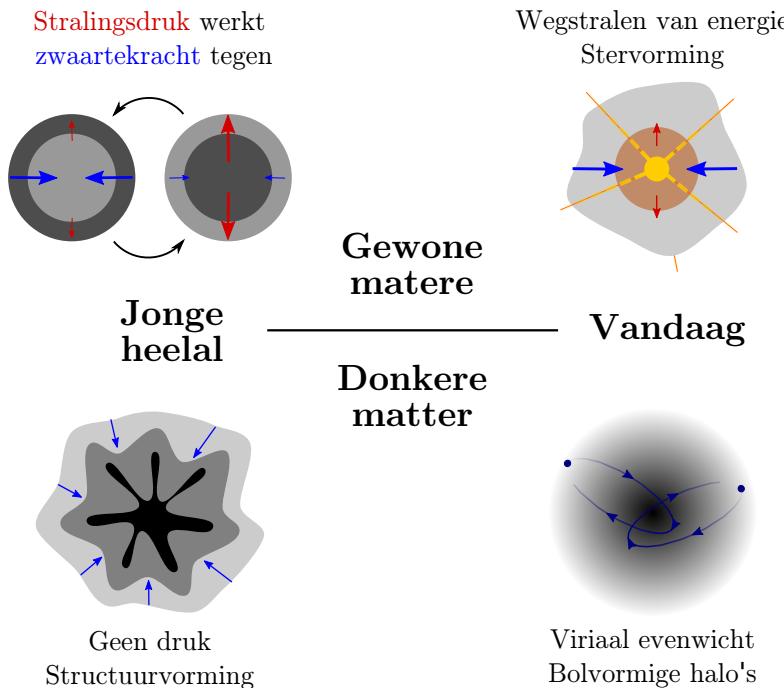


Figure Z.1: Versimpelde weergave van de geschiedenis van donkere en gewone materie in het heelal.

die onmiskenbare voetafdrukken van donkere materie bevatten. Sterker nog: de CMB laat zien dat het vroege universum zo uniform was dat sterrenstelsels niet gevormd hadden kunnen worden zonder de invloed van donkere materie (linksonder in figuur Z.1). Zonder donkere materie was de sterrenhemel duister, en was er niemand om haar te aanschouwen.

Dit astrofysisch en kosmologisch bewijs laat zien dat donkere materie en materie elkaar via de zwaartekracht beïnvloeden. Helaas is zwaartekracht te zwak om individuele deeltjes te meten (op zwarte gaten na). Experimenten proberen daarom donkeremateriedeeltjes te vinden op drie verschillende manieren:

- maak donkere materie bij botsingen van gewone materie (in deeltjesversnellers),
- zoek zichtbare producten van annihilatie van donkere materie en antideeltjes van donkere materie, of
- zoek naar de minuscule bewegingen van gewone atomen door zeldzame botsingen met donkere materie (verstrooing).

Elke methode heeft beperkingen: versnellers kunnen donkere materie niet maken als daarvoor eerst een zeer zwaar deeltje gemaakt moet worden; annihilatie veronderstelt het bestaan van donkere antideeltjes en moet onderscheiden

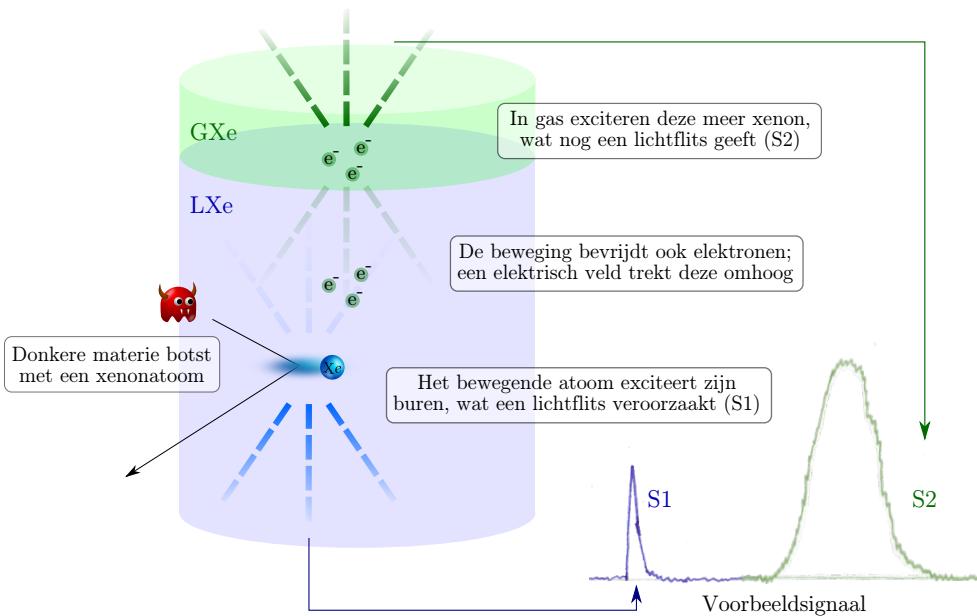


Figure Z.2: Principe van een twee-fase vloeibaar-xenon (LXe) gasvormig-xenon (GXe) tijdprojectiekamer (TPC) zoals XENON1T.

worden van een veelheid aan andere astrofysische signalen; verstrooiingsexperimenten tenslotte zijn blind voor donkere materie als dat uit te lichte deeltjes bestaat, en botsingen met atomen geen meetbare energie overdragen.

Dit proefschrift beschrijft een zoektocht naar donkere materie met XENON1T, 's werelds meest gevoelige donkeremateriedetector. Figuur Z.2 laat zien hoe XENON1T werkt. Eerst botst donkere materie met een xenonatoom in de detector (specifieker, de atoomkern). De beweging van dit atoom exciteert en ioniseert naburige atomen. Geëxciteerde atomen sturen (na het vormen van instabiele Xe_2 moleculen) fotonen uit; een lichtflits (S1) die wordt opgevangen door gevoelige sensoren genaamd fotomultiplicatoren (PMTs) boven- en onderin de detector. De elektronen die bij ionisatie vrijkomen worden door een elektrisch veld naar een gedeelte van de detector met gasvormig xenon getrokken. Hier worden ze door een nog sterker elektrisch veld versneld, waardoor meer xenonatomen geëxciteerd worden, wat vervolgens een tweede, grotere lichtflits (S2) veroorzaakt. De verdeling van licht over de PMTs en de tijd tussen S1 en S2 geven samen de 3D-positie van de oorspronkelijke botsing in de detector aan.

Helaas zijn er vele andere processen die dit soort signalen veroorzaken, en waar XENON1T zich dus tegen moet wapenen. Om te beschermen tegen kosmische straling staat XENON1T diep ondergronds, in het *Laboratori Nazionali del Gran Sasso* onder de Italiaanse Apennijnen. Natuurlijke radioactiviteit, de aartsvijand van gevoelige astrodeeltjesfysica-experimenten, wordt bestreden door speciaal gescreende materialen en een 10 m hoge waterbad wat de detector

omsluit. Bovendien wordt alleen de binnenste ton van de drie ton xenon in XENON1T gebruikt voor de donkere materie zoektocht; de andere twee ton geeft extra stralingsbescherming. De grootste overgebleven achtergrond komt van radioactieve verontreinigingen (in het bijzonder radon) die in het xenon oplossen uit materialen in de detector. Deze veroorzaken meestal botsingen met individuele elektronen in plaats van hele xenon atomen, en kunnen daarom goed (maar niet perfect) onderscheiden worden van botsingen met donkere materie door een andere verhouding tussen S1- en S2-licht.

Figuur Z.3 laat de signalen van XENON1T's eerste donkeremateriezoektocht van 34.2 dagen zien (in detail beschreven in hoofdstuk 5). Bijna alle signalen komen overeen met achtergrondmodellen die zijn afgeleid van ijkmetingen waarbij een radioactieve bron expres naast de detector was geplaatst. Zelfs het meest donkerematerieachtige signaal (het rode punt midden in figuur Z.3) komt overeen met achtergrond van elektron-botsingen. Één signaal met grote S1 maar kleine S2 heeft geen eenvoudige verklaring, maar ligt ver buiten de regio waar donkere materie verwacht wordt. Dit komt waarschijnlijk van een zeer zeldzame achtergrond, zoals het toevallig samenkomen van ongerelateerde S1- en S2-signalen.

Deze eerste resultaten bewijzen dat XENON1T's achtergrond lager is dan enig ander experiment ooit heeft bereikt, althans in de energieregio relevant voor donkeremateriebotsingen. Aangezien de resultaten geen aanwijzing voor donkere materie bevatten, kan XENON1T nieuwe grenzen stellen aan de sterkte van de interactie tussen donkere en gewone materie, zoals getoond in figuur Z.4. Grenzen van de LUX en PandaX-II experimenten zijn vergelijkbaar, maar dit zijn kleinere experimenten met een hogere achtergrond die hun maximale capaciteit hebben bereikt. XENON1T daarentegen zal binnenkort resultaten van een jaar durende zoektocht naar donkere materie bekend maken, wat tot veel sterkere grenzen zal leiden – of een ontdekking.

De komende jaren wordt XENON1T verbouwd tot XENONnT, een detector met driemaal zoveel xenon [2]. De daaropvolgende generatie donkerematerie-experimenten zal bestaan uit detectoren met meer dan tien ton xenon, zoals DARWIN [3]. Het komend decennium geeft ons ofwel hard bewijs tegen gangbare donkerematerietheorieën, ofwel een sensationele ontdekking. Donkere materie opgelet – XENON1T komt je halen!

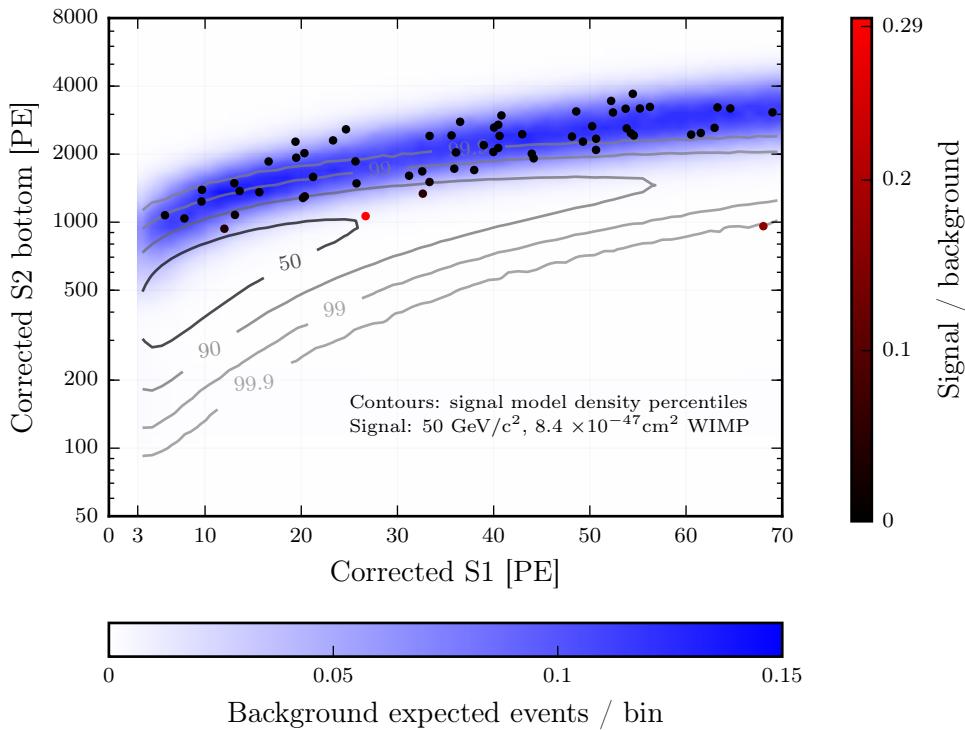


Figure Z.3: Punten: gevonden signalen in XENON1T's eerste donkere materie zoektocht, als functie van de S1- en S2-signalgroote gemeten in fotoelektronen (PE). Blauwe tinten laten zien waar we achtergrondsignalen verwachten; contouren waar donkeremateriesignalen zouden liggen (voor een bepaald donkeremateriemodel zoals aangegeven in de figuur). De contouren geven dichtheidpercentielen aan, zodat de helft van de donkeremateriesignalen binnen de contour '50' zouden liggen. De rode tinten op de punten geven de signaal/achtergrond waarschijnlijkhedsverhouding voor de gevonden signalen aan: een waarde van 1 zou betekenen dat een signaal evengoed achtergrond of donkere materie zou kunnen zijn (als het aangegeven donkeremateriemodel klopt).

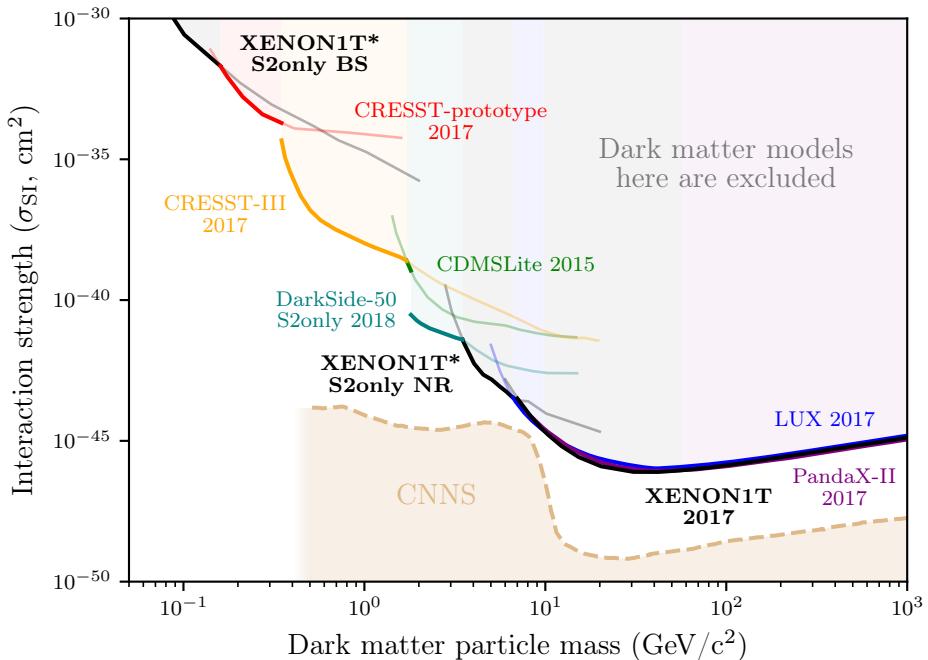


Figure Z.4: Grenzen op de interactiesterkte tussen materie en donkere materie, voor verschillende mogelijke massa's van donkere materie. Wereldrecords zijn aangegeven met dikke lijnen (met kleurtint erboven), andere resultaten als dunne lijnen. De zwarte grenzen (van XENON1T) worden in dit proefschrift afgeleid, inclusief de met een asterisk gemarkeerde *proof-of-concept* resultaten. Donkeremateriemodellen in de met CNNS gemarkeerde regio zullen lastig te vinden zijn door een neutrino-achtergrond van voornamelijk de zon [4] (voor de laagste massa's is dit niet in detail doorgerekend omdat het ver buiten bereik is van hedendaagse technologie).

Acknowledgements

Scientists say they stand on the shoulders of giants, but on close inspection, each giant is a human pyramid of like-minded colleagues supporting each other. I would not have completed this thesis without them.

First and foremost, I owe great thanks to my supervisor, Patrick Decowski, for giving me the opportunity to do this work, and for very valuable advice and detailed feedback during all stages of its growth. He and Auke-Pieter Colijn can be especially proud of creating and sustaining the motivating and supportive atmosphere in the Nikhef dark matter group.

I am grateful to my other promotor, Stan Bentvelsen, and *C3-member* Paul de Jong, for useful discussions and advice at several moments during my PhD. Thank you to Gianfranco Bertone, Jan Conrad, Auke-Pieter Colijn, Paul de Jong, Marieke Postma, Wouter Verkerke and Ivo van Vulpen for reading my thesis and forming my defense committee.

This research was supported by the University of Amsterdam's Gravitation and Astroparticle Physics (GRAPPA) Institute's PhD fellowship. I thank all my colleagues at GRAPPA for providing a stimulating environment.

When I was younger, I thought experimental physics was the domain of loners in dark cellars tangled in odd tubes and wires. Thanks to my colleagues at Nikhef, I discovered it is actually a very cooperative enterprise, where you continuously bounce ideas, plots and mad hypotheses off each other.

Christopher Tunnell switched between roles as software guru, collaborator, and mentor throughout my PhD. We worked on many projects together, such as the XENON1T eventbuilder (chapter 4) and first results paper (chapter 5). I especially enjoyed our trips to New York and Heidelberg to convert our collaboration to the Church of Python.

Andrew Brown provided his substantial expertise on many topics, especially the coverage problem in the XENON1T SR0 result (see section 6.3.2). I greatly enjoyed our conversations on life and science, and the various board games that accompanied them.

Andrea Tiseni showed us how to cook good pasta and win at poker (though it remains doubtful how much either skill has transferred). His and Chris' work on the XENON100 S2-only analysis inspired the S2-only analysis of section 7.4.

Sander Breur and Erik Hogenbirk have been great colleagues for the last

four years. Whether it is doing something in the lab, coding software, styling plots, trying to understand statistics, surviving in Italy, or identifying peaks in spectra, discussions with Sander and Erik were always the best place to start. Listing all our collaborations would go to far here, but I was especially glad to work with Erik on the S1 pulse shape analysis and Sander on the first commissioning data of XENON1T.

I am also grateful to all other current and former students and colleagues in the Nikhef dark matter group, especially Matteo Alfonsi, Rolf Schön, Alper Topçuoğlu, Bart Pelssers, Maria Bader, Froukje Gjaltema, Kiefer van Teutem, Luuk den Uijl, Dorine Schenk, Arjen Wildeboer, Joran Angevaare, Jasper Nobelen, Mitch Vonk, Katherine McEwan, and Peter Gaemers.

Sander and Bart deserve special thanks for their role as house-mates during the Chicago analysis workshop. If this was *strafkamp*, I vote for more totalitarianism.

I also had the privilege of working with many great XENON colleagues outside Nikhef. Special thanks go to Stefan Schindler, Jörn Mahlstedt, and Francesco Piastra for good times on shift in Italy; Daniel Coderre for collaboration on the XENON1T DAQ; and Knut Morå for collaboration on the statistical inference software.

I would never have started my PhD without crucial academic guidance and encouragement from Floris van der Burg and Frank Witte during my undergraduate years at University College Utrecht, and Gerard 't Hooft during my MSc thesis in Utrecht. I am indebted to my colleagues from the *Eerst the Klas* traineeship and my colleagues and students at the *Gymnasium Apeldoorn* for the opportunity and privilege to become a physics teacher. And going somewhat further back, to the volunteers at the *Volkssterrenwacht Bussloo*, for encouraging an inquisitive young kid to take an interest in astrophysics.

I would like to thank my good friends Inge, Eelke and Brigitte for continuing to put up with me even when I was deeply absorbed in scientific arcana. Finally, my deepest thanks go to my parents, for their support and encouragement in countless ways for close to thirty years.

