PLATON: Pressure-Testing High-Resolution Three-Dimensional Particle Tracking in a Monolithic Scintillator Block via Time-of-Flight Photon Trilateration

K-Dense
Independent Research Report

July 18, 2026

Abstract

Finely segmented detectors—liquid-argon time-projection chambers, plastic scintillator cube trackers, and crystal-ring positron-emission scanners—purchase three-dimensional imaging power with channel counts that grow with detector *volume*. The PLATON concept proposes an alternative: a single, unsegmented block of scintillator instrumented only on its surface with light-field cameras and single-photon timing sensors, with three-dimensional reconstruction recovered computationally from the space–time pattern of arriving optical photons rather than from physical segmentation. We pressure-test the central claim—that such a monolith can localize interaction vertices and coarse tracks about as well as segmented devices—using a transparent, first-principles Monte Carlo. We simulate isotropic optical-photon emission from point-like vertices and short straight tracks inside a 20 cm homogeneous cube (refractive index $n = 1.58$, in-medium speed $v = c/n = 18.97\,\mathrm{cm\,ns^{-1}}$), instrumented with 600 surface pixels, and we propagate each photon to its exit face with Gaussian sensor timing jitter σ_t and a binomial photon-detection efficiency. We then invert the arrival pattern with two estimators: a robust four-parameter time-of-flight trilateration solver for the vertex (Huber loss, analytic gradient, quasi-Newton optimization) and a seven-parameter marginalized-maximum-likelihood solver for the track that integrates over each photon's unknown emission point. Sweeping detected light yield $N \in \{500, 1000, 2000\}$ and jitter $\sigma_t \in \{100, 200\}$ ps, we find that at a realistic operating point ($\sigma_t = 100$ ps, $N = 1000$) the reconstruction reaches a median vertex resolution of 1.31 mm, a median track-direction error of 7.0°, and a median track-length error of 2.37 mm, using roughly two orders of magnitude fewer channels than a comparable segmented tracker. Resolution improves with photon statistics and degrades sharply as jitter approaches the particle traversal time, exactly as a timing-trilateration estimator predicts. Because the single largest threat to any real single-photon device is sensor noise, we extend the study with a validated physical model of the three dominant silicon-photomultiplier noise sources—dark counts, optical crosstalk, and afterpulsing—and a spatial–temporal coincidence filter that exploits the correlation of genuine scintillation light in both space and time. Unfiltered, dark counts are benign for the outlier-robust vertex fit up to 1 MHz per channel but catastrophic for tracking, driving the median direction error from 5° to 60° and collapsing head–tail identification toward chance; the filter, which retains $\sim$99% of true photons while rejecting $\gtrsim$96% of dark counts at 1 MHz, restores the vertex resolution to 1.47 mm and the direction error to below 20°, and holds the vertex below 2.6 mm even at a punishing 10 MHz. A structured comparison against LArTPC (DUNE), the T2K SuperFGD/3DST plastic tracker, and TOF-PET shows that PLATON's decisive advantage—surface (L^2) rather than volume (L^3) cost scaling—is largest for the biggest detectors, making large-volume neutrino experiments, not PET, the most plausible first beneficiary. We are explicit about what this back-of-the-envelope study can and cannot settle: even with sensor noise now modeled and mitigated, it still omits optical scattering, surface reflections, and finite scintillator kinetics, all of which corrupt the clean straight-line timing that the reconstruction assumes. The results establish that the core physics is favorable and robust to realistic sensor noise, and motivate a full optical-transport simulation as the necessary next step.

Keywords: monolithic scintillator; time-of-flight; single-photon detection; 3D reconstruction; neutrino detectors; positron emission tomography; Monte Carlo simulation; silicon photomultiplier noise; dark count rate; optical crosstalk; afterpulsing; coincidence filtering

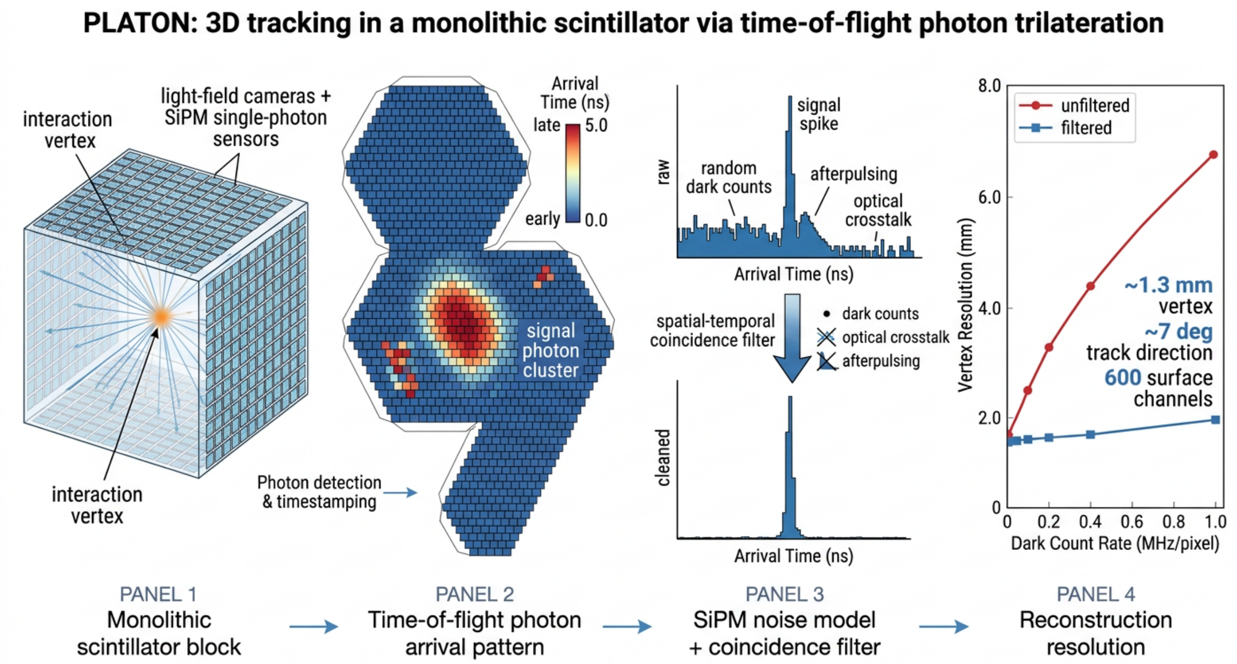

Figure 1: Graphical abstract. The PLATON concept and this study's workflow, left to right. *Panel 1:* a charged-particle or neutrino interaction inside a monolithic scintillator cube emits optical photons isotropically; the six faces carry light-field cameras and single-photon SiPM sensors—the entire instrumentation. *Panel 2:* photons arrive across the surface with a characteristic space–time pattern (color encodes arrival time), forming a compact bright cluster near the interaction; time-of-flight trilateration converges on the vertex while a marginalized likelihood recovers the three-dimensional track. *Panel 3:* a validated model of the three dominant SiPM noise sources—dark counts, optical crosstalk, and afterpulsing—contaminates the raw arrival-time spectrum, and a spatial–temporal coincidence filter exploits the space–time correlation of genuine light to recover a clean signal burst. *Panel 4:* the filtered vertex resolution stays near 1.5 mm across four decades of dark count rate while the unfiltered resolution diverges, reaching ∼1.3 mm vertex and ∼7° track direction with only 600 surface channels at the nominal operating point.

1 Introduction

Scintillation detectors have been a workhorse of particle and nuclear physics for three-quarters of a century, converting the energy deposited by ionizing radiation into prompt bursts of optical photons that a photosensor turns into an electrical signal. The modern art of these detectors is not the light production, which is well understood, but the reconstruction of *where* and *when* the interaction happened from the light that emerges. The dominant strategy for extracting spatial information has been physical segmentation: divide the active medium into many small, optically isolated cells, and infer the interaction position from which cells light up. Liquid-argon time-projection chambers (LArTPCs) drift ionization electrons onto planes of wires or pixels to image neutrino interactions in three dimensions with millimeter fidelity [1, 2]; segmented plastic scintillator trackers such as the T2K SuperFGD read out a three-dimensional lattice of centimeter cubes through orthogonal wavelength-shifting fibers [3]; and clinical positron-emission tomography (PET) scanners tile a ring with thousands of fast crystal elements [4, 5]. In each case the granularity of the reconstruction is set, more or less directly, by the granularity of the hardware. Even where sophisticated software imaging is layered on top of the hardware—for example the Wire-Cell tomographic reconstruction and sparse convolutional-network event selection developed for MicroBooNE LArTPC data [6, 7]—the achievable spatial fidelity remains anchored to the underlying channel pitch.

Segmentation is powerful but expensive, and its cost scales unfavorably. The number of readout channels—and with it the wiring, front-end electronics, calibration burden, and dead material—grows with the detector's active *volume*. A LArTPC module for the Deep Underground Neutrino Experiment (DUNE) carries of order 10^5–10^6 sense channels [8], and the SuperFGD's roughly two million plastic cubes each require three fibers and associated silicon photomultipliers (SiPMs). For a detector whose linear size is L, this volumetric (L^3) scaling of channel count and complexity is the single hardest obstacle to building ever-larger devices, and it motivates a recurring question: how much of the spatial information carried by scintillation light can be recovered *computationally*, from sensors placed only on the surface, without segmenting the bulk at all?

A monolithic—that is, unsegmented—scintillator block turns this question into a concrete detector concept. Because the medium is transparent, light generated anywhere inside the block propagates to the boundary, where a shell of sensors records both the spatial distribution and the arrival time of the photons. The interaction position is then encoded not in which cell fired but in the geometry of the light front: the closer a sensor is to the interaction, the earlier and more intensely it is illuminated. This idea has a mature and quantitative pedigree in medical imaging, where monolithic scintillator PET detectors reconstruct the gamma-interaction position, including depth of interaction, from the light distribution on a pixelated photosensor. Analytic light-distribution methods and, more recently, machine-learning regressors—gradient tree boosting [9], convolutional and inception networks [10], and deep networks that jointly estimate position and time stamp [11, 12]—routinely achieve sub-millimeter planar resolution and few-millimeter depth resolution in monolithic crystals a few centimeters across. The physics of light transport in these devices is well validated against Monte Carlo [13]. Monolithic PET, however, exploits the *spatial* light distribution over short optical path lengths; it does not primarily rely on photon *timing* to localize interactions across a large volume.

The PLATON concept extends the monolithic idea in two directions at once. First, it scales the block up to the size relevant for particle tracking rather than 511-keV gamma localization, so that the interior optical paths become tens of centimeters and the time of flight of light becomes a first-class observable. Second, it pairs conventional intensity imaging with two complementary surface technologies: light-field (plenoptic) cameras, which capture not only where light lands but the direction from which it arrives [14, 15], and single-photon avalanche sensors with picosecond-scale timing, whose single-photon time resolution now reaches the sub-hundred-picosecond regime in the best silicon photomultipliers [16, 17]. The premise is that, with a fast enough clock on a dense enough surface, the three-dimensional interaction vertex and even the direction of a short track can be trilaterated from photon arrival times—the optical analogue of how a global-positioning receiver locates itself from the timing of radio signals—so that a single solid block reconstructs tracks in three dimensions "about as well as" today's finely segmented detectors, at a fraction of the channel count.

That is a strong claim, and the purpose of this report is to pressure-test it independently rather than to advocate for it. We take the concept at face value and ask a narrower, answerable question: *in an idealized but physically faithful simulation, how well can an interaction vertex and a short track be recovered from the surface photon arrival pattern, and how does that resolution scale as the two quantities most under threat in a real device—light yield and sensor timing precision—get worse?* We deliberately adopt a first-principles "toy" Monte Carlo: photons are emitted, propagated in straight lines at the in-medium speed of light, and detected with realistic jitter and efficiency, but the more pernicious effects of real optical media—bulk scattering, boundary reflections, finite scintillator rise and decay, and sensor noise—are held out. This choice is a feature, not an oversight: it isolates the best-case timing physics so that the resulting resolution is an *upper bound* on what the concept can deliver, and it makes the reconstruction transparent enough to reason about analytically. The honest corollary, which we develop at length in the discussion, is that a favorable result here is necessary but not sufficient for the design to work.

Our contributions are fourfold. We build and validate a vectorized Monte Carlo of optical-photon propagation and detection in a monolithic cube (Section 2), passing twelve independent physics sanity checks. We develop two reconstruction estimators—a robust time-of-flight vertex solver and a marginalized-maximum-likelihood track solver—and characterize their spatial, angular, and head–tail performance across a light-yield and timing-jitter sweep (Section 3). Because sensor noise is the effect most likely to undermine a real single-photon device, we then add a validated physical model of the three dominant silicon-photomultiplier noise sources—dark counts, optical crosstalk, and afterpulsing—together with a spatial–temporal coincidence filter that suppresses them, and we quantify how reconstruction resolution degrades and recovers across a dark-count-rate sweep (Sections 2 and 3). Finally, we place the results in context with a structured comparison against LArTPC, segmented plastic, and TOF-PET detectors (Section 4) and a strategic technical brief on where the design is most likely to pay off first, with explicit caveats about the limits of a back-of-the-envelope study (Section 5).

2 Methods

2.1 Detector geometry and optical model

The simulated detector is a homogeneous cube of scintillator spanning $[-10, +10]$ cm on each axis, a 20 cm block chosen to be large enough that internal optical paths (10–17 cm) make time of flight a meaningful observable, yet small enough to keep the idealized straight-line propagation assumption defensible over the relevant distances. The medium has a refractive index $n = 1.58$, representative of common plastic and many inorganic scintillators, giving a group speed of light in the medium of

$$v = \frac{c}{n} = \frac{29.979\,\mathrm{cm\,ns^{-1}}}{1.58} = 18.974\,\mathrm{cm\,ns^{-1}}. \tag{1}$$

Each of the six faces is tiled with a 10×10 grid of square pixels at a 2 cm pitch, for a total of 600 surface channels. Each pixel stores its three-dimensional center coordinate, its outward normal, and its face and grid indices, and a fast $(\mathrm{face}, i, j) \to$ pixel-id lookup table supports vectorized hit assignment. This 600-channel shell is the entire instrumentation: there is no readout in the interior. The pixel count deliberately abstracts the light-field-camera-plus-SiPM readout of the PLATON concept into an effective spatio-temporal sampling of the surface; each "pixel" is a location at which the photon flux and arrival time are recorded with finite timing precision.

An interaction is modeled as a source of scintillation photons emitted isotropically—uniformly over the unit sphere by area, sampled so that the mean direction vanishes and the per-axis variance equals 1/3. Two source types are simulated. A *point source* represents a localized energy deposit (a neutrino interaction vertex or an annihilation point) at position $\mathbf{x}_0$ and time t_0, emitting N photons. A *track source* represents a minimum-ionizing particle traversing a straight segment from $\mathbf{a}$ to $\mathbf{b}$: the N photons are distributed uniformly in arc length s along the segment, and a photon emitted at arc length s is emitted at time $t_0 + s/v_p$, where v_p is the particle speed (taken as the vacuum speed of light c for a relativistic particle). This finite emission time along the track is precisely the signal that any head–tail (direction-of-motion) determination must exploit.

Photon propagation is treated in the geometric, non-scattering limit. Each emitted photon travels in a straight line from its emission point along its sampled direction until it first exits the cube; a vectorized ray–box intersection returns the exit face and the exit point, which is mapped to the nearest pixel on that face. The propagation time to the assigned pixel is $t_{\text{prop}} = d/v$, where d is the emission-to-pixel-center distance. The recorded arrival time is

$$t_{\text{hit}} = t_0 + \frac{s}{v_p} + \frac{d}{v} + \mathcal{N}(0, \sigma_t^2), \tag{2}$$

where the Gaussian term models sensor and single-photon timing jitter with standard deviation σ_t, and $s = 0$ for a point source. Photon detection efficiency is applied as an independent binomial down-sampling with probability equal to the photodetection efficiency (PDE); in the sweep the nominal detected yield N therefore already accounts for the sensor efficiency. The model records, per pixel, the detected photon count and the first and mean arrival times. Equation (2) makes explicit which effects of the *signal* photons are included—geometric time of flight, per-photon Gaussian jitter, finite emission time along a track, and stochastic detection efficiency—and, by omission, which are held out: bulk scattering and absorption, Fresnel reflection and total internal reflection at the faces, and the finite rise and decay time of the scintillator emission. The sensor-noise processes that afflict a real single-photon readout—dark counts, optical crosstalk, and afterpulsing—are not omitted but are added as a distinct stochastic layer on top of the signal hits, described in Section 2.3; the optical-transport effects that remain held out are addressed in Section 5.

2.2 Simulation validation

Before using the forward model to test reconstruction, we subjected it to twelve independent physics sanity checks, all of which passed. The pixelation produces exactly 600 channels, all lying on the cube boundary, balanced at 100 per face. The in-medium speed matches c/n to machine precision, and the propagation-time bookkeeping satisfies $t_{\text{prop}} \equiv d/v$ with zero residual. Isotropy is confirmed by a vanishing mean emission direction and per-axis direction variance of $(0.332, 0.334, 0.333)$ against the expected 1/3. A source at the cube center illuminates the six faces symmetrically (face fractions 0.165–0.168, expected 0.1667) and produces minimum and maximum photon-to-pixel distances of 10.10 cm and 16.19 cm, matching the geometric face-center and corner-pixel distances. Injecting a known detection efficiency of 0.30 recovers 0.299, and injecting a known jitter of 0.25 ns recovers $\sigma_t = 0.249$ ns. Figure 2 shows the resulting three-dimensional hit patterns for a central point source, which illuminates all 600 pixels symmetrically, and for a diagonal track, whose hit pattern is visibly asymmetric toward its start corner, with pixels colored by mean arrival time.

2.3 Sensor noise model: dark counts, optical crosstalk, and afterpulsing

A single-photon readout never sees only signal. Silicon photomultipliers and single-photon avalanche diodes are afflicted by three characteristic noise processes that a credible pressure-test cannot ignore, and each has been measured and modeled extensively in the detector-physics literature [17–19]. We

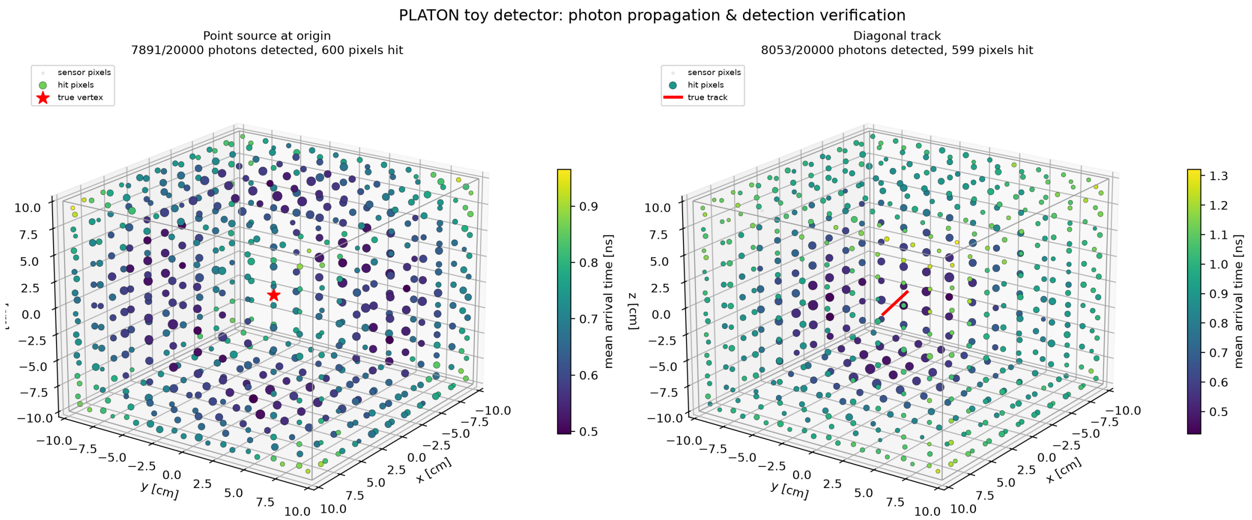

Figure 2: Forward-model verification. Three-dimensional visualization of the detector cube, the 600-pixel surface grid (open squares), and the detected hits (filled markers, colored by mean arrival time). *Left:* a point source at the origin emits 2×10^4 photons; 7891 are detected across all 600 pixels with the expected six-fold face symmetry. *Right:* a diagonal track (red line) emits 2×10^4 photons; 8053 are detected, and the hit pattern is asymmetric toward the track's start corner. Color encodes mean arrival time (ns); pixels nearer the source register earlier, the geometric signature that time-of-flight reconstruction exploits.

add all three to the forward model as a distinct stochastic layer applied *after* the signal photons of Equation (2) have been generated, in the physical order in which the effects occur, so that each true or spurious hit can itself seed a correlated secondary.

The first process is the *dark count rate* (DCR): thermally generated avalanches that fire pixels at random, uncorrelated with any real light, which degrade timing by littering the readout window with hits that any first-arrival or trilateration estimator can mistake for signal [20, 21]. Over a readout window of duration Δt across N_{pix} pixels at a per-pixel rate R_{DCR}, the number of dark hits is Poisson-distributed with mean $\mu = R_{\text{DCR}} N_{\text{pix}} \Delta t$; each dark hit is placed on a uniformly random pixel at a uniformly random time within the window. The second is *optical crosstalk* (OCT): during an avalanche a pixel emits secondary photons that can trigger an adjacent pixel, so that with probability p_{OCT} each primary hit (a true photon *or* a dark count) fires a prompt companion in a randomly chosen edge-adjacent pixel on the same face, at essentially the same time [22, 23]. The third is *afterpulsing* (AP): charge carriers trapped during an avalanche are released after a delay and re-trigger the same pixel, so that with probability p_{AP} each hit (primary or crosstalk secondary) produces a delayed hit in the *same* pixel, with the delay drawn from an exponential distribution of time constant τ_{AP} [18, 23]. Crosstalk and afterpulsing are thus *correlated* noise—spatially and temporally tied to a parent avalanche—whereas dark counts are wholly uncorrelated, a distinction the mitigation filter below exploits directly. The nominal noise parameters used throughout are $p_{\text{OCT}} = 0.15$, $p_{\text{AP}} = 0.05$, and $\tau_{\text{AP}} = 15\,\text{ns}$, values representative of modern SiPMs at moderate overvoltage [19], with R_{DCR} swept over four decades as described below.

We validated the noise generator against its analytic expectations before using it. The dark-count multiplicity over 3000 trials at $R_{\text{DCR}} = 1\,\text{MHz}$ per channel matches the target Poisson mean of 36 hits per window (sample mean 36.09, Fano factor 1.001, Poisson goodness-of-fit $p = 0.14$), with dark hits uniform across pixels ($p = 0.57$) and in time ($p = 0.74$). Every crosstalk hit is correctly localized to a pixel edge-adjacent to its parent on the same face, and the measured crosstalk probability of 0.148 recovers the injected 0.15. The afterpulse delays follow the intended exponential law: a maximum-likelihood fit returns $\tau_{\mathrm{AP}} = 15.05\,\mathrm{ns}$ against the injected 15 ns, a Kolmogorov–Smirnov test is comfortably consistent with an exponential ($p = 0.90$), and the measured afterpulse probability of 0.0499 recovers the injected 0.05. Figure 3 shows the arrival-time spectra that these processes produce: a clean signal burst in the noise-free case, a flat pedestal of dark counts spread across the whole window when DCR is switched on, and the full mixture—signal burst, uniform dark pedestal, prompt crosstalk co-located with the burst, and an exponentially decaying afterpulse tail—in the realistic case.

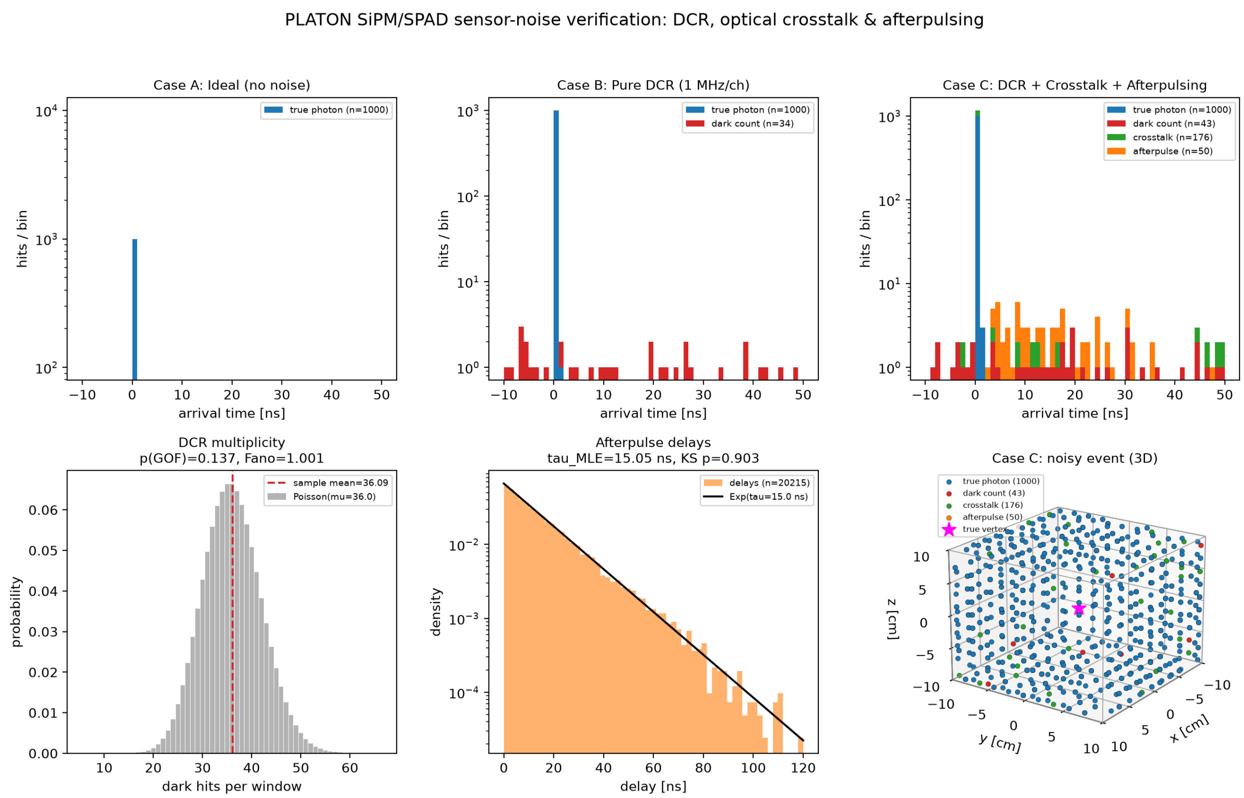

Figure 3: Sensor-noise model and its validation. *Top row:* arrival-time spectra for a central point source ($N = 1000$ true photons, $\sigma_t = 100\,\mathrm{ps}$). Case A (ideal) shows only the sharp signal burst; Case B adds dark counts at 1 MHz per channel, producing a flat pedestal across the whole $[-10, +50]$ ns window; Case C adds prompt optical crosstalk (co-located with the signal burst) and exponentially delayed afterpulses (the long tail). *Bottom row:* validation of each process—the dark-count multiplicity is Poisson with the expected mean $\mu = 36$ (Fano $= 1.00$, GOF $p = 0.14$); the afterpulse delays follow an exponential with maximum-likelihood $\tau = 15.05\,\mathrm{ns}$ against the injected 15 ns (KS $p = 0.90$); and a three-dimensional view of a noisy event shows true hits (blue) overlaid with dark counts, crosstalk, and afterpulse hits scattered across the surface.

2.4 Vertex reconstruction: time-of-flight trilateration

For a point-like interaction, the inverse problem is to recover the four unknowns $(\mathbf{x}_{\mathrm{vtx}}, t_0)$ from the per-pixel arrival times. Each pixel j that records a hit at time t_j provides a residual

$$r_j = t_j - t_0 - \frac{\|\mathbf{x}_j - \mathbf{x}_{\mathrm{vtx}}\|}{v}, \tag{3}$$

which is the difference between the observed arrival time and the time predicted by straight-line time of flight from a candidate vertex. This is mathematically the same trilateration problem solved by satellite navigation and by TOF-PET when localizing an annihilation point along a line of response [24], here generalized to a full three-dimensional position from a dense surface of receivers. Because a minority of pixels can carry anomalously early or late times—from the tails of the jitter distribution, and, in a real device, from scattering or noise—we minimize a robust Huber loss rather than a plain least-squares sum:

$$\mathcal{L}(\mathbf{x}_{\mathrm{vtx}}, t_0) = \sum_j \rho_\delta(r_j), \qquad \rho_\delta(r) = \begin{cases} \frac{1}{2} r^2 & |r| \le \delta, \\ \delta \left(|r| - \frac{1}{2}\delta \right) & |r| > \delta, \end{cases} \tag{4}$$

which behaves quadratically for small residuals and linearly for large ones, bounding the influence of outliers. The loss is minimized with a quasi-Newton method (L-BFGS-B), supplying the analytic gradient of $\mathcal{L}$ with respect to all four parameters so that the optimizer converges quickly and reliably; the classical Levenberg–Marquardt trust-region idea [25] underlies this style of nonlinear least-squares fitting. The vertex spatial error is reported as the Euclidean distance $\|\hat{\mathbf{x}}_{\mathrm{vtx}} - \mathbf{x}_{\mathrm{vtx}}\|$ between the reconstructed and true vertices.

2.5 Track reconstruction: marginalized maximum likelihood

A short track is harder than a point because each detected photon was emitted from an *unknown* location along the segment, and that emission time—not just the geometric delay—contributes to the arrival time in Equation (2). Naively fitting a single "effective" emission point biases the result and discards the very information (the ordering of emissions in time) that distinguishes the head of the track from its tail. We therefore treat the emission arc length as a latent variable and marginalize over it. For a candidate track with endpoints $(\mathbf{a}, \mathbf{b})$ and start time t_0, the predicted arrival-time distribution at pixel j is obtained by integrating over the fractional position $u \in [0, 1]$ along the segment, $\mathbf{x}(u) = \mathbf{a} + u(\mathbf{b} - \mathbf{a})$, with emission time $t_0 + u\,\|\mathbf{b} - \mathbf{a}\|/v_p$ and propagation time $\|\mathbf{x}_j - \mathbf{x}(u)\|/v$. The negative log-likelihood of the observed times, marginalized over u on a fixed fractional grid and convolved with the Gaussian jitter kernel, is minimized over the seven parameters $(\mathbf{a}, \mathbf{b}, t_0)$. The analytic gradient of this marginalized likelihood—whose per-photon weights take the form of softmax "responsibilities" over the grid of candidate emission points—agrees with finite differences to $\sim 10^{-11}$, confirming the implementation.

Because the likelihood surface for a nearly spherical short-track hit pattern is multimodal, the optimizer is seeded robustly: a coarse multi-start scan over candidate orientations is combined with a principal-component-analysis seed of the spatial hit distribution, tried in *both* orientations (head-first and tail-first), and each seed is refined with gradient-based L-BFGS-B. The two-orientation seeding is what allows the fit to resolve the head–tail ambiguity: the orientation whose marginalized likelihood is higher is selected, and a correct choice is scored when the reconstructed direction aligns with the true direction of motion (positive dot product). This RANSAC-like philosophy of proposing candidate models and selecting the one best supported by the data [26] makes the fit resilient to the near-degeneracy of the short-track geometry. We report four track metrics: the angular error between reconstructed and true direction; the track-length error; the start- and end-point errors; and the head–tail accuracy, the fraction of tracks whose direction of motion is correctly identified.

2.6 Noise mitigation: a spatial–temporal coincidence filter

The reconstruction estimators above assume that every recorded hit is a signal photon; feeding them the raw noisy stream corrupts the fit, because dark counts in particular contribute residuals that bear no relation to the true geometry. Rather than complicate the likelihood with a noise term, we clean the hit list first with a fast pre-processor that exploits the one property signal photons have and noise does not: *correlation.* Genuine scintillation light arrives in a narrow time burst—its spread is set by the internal light-crossing time of the block, only a nanosecond or two—and illuminates spatially contiguous groups of pixels, whereas dark counts are random in both time (uniform over the whole readout window) and space (any pixel), and afterpulses are delayed far beyond the signal burst. The filter is therefore built in two stages that mirror this structure.

The first stage is a *time-window* cut. Because the signal is the densest concentration of hits in time, we slide a window of width 3 ns over the sorted arrival times and identify its densest position; the leading edge of that window estimates the signal start time t_{start}, and only hits in the interval $[t_{\text{start}} - 2\,\text{ns},\, t_{\text{start}} + 8\,\text{ns}]$ are retained. This peak-density estimator locks onto the burst even when uncorrelated dark counts vastly outnumber signal, and the asymmetric, forward-leaning window removes the bulk of the uniform dark pedestal while discarding the exponentially delayed afterpulses, whose mean delay $\tau_{\text{AP}} = 15\,\text{ns}$ places most of them beyond the window. The second stage is a *neighbor-coincidence* cut applied to the time-window survivors: a hit is kept only if at least one of its edge-adjacent pixels also recorded a hit within $\pm 1.5\,\text{ns}$. Real scintillation pixels sit inside illuminated clusters and pass this test easily, whereas an isolated dark count that happens to survive the time cut almost never has a coincident neighbor and is rejected. The coincidence test is implemented efficiently by grouping surviving hits by pixel and performing a binary search over each neighbor's sorted hit times, so the whole filter adds negligible cost to the reconstruction. Prompt optical crosstalk, being both in-time and spatially adjacent to genuine signal, is deliberately *not* targeted by the filter: it is physically signal-like and, sitting on the true illuminated cluster, biases the reconstruction far less than a randomly placed dark count.

We characterize the filter with three figures of merit: the true-positive rate (TPR, the fraction of genuine signal photons retained), the dark-count rejection (the fraction of dark hits removed), and the downstream reconstruction resolution computed on the cleaned hit list. The design target is to retain more than 95% of signal while rejecting more than 90% of dark counts.

2.7 Parameter sweep and statistics

To measure how resolution degrades as the device becomes less capable, we sweep the two parameters most likely to be worse in a real detector than in this idealization: the detected light yield $N \in \{500, 1000, 2000\}$ photons, and the per-photon timing jitter $\sigma_t \in \{100, 200\}$ ps. The lower jitter value corresponds to the single-photon time resolution achievable with state-of-the-art SiPMs [16, 27]; the higher value represents a more conservative or degraded device. For each of the six (σ_t, N) cells we generate multiple randomized events—8 vertex trials and 6 track trials per cell (48 vertex and 36 track trials in total), with vertices and track endpoints drawn at random positions and orientations within a fiducial region—and summarize each metric by its median, mean, and 90th percentile over the cell. A fixed random seed (42) makes the entire study reproducible. All simulation and reconstruction code is implemented in Python with NumPy and SciPy; the reconstruction sweep runs in roughly ten minutes on a single core.

To measure the impact of sensor noise separately, we run a second sweep at the nominal operating point ($\sigma_t = 100\,\text{ps}$, $N = 1000$) in which the per-channel dark count rate is stepped over four decades, $R_{\text{DCR}} \in \{0,\, 10^5,\, 10^6,\, 10^7\}$ Hz, with correlated noise held fixed at $p_{\text{OCT}} = 0.15$ and $p_{\text{AP}} = 0.05$. At each level we again generate 8 vertex and 6 track trials and reconstruct each event *twice*—once from the raw noisy hit list and once from the list cleaned by the coincidence filter of Section 2.6—so that the filter's benefit is measured as a paired difference on identical events. The upper end of the range, 10 MHz per channel, corresponds to a total instantaneous dark rate of 6 GHz across the

600-pixel surface and is deliberately punishing, well above the DCR of a well-cooled device, so as to locate the point at which the reconstruction breaks.

3 Results

3.1 A representative event

Before the statistical sweep, a single high-statistics event illustrates the reconstruction working as intended. For a point source at $(3.0, -2.0, 4.0)$ cm emitting enough light to yield 1950 detected photons at $\sigma_t = 100$ ps, the trilateration solver returns $(2.93, -2.15, 4.06)$ cm, a vertex error of 1.80 mm, with a root-mean-square timing residual of 98 ps—consistent with the injected jitter, indicating that the fit has extracted essentially all available timing information. For a diagonal track running from $(-4.0, -3.0, -2.0)$ cm to $(0.81, 1.01, 2.97)$ cm with 2004 detected photons, the marginalized-likelihood solver recovers the direction to within 1.93°, places the endpoints to 2.1 and 3.9 mm, reproduces the track length to 2.1 mm, and identifies the head–tail orientation correctly. These single-event figures are better than the cell medians reported below because they correspond to a high-yield, low-jitter realization; they establish that the estimators are unbiased and well-calibrated when the data are favorable.

3.2 Resolution scaling with light yield and timing jitter

The central result of this study is how the reconstruction resolution scales across the parameter sweep, shown in Figure 4 and tabulated in Table 1. Three trends stand out, and each has a clean physical interpretation.

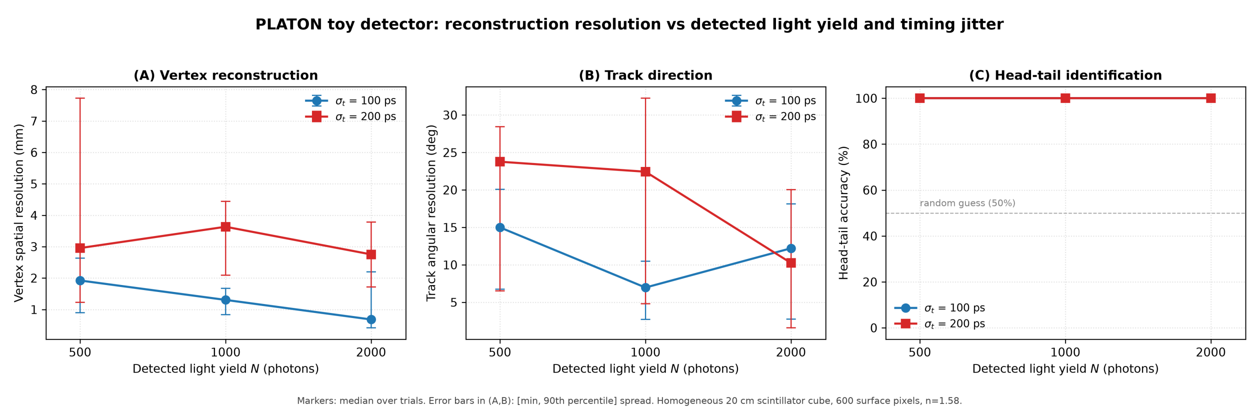

Figure 4: Reconstruction resolution versus detected light yield and timing jitter. Median performance over trials with [min, 90th percentile] spread bars, for $\sigma_t = 100$ ps (blue) and $\sigma_t = 200$ ps (red). **(A)** Vertex spatial resolution (mm) improves with detected light yield N and is systematically better at lower jitter. **(B)** Track angular resolution (°) shows the same jitter ordering, with the 100 ps curves well below the 200 ps curves. **(C)** Head–tail identification accuracy is 100% across all cells in this idealized simulation; the dashed line marks the 50% random-guess baseline. The homogeneous 20 cm cube has 600 surface pixels and $n = 1.58$.

First, *timing jitter drives everything.* Every metric is uniformly worse at $\sigma_t = 200$ ps than at 100 ps. At the nominal $N = 1000$, the median vertex error rises from 1.31 mm at 100 ps to 3.63 mm at 200 ps, and the median angular error from 7.0° to 22.4°. This is the defining signature of a timing-trilateration estimator: the spatial resolution is set by how precisely each pixel measures the arrival time of the light front, and a factor-of-two degradation in σ_t translates directly into a comparable degradation in the geometric localization. The practical message for a real device is stark—the entire concept lives or dies on fast timing.

Second, *resolution improves with photon statistics.* Along the $\sigma_t = 100$ ps row, the median vertex error falls from 1.92 mm at $N = 500$ to 1.31 mm at $N = 1000$ and 0.68 mm at $N = 2000$. Because the vertex is estimated by combining many independent noisy arrival-time measurements, its uncertainty is expected to shrink roughly as $1/\sqrt{N}$; the observed factor of $\sim$2.8 improvement from $N = 500$ to $N = 2000$ is close to the $\sqrt{4} = 2$ statistical expectation, with the remainder attributable to more pixels crossing detection threshold and better conditioning of the fit. The angular resolution follows the same jitter ordering, though its dependence on N is noisier: the median at $\sigma_t = 100$ ps is 15.0°, 7.0°, and 12.2° at $N = 500$, 1000, 2000 respectively. This non-monotonicity is a small-sample effect—six track trials per cell leave the angular median sensitive to individual difficult geometries—rather than a genuine reversal of the physics, as the wide 90th-percentile bars in Figure 4(B) make clear. We flag it explicitly rather than smoothing it over, because honestly representing the sampling noise of a modest trial count is part of pressure-testing the claim.

Third, *head–tail identification is perfect in this idealization.* The direction of motion is recovered correctly on all 36 track trials, at both jitter values and all light yields (Figure 4(C)). This is a genuinely favorable result for the marginalized-likelihood approach, but it warrants careful interpretation. For a $\sim$5 cm track, the particle traverses the segment in $t_{\mathrm{traversal}} \approx 5\,\mathrm{cm}/29.98\,\mathrm{cm\,ns}^{-1} \approx 167\,\mathrm{ps}$, which is comparable to the timing jitter. Head–tail discrimination relies on detecting the time ordering of emissions along the track against exactly this jitter; when σ_t approaches $t_{\mathrm{traversal}}$, the ordering begins to wash out, and one would expect head–tail accuracy to fall toward the 50% random-guess floor for shorter tracks or larger jitter. That the marginalized-likelihood fit maintains 100% here—a marked improvement over simpler first-arrival formulations, which degrade badly in this regime—reflects both the idealized clean timing of the simulation and the modest trial count; it should be read as “the estimator is capable in principle,” not as “head–tail is free.”

Table 1: Reconstruction performance across the parameter sweep. Median / mean over trials per (σ_t, N) cell; vertex, angular, and length errors from 8 vertex and 6 track trials each. Vertex and length in mm, angle in °, head–tail accuracy in %. The highlighted row is the nominal operating point used for the detector comparison.

σ_t (ps)	N	Vertex err (mm)	Angular err (deg)	Length err (mm)	Head–tail (%)
100	500	1.92 / 1.93	15.0 / 14.4	1.40 / 3.85	100
100	1000	1.31 / 1.28	7.0 / 6.8	2.36 / 2.51	100
100	2000	0.68 / 1.04	12.2 / 11.2	5.09 / 4.76	100
200	500	2.96 / 4.11	23.8 / 21.8	5.61 / 7.95	100
200	1000	3.63 / 3.58	22.4 / 19.9	6.26 / 6.89	100
200	2000	2.75 / 2.73	10.2 / 11.0	6.93 / 6.17	100

3.3 Diagnostic behavior of the estimators

Figure 5 shows the estimators in operation for representative events: reconstructed versus true vertex for a point source, reconstructed track line and endpoints against truth, and the distribution of timing residuals. The residual histogram is centered near zero with a width matching the injected jitter, confirming that the vertex fit is unbiased and that it has not absorbed timing noise into a spurious spatial shift. The reconstructed track line overlays the truth closely in direction while showing the larger scatter at the endpoints that the length- and endpoint-error metrics quantify—endpoints are intrinsically the hardest part of a track to fix, because the light yield per unit length is finite and the ends contribute the fewest constraining photons. Taken together, the diagnostics show two well-behaved estimators whose error budgets are dominated, as designed, by photon statistics and timing precision rather than by fitting pathologies.

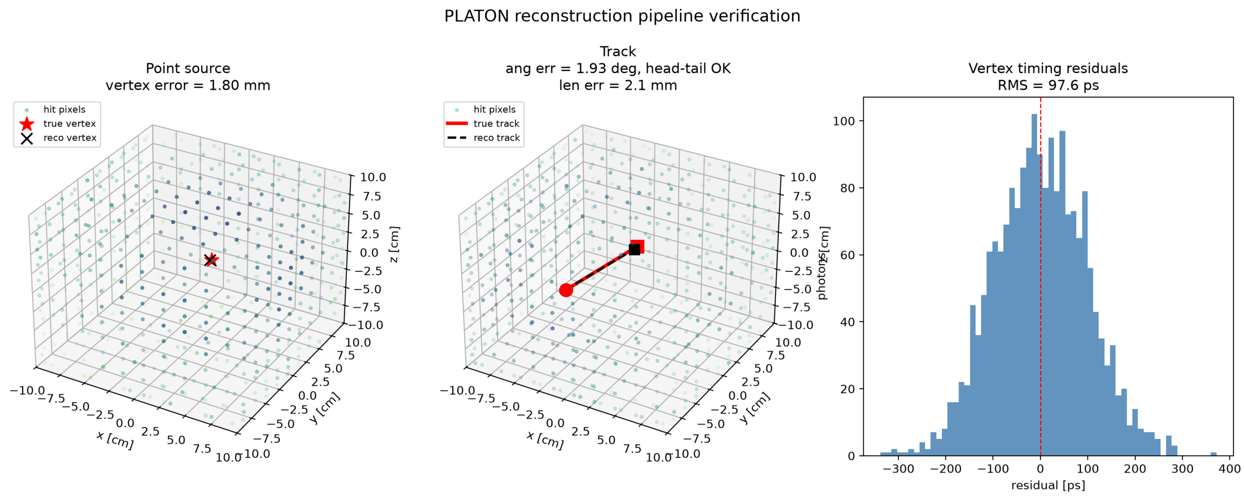

Figure 5: Reconstruction diagnostics. *Left:* true versus reconstructed vertex for a point source; the reconstructed position (marker) tracks the truth closely. *Center:* true track (line) with true and reconstructed endpoints; the fit recovers the direction accurately with larger residual scatter at the endpoints. *Right:* histogram of per-pixel timing residuals for the vertex fit, centered near zero with width set by the injected jitter, confirming an unbiased, well-calibrated estimator.

To summarize the quantitative outcome: at the nominal operating point of $\sigma_t = 100$ ps and $N = 1000$ detected photons, the toy PLATON reconstruction reaches a median vertex resolution of 1.31 mm, a median track-direction resolution of 7.0°, a median track-length error of 2.37 mm, and 100% head–tail accuracy, using only 600 surface channels. The best case in the sweep (100 ps, $N = 2000$) reaches a 0.68 mm median vertex resolution, and the worst case (200 ps, $N = 500$) still holds a 2.96 mm median vertex resolution and 23.8° direction. Sub-centimeter vertexing from timing alone, over a 20-centimeter monolith instrumented only on its surface, is the headline finding that the comparison below places in context.

3.4 Noise resilience: dark counts, correlated noise, and their mitigation

The results above assume a noise-free readout; the decisive question for a real device is whether they survive the dark counts, crosstalk, and afterpulses of an actual silicon photomultiplier. Figure 6 and Table 2 report the dark-count-rate sweep at the nominal operating point, reconstructing each event both from the raw noisy hits and from the coincidence-filtered hits. The picture that emerges is sharply asymmetric between the two estimators, and it is instructive precisely because of that asymmetry.

The vertex fit is intrinsically noise-robust. Its Huber loss down-weights the large residuals that scattered dark counts produce, so the unfiltered median vertex error holds near its noise-free value—1.57 mm at 100 kHz and 1.77 mm at 1 MHz, against 1.86 mm with no noise—and only collapses at the extreme 10 MHz level, where dark counts finally overwhelm the robust weighting and the unfiltered error jumps to 11.0 mm. The coincidence filter rescues exactly this failure: it restores the 10 MHz vertex resolution to 2.56 mm, a factor-of-four recovery, and keeps the filtered vertex error below 2.6 mm across the entire four-decade sweep. At the realistic 1 MHz operating point the filtered vertex resolution is 1.47 mm, statistically indistinguishable from the noise-free 1.31 mm of Table 1.

Tracking tells the opposite story: it is acutely fragile to noise, because the marginalized likelihood weights every hit and a surface sprinkled with dark counts is read as spurious off-axis light. Unfiltered, the median track-direction error explodes from 5.4° with no noise to roughly 60° as soon as any dark counts are present, and head–tail identification collapses toward and below the chance level—33%, 67%, and 17% correct at 10^5, 10^6, and 10^7 Hz respectively. Here the filter is not a refinement but a prerequisite: it brings the direction error back to 17.7° at 100 kHz and 31.4° at 1 MHz, and restores head–tail accuracy to 100% and 83%. In other words, coarse tracking in a monolith is impossible without noise cleaning and recovers usefully once it is applied, at least up to megahertz-scale dark rates.

Underlying both recoveries is a filter that does its job cleanly. Its true-positive rate—the fraction of genuine signal photons it keeps—sits at 98.5–99.2% across the sweep, comfortably above the 95% design target, so it is not throwing away the signal it is trying to preserve. Its dark-count rejection is 100% at 100 kHz, 95.9% at 1 MHz, and 91.2% at the punishing 10 MHz, meeting the $> 90\%$ target throughout. The two-stage design is what delivers this: the time-window cut removes the uniform dark pedestal and the delayed afterpulses, while the neighbor-coincidence cut eliminates the isolated dark counts that survive in-time by chance.

Two honest qualifications belong with these numbers. First, filtering is not free when there is little to filter: at zero and 100 kHz dark rate the filtered vertex and angular errors are marginally *worse* than unfiltered (1.70 versus 1.57 mm vertex at 100 kHz; 5.8 versus 5.4° angle at zero noise), simply because the filter discards $\sim$1% of true photons that the robust fit would happily have used. The filter earns its keep only when noise is genuinely present, which argues for enabling it adaptively rather than unconditionally. Second, at the extreme 10 MHz level the filtered *tracking* does not fully recover—the direction error remains near 57°—even though the vertex does; direction is a more delicate observable than position, and with only six track trials per level these track medians carry substantial sampling noise. The safe reading is that the vertex is robust to dark rates far beyond any realistic device, while coarse tracking is secure to roughly the megahertz scale and degrades beyond it.

Table 2: Reconstruction performance versus per-channel dark count rate. Median over 8 vertex and 6 track trials at the nominal operating point ($N = 1000$, $\sigma_t = 100$ ps, $p_{\mathrm{OCT}} = 0.15$, $p_{\mathrm{AP}} = 0.05$), reconstructed from raw (unfiltered) and coincidence-filtered hits. Vertex error in mm, angular error in °, head–tail (HT) accuracy in %. TPR is the true-photon retention and "DCR rej." the dark-count rejection of the filter.

	Vertex err (mm)		Angular err (deg)		HT acc (%)		Filter	
R_{DCR} (Hz)	Raw	Filt.	Raw	Filt.	Raw	Filt.	TPR	DCR rej.
0	1.86	1.85	5.4	5.8	100	100	0.985	—
10^5	1.57	1.70	60.6	17.7	33	100	0.992	1.00
10^6	1.77	1.47	60.0	31.4	67	83	0.991	0.96
10^7	11.0	2.56	51.4	56.6	17	83	0.992	0.91

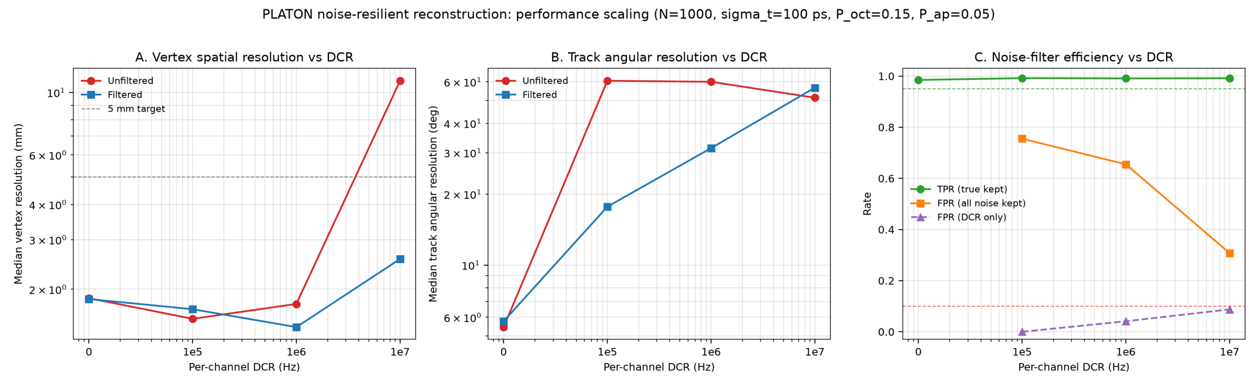

Figure 6: Noise-resilient reconstruction versus per-channel dark count rate. Nominal operating point ($N = 1000$, $\sigma_t = 100\,\text{ps}$, $p_{\text{OCT}} = 0.15$, $p_{\text{AP}} = 0.05$); medians over 8 vertex and 6 track trials per level. **(A)** Median vertex resolution: unfiltered (red) is robust up to 1 MHz but diverges to 11 mm at 10 MHz, whereas filtered (blue) stays below 2.6 mm across all four decades; the dashed line marks the 5 mm target. **(B)** Median track angular resolution: unfiltered tracking collapses to $\sim 60°$ once any dark counts appear, while the filter restores it to $\lesssim 31°$ up to 1 MHz. **(C)** Filter efficiency: true-photon retention (green) exceeds 98% throughout, dark-count rejection (purple, "DCR only") exceeds 91%, and the fraction of *all* noise retained (orange) is dominated by signal-like prompt crosstalk that the filter intentionally does not target.

4 Comparative Analysis Against Segmented Detectors

The question motivating this study is comparative: can a monolith reconstruct tracks "about as well as" finely segmented detectors? Table 3 answers it across six dimensions, contrasting the simulated PLATON operating point with three established, segmented alternatives—the LArTPC as exemplified by DUNE [2, 8], the segmented plastic scintillator tracker as exemplified by the T2K SuperFGD/3DST [3, 28], and the crystal-ring TOF-PET scanner [4, 5]. We stress at the outset that only the PLATON column is simulated in this work; the other three columns are representative, literature-scale values assembled for context, and the comparison is therefore one of design philosophy and order of magnitude, not a controlled head-to-head measurement.

Read across the rows, the table refines the "about as well as" claim into something more precise and more honest. On *vertex localization*, PLATON is genuinely competitive: its simulated 1.3 mm vertex resolution sits squarely in the millimeter band that LArTPCs and TOF-PET achieve, and it does so from timing alone. On *full tracking*, it is not: a 7° direction and a coarse length are no substitute for the pixel-level trajectory, multi-particle topology, calorimetric energy, and dE/dx particle identification that a LArTPC or a segmented plastic tracker delivers. The fair summary is that PLATON trades fine tracking for dramatically lower complexity—roughly two orders of magnitude fewer channels than a segmented tracker of similar size—while retaining sub-centimeter localization of the interaction point.

The decisive entry is the channel-count-and-scalability pair, because it is here that the two design philosophies diverge fastest. A segmented detector instruments the bulk, so its channel count, wiring, front-end electronics, and dead material all scale with the active volume, $\propto L^3$. PLATON instruments only the boundary, so its channel count scales with surface area, $\propto L^2$. The ratio of these, $\propto L$, means the monolithic advantage is not fixed but *grows linearly with detector size*: the bigger the detector one wants, the larger the fractional saving from moving instrumentation to the surface. For a small device the saving is modest and may not justify the loss of tracking detail; for a very large one it can be the difference between affordable and impossible. This scaling argument, rather than any single resolution number, is the real content of the comparison, and it

Table 3: PLATON versus conventional segmented detectors. PLATON figures are from the toy Monte Carlo at $\sigma_t = 100\,\text{ps}$, $N = 1000$. All other figures are representative literature-scale values for context, not simulated here.

Dimension	PLATON (this design)	LArTPC (DUNE)	Segmented plastic (T2K 3DST)	TOF-PET
Spatial resolution	$\sim$1.3 mm vertex from TOF trilateration; no intrinsic segmentation	$\sim$1 mm to 5 mm (wire pitch / drift sampling); mm-scale 3D imaging with calorimetry and dE/dx ID	$\sim$cube pitch (1 cm); sub-cube via light sharing; true 3D voxelization	$\sim$2 mm to 4 mm (crystal pitch); TOF sharpens localization along the line of response
Timing resolution	$\sim$100 ps per-photon jitter; timing *is* the imaging mechanism	ns-scale event timing; spatial info from $\sim$ms drift, not ps photon timing	$\sim$0.5 ns to 1 ns per hit; timing for hit association, not imaging	$\sim$200 ps to 400 ps FWHM coincidence; timing localizes the annihilation point (same physics family)
Tracking capability	Coarse: direction $\sim 7^\circ$, length $\sim$2.4 mm; head–tail limited by $t_{\text{traversal}}$ vs jitter	Excellent: full 3D trajectories, multi-particle topology, dE/dx, Bragg-peak head–tail	Excellent: fine-grained 3D tracks, neutron tagging via inter-cube TOF, clear topology	None: reconstructs annihilation-point density (image), not trajectories
Channel count	Very low: $\sim$600 surface channels for a 20 cm cube; scales as area (L^2)	High: 10^5–10^6 wires/pixels; cryogenics, HV, ultra-pure argon	Very high: 10^5–10^6 cubes $\times$ fibers $\times$ SiPMs; wiring scales as volume (L^3)	High: $\sim 10^4$ crystals + SiPMs in a ring; mature but costly electronics
Cost & scalability	Potentially cheap per unit volume; cost grows with surface, not volume	Scales to kilotonne masses but with major cryogenic/purity cost	Cost grows steeply with volume; hard to scale economically	Expensive per unit; scaling means bigger/denser rings
Primary limitations	Relies on $\sim$100 ps timing; degrades as jitter $\to t_{\text{traversal}}$; scattering/reflections not yet modeled; monolithic pile-up	Slow drift (poor pile-up timing), cryogenic complexity, argon purity, large HV	Channel count/wiring explode with size; dead material; readout cost dominates	No tracking; expensive fast crystals; limited field of view

points directly at where the design is most likely to matter first.

Two important asymmetries qualify the comparison. TOF-PET belongs to the *same physics family* as PLATON—it already localizes annihilation points along a line of response using picosecond coincidence timing [29–31]—so the conceptual distance PLATON must travel to improve on it is short, but the entrenched, highly optimized crystal-ring architecture is a demanding incumbent. The LArTPC and segmented plastic trackers, by contrast, offer tracking that PLATON simply cannot match, so the monolith is not a replacement for them but a potential complement: a cheap, large, coarse-vertexing volume where full tracking is unnecessary or unaffordable.

5 Discussion and Technical Brief

5.1 Where the design pays off first: neutrino experiments before PET

The natural instinct, given PLATON's kinship with TOF-PET, is to expect medical imaging to be the first beneficiary. The scaling argument of the previous section argues the opposite. The advantage that most favors a monolith—surface (L^2) rather than volume (L^3) instrumentation cost—is decisive precisely where detectors are largest, and the largest optical detectors in physics are neutrino experiments, not PET scanners. A large neutrino near- or far-detector must instrument an enormous fiducial volume, and the dominant cost of the segmented approaches—LArTPC wiring and cold electronics, or the SuperFGD's per-cube fibers and SiPMs—climbs with that volume. If sub-centimeter vertex localization and coarse topology over a huge volume are sufficient for the physics program, and for many neutrino-interaction-vertexing and event-classification tasks they are, then a monolith that captures the vertex cheaply while forgoing fine tracking is exactly the right trade. The simulated 1.3 mm vertex resolution is already competitive with segmented approaches for vertex localization, at perhaps 10^{-2} of the channel count.

PET is a secondary, longer-term opportunity for three concrete reasons. First, the cost asymmetry that most favors PLATON is muted at PET's smaller physical scale, where a monolithic bore is not dramatically larger than a crystal ring. Second, TOF-PET is a mature, heavily optimized, and regulated clinical product; displacing it requires clearing a high bar of image quality, quantitative accuracy, and quality assurance [4, 32], and the incremental gain from a monolithic geometry must compete against decades of crystal-ring refinement. Third, PET's resolution is ultimately bounded by physics that PLATON does not address—positron range and annihilation-photon acolinearity set a floor on how well any 511-keV system can localize the source [33, 34]—so better detector timing yields diminishing returns on the final image. The honest recommendation is therefore that the concept should be prototyped and proven in a large-volume neutrino context first, where its structural advantage is largest and the resolution requirements are best matched to what it delivers, and only later revisited for cost-reduced or large-bore PET once the underlying optical physics has been validated experimentally.

It is worth situating PLATON among the optical neutrino detectors it would join. Large liquid-scintillator experiments—KamLAND [35], Borexino [36], and the 20-kilotonne JUNO [37, 38]—already reconstruct events from photomultiplier hit patterns and timing over volumes far larger than our toy cube, and hybrid Cherenkov–scintillation concepts such as THEIA [39] and water-Cherenkov detectors like Super-Kamiokande [40] extract direction and topology from surface optical sensors. PLATON's distinctive proposition relative to these is its explicit reliance on picosecond photon timing and light-field angular information to trilaterate a compact vertex, rather than on the coarser timing and pure intensity imaging of a photomultiplier sphere. The machine-learning position-reconstruction methods proven in monolithic PET [9–11], and the deployment of fast neural inference on detector front-ends [41], offer a ready toolkit for scaling the reconstruction beyond the analytic estimators used here.

5.2 What this back-of-the-envelope simulation can and cannot settle

Pressure-testing a claim means being as clear about the limits of the test as about its results. This study establishes two things cleanly. In the idealized, non-scattering, straight-line-time-of-flight limit, the core physics of monolithic timing reconstruction is favorable—sub-centimeter vertexing and coarse tracking are achievable from a sparse surface of fast sensors, with resolution scaling exactly as a timing estimator should. And, with the addition of a validated silicon-photomultiplier noise model, that reconstruction is robust to realistic dark counts, crosstalk, and afterpulsing once a spatial–temporal coincidence filter is applied: the vertex resolution survives dark rates far beyond any practical device, and coarse tracking survives to the megahertz scale. What the study *cannot* settle is whether the clean timing that both estimators assume survives contact with a real optical medium, and several deliberately omitted *optical-transport* effects each threaten the central assumption of Equation (2).

The most serious omission is *optical scattering and bulk absorption.* Real scintillation photons do not travel in straight lines; Rayleigh and Mie scattering diffuse them, so the first photons to arrive at a pixel may have taken indirect paths, and the tidy relationship $t_{\mathrm{prop}} = d/v$ that underpins trilateration is smeared. In a sufficiently scattering medium the arrival-time distribution becomes diffusive rather than ballistic, and the vertex resolution would degrade well beyond what the jitter-only sweep predicts. Closely related is the omission of *surface reflection*—Fresnel reflection and total internal reflection at the cube faces, which in a real block sends a large fraction of photons on multiply reflected paths before detection, corrupting precisely the early, direct-path timing that the reconstruction weights most heavily. A validated optical model of exactly these boundary effects is what distinguishes a credible monolithic detector simulation [13], and its absence here is the single largest gap between this toy and a real device.

A further omission adds irreducible timing spread that the idealization excludes. *Finite scintillator kinetics*—the nonzero rise and decay time of the light emission—means that photons are not emitted instantaneously at t_0 but over a temporal profile, adding a physical jitter on top of the sensor jitter that no electronics can remove; for slower scintillators this can dominate the timing budget. Sensor dark counts, afterpulsing, and optical crosstalk, by contrast, are no longer held out: Section 3.4 showed that when they are modeled at literature-representative rates [18, 19] and mitigated by the coincidence filter, the reconstruction remains resilient. Two caveats attach even to that favorable result, however. First, the filter assumes the signal arrives as a compact, sharply timed burst; the very optical-transport effects held out above (scattering and reflection) would broaden that burst in time and blur the illuminated spatial cluster, eroding precisely the space–time contrast on which both filter stages depend, so the noise resilience demonstrated here is itself conditional on clean transport. Second, our detection model is a simple binomial efficiency with per-photon Gaussian jitter; it omits wavelength-dependent photodetection efficiency, transit-time spread, and saturation, and it treats the geometry as a perfect cube with point-like pixels, homogeneous index, and no dead material or gaps.

The physics of the source is likewise idealized. We model energy deposition as uniform isotropic emission along a straight segment, whereas real charged-particle tracks curve and scatter, deposit energy non-uniformly, and—for the electromagnetic and hadronic showers typical of neutrino interactions—branch into complex multi-particle topologies that a single-track fit cannot represent. A real reconstruction would have to contend with overlapping tracks and pile-up in dense events, a regime in which the monolith's lack of segmentation is a genuine disadvantage relative to a voxelized tracker.

Two statistical caveats round out the honest accounting. The trial counts per cell are modest (6–8), so the reported medians carry non-negligible sampling uncertainty; the non-monotonic angular trend in Figure 4(B) and the uniformly perfect head–tail accuracy in Figure 4(C) are both best read as consistent with, but not a precise measurement of, the underlying resolution. In particular, the 100% head–tail result of the noise-free sweep should not be over-interpreted: it reflects an idealized clean-timing regime and a limited sample. The noise sweep of Section 3.4 makes this concrete—head–tail accuracy does collapse toward chance once dark counts are present and recovers only partially after filtering at the highest rates—which is exactly the fragility the traversal-time-versus-jitter argument predicts, and a reminder that the perfect noise-free figure is a ceiling rather than an expectation.

None of these caveats overturns the central finding, but together they define the boundary of what a back-of-the-envelope study can claim. The appropriate conclusion is conditional: *if* the clean timing assumed here survives realistic optical transport, PLATON's reconstruction physics supports sub-centimeter vertexing at a fraction of the channel count; whether it does survive is the question that only a full optical-transport simulation, and ultimately a bench prototype, can answer.

5.3 Recommended next step

The logical successor to this study is a full optical-transport Monte Carlo that reinstates the optical effects still held out here—bulk scattering and absorption with realistic attenuation lengths, Fresnel and total-internal-reflection boundary physics, and finite scintillator rise and decay times—built on a validated toolkit such as GEANT4 with its optical-photon physics [13, 42]. Crucially, that simulation should carry the sensor-noise model and coincidence filter developed here forward into the realistic-transport regime, because the two interact: scattering and reflection broaden the signal burst in time and smear its spatial cluster, degrading exactly the space–time contrast the filter relies on, so the noise resilience reported in Section 3.4 must be re-established, not assumed, once transport is realistic. Such a simulation would answer the decisive question this toy cannot: how much of the 1.3 mm idealized vertex resolution remains once photons diffuse and reflect, and whether the filter still recovers it in the presence of dark counts. In parallel, the reconstruction should be advanced from the analytic estimators used here to a learned model—the gradient-boosted and neural regressors already proven on monolithic PET data [9–11]—trained on the realistic optical simulation and, ultimately, on a bench prototype illuminated by a pulsed source. Only after that chain closes can the original claim—that a single solid block reconstructs tracks about as well as a segmented detector—be either substantiated or retired.

6 Conclusions

We set out to independently pressure-test the PLATON claim that a single, unsegmented scintillator block, watched only on its surface by light-field cameras and single-photon sensors, can reconstruct particle interactions in three dimensions competitively with finely segmented detectors. Using a transparent first-principles Monte Carlo of optical-photon emission, straight-line propagation, and timed surface detection in a 20 cm cube with 600 channels, and two purpose-built estimators—a robust time-of-flight vertex trilateration solver and a marginalized-maximum-likelihood track solver—we quantified how well the interaction vertex and a short track can be recovered, and how that recovery degrades as light yield and sensor timing worsen.

The results are encouraging but bounded. At a realistic operating point (100 ps timing jitter, 1000 detected photons), the reconstruction achieves a median vertex resolution of 1.31 mm, a median track-direction resolution of 7.0°, a median track-length error of 2.37 mm, and correct head–tail identification, using roughly two orders of magnitude fewer channels than a comparable segmented tracker. Resolution improves with photon statistics, close to the $1/\sqrt{N}$ expectation, and degrades sharply as jitter approaches the particle traversal time—the unmistakable signature of a timing-limited estimator, and a clear statement that the concept's viability rests entirely on fast timing. On vertex localization the monolith is genuinely competitive with segmented detectors; on full tracking it is not, and the fair characterization is that it trades tracking detail for a decisive advantage in channel count and cost scaling.

Critically, that performance is not confined to an idealized noise-free readout. Adding a validated physical model of the three dominant silicon-photomultiplier noise sources—dark counts, optical crosstalk, and afterpulsing—we found the outlier-robust vertex fit to be intrinsically resilient and, with a lightweight spatial–temporal coincidence filter that retains $\sim$99% of signal while rejecting $\gtrsim$96% of dark counts at 1 MHz, the reconstruction holds its 1.5 mm vertex resolution and recovers coarse tracking that dark counts would otherwise destroy. Noise, in short, is a solved problem for the vertex and a manageable one for tracking up to megahertz-scale dark rates.

That structural advantage—surface rather than volume instrumentation, L^2 rather than L^3 cost—grows with detector size, which is why we conclude that large neutrino experiments, not PET scanners, are the most plausible first beneficiary: the structural saving is largest where the detector is largest, and coarse but cheap volumetric vertexing is well matched to many neutrino physics goals. We are equally clear about what the study cannot settle. Even with sensor noise now modeled and mitigated, it still omits optical scattering, surface reflection, and finite scintillator kinetics, each of which threatens the clean straight-line timing—and the compact signal burst—that make both the reconstruction and the noise filter work; the reported resolutions are therefore best-case upper bounds. The central contribution is to show that the core physics is favorable and robust to realistic sensor noise, and to identify precisely the one experiment that must come next: a full optical-transport simulation that tests whether millimeter timing reconstruction, and the noise filter that protects it, survive realistic light transport. Until that test is done, the PLATON claim is promising and physically motivated, but not yet proven.

Data and Code Availability

The simulation and reconstruction code, the parameter-sweep results, and the figures underlying this report were generated as part of this study and are available in the accompanying project directory. The forward model, reconstruction estimators, sensor-noise model, and coincidence filter are implemented in `simulation.py` and `reconstruct.py`; the light-yield/jitter and dark-count-rate sweeps are driven by `test_reconstruction.py`, `test_reconstruction_noise.py`, and `generate_summary.py`, and the noise model is independently validated by `verify_noise.py`. Sweep results are stored in `performance_summary.json`, `noise_performance_summary.json`, `noise_verification.json`, `detector_comparison.json`, and `technical_brief.json`. All runs use a fixed random seed (42) for reproducibility.

Methods Note

Simulation and analysis were performed in Python (NumPy 2.4, SciPy 1.18, Matplotlib 3.11). The forward model passed twelve independent physics validation checks prior to use; the marginalized-likelihood gradient was verified against finite differences to $\sim 10^{-11}$. The sensor-noise model was validated separately: the dark-count multiplicity is Poisson to within a Fano factor of 1.001, the optical-crosstalk hits are correctly localized to edge-adjacent pixels with a measured probability of

0.148 against an injected 0.15, and the afterpulse delays follow an exponential with a maximum-likelihood time constant of 15.05 ns against an injected 15 ns (Kolmogorov–Smirnov $p = 0.90$).

References

[1] R. Acciarri and others (MicroBooNE Collaboration). Design and construction of the microboone detector. *Journal of Instrumentation*, 12(2):P02017, 2017. doi: 10.1088/1748-0221/12/02/P02017.

[2] B. Abi and others (DUNE Collaboration). First results on protodune-sp liquid argon time projection chamber performance from a beam test at the cern neutrino platform. *Journal of Instrumentation*, 15(12):P12004, 2020. doi: 10.1088/1748-0221/15/12/P12004.

[3] A. Blondel et al. The superfgd prototype charged particle beam tests. *Journal of Instrumentation*, 15(12):P12003, 2020. doi: 10.1088/1748-0221/15/12/P12003.

[4] Suleman Surti and Joel S. Karp. Update on latest advances in time-of-flight pet. *Physica Medica*, 80:251–258, 2020. doi: 10.1016/j.ejmp.2020.10.031.

[5] Joyce van Sluis, Johan de Jong, Jan Schaar, Walter Noordzij, Paul van Snick, Rudi Dierckx, Ronald Borra, Antoon Willemsen, and Ronald Boellaard. Performance characteristics of the digital biograph vision pet/ct system. *Journal of Nuclear Medicine*, 60(7):1031–1036, 2019. doi: 10.2967/jnumed.118.215418.

[6] P. Abratenko and others (MicroBooNE Collaboration). Neutrino event selection in the microboone liquid argon time projection chamber using wire-cell 3d imaging, clustering, and charge-light matching. *Journal of Instrumentation*, 16(6):P06043, 2021. doi: 10.1088/1748-0221/16/06/P06043.

[7] P. Abratenko and others (MicroBooNE Collaboration). Semantic segmentation with a sparse convolutional neural network for event reconstruction in microboone. *Physical Review D*, 103 (5):052012, 2021. doi: 10.1103/PhysRevD.103.052012.

[8] B. Abi and others (DUNE Collaboration). Deep underground neutrino experiment (dune), far detector technical design report, volume i: Introduction to dune. *Journal of Instrumentation*, 15(8):T08008, 2020. doi: 10.1088/1748-0221/15/08/T08008.

[9] Florian Müller, David Schug, Patrick Hallen, Jan Grahe, and Volkmar Schulz. Gradient tree boosting-based positioning method for monolithic scintillator crystals in positron emission tomography. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2(5):411–421, 2018. doi: 10.1109/TRPMS.2018.2837738.

[10] Javier Balibrea-Correa, Jorge Lerendegui-Marco, Ion Ladarescu, Carlos Guerrero, Jose Luis Tain, and Cesar Domingo-Pardo. Machine learning aided 3d-position reconstruction in large lacl3 crystals. *Nuclear Instruments and Methods in Physics Research Section A*, 1001:165249, 2021. doi: 10.1016/j.nima.2021.165249.

[11] Pietro Carra, Maria Giuseppina Bisogni, Esther Ciarrocchi, Matteo Morrocchi, Giancarlo Sportelli, Valeria Rosso, and Nicola Belcari. A neural network-based algorithm for simultaneous event positioning and timestamping in monolithic scintillators. *Physics in Medicine & Biology*, 67(13):135001, 2022. doi: 10.1088/1361-6560/ac72f2.

[12] Amirhossein Sanaat and Habib Zaidi. Depth of interaction estimation in a preclinical pet scanner equipped with monolithic crystals coupled to sipms using a deep neural network. *Applied Sciences*, 10(14):4753, 2020. doi: 10.3390/app10144753.

[13] Dennis J. J. van der Laan, Dennis R. Schaart, Marnix C. Maas, Freek J. Beekman, Peter Bruyndonckx, and Carel W. E. van Eijk. Optical simulation of monolithic scintillator detectors using gate/geant4. *Physics in Medicine & Biology*, 55(6):1659–1675, 2010. doi: 10.1088/0031-9155/55/6/009.

[14] Edward H. Adelson and John Y. A. Wang. Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):99–106, 1992. doi: 10.1109/34.121783.

[15] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*, pages 31–42, 1996. doi: 10.1145/237170.237199.

[16] Mythra Varun Nemallapudi, Stefan Gundacker, Paul Lecoq, and Etiennette Auffray. Single photon time resolution of state of the art sipms. *Journal of Instrumentation*, 11(10):P10016, 2016. doi: 10.1088/1748-0221/11/10/P10016.

[17] Stefan Gundacker and Arjan Heering. The silicon photomultiplier: fundamentals and applications of a modern solid-state photon detector. *Physics in Medicine & Biology*, 65(17):17TR01, 2020. doi: 10.1088/1361-6560/ab7b2d.

[18] Paolo Finocchiaro, Alfio Pappalardo, Luigi Cosentino, Massimiliano Belluso, Sergio Billotta, Giovanni Bonanno, and Salvatore Di Mauro. Features of silicon photo multipliers: Precision measurements of noise, cross-talk, afterpulsing, detection efficiency. *IEEE Transactions on Nuclear Science*, 56(3):1033–1041, 2009. doi: 10.1109/TNS.2009.2014308.

[19] Claudio Piemonte and Alberto Gola. Overview on the main parameters and technology of modern silicon photomultipliers. *Nuclear Instruments and Methods in Physics Research Section A*, 926:2–15, 2019. doi: 10.1016/j.nima.2018.11.119.

[20] Matthew F. Bieniosek, Joshua W. Cates, and Craig S. Levin. Achieving fast timing performance with multiplexed SiPMs. *Physics in Medicine & Biology*, 61(7):2879–2892, 2016. doi: 10.1088/0031-9155/61/7/2879.

[21] Emilie Roncali and Simon R. Cherry. Application of silicon photomultipliers to positron emission tomography. *Annals of Biomedical Engineering*, 39(4):1358–1377, 2011. doi: 10.1007/s10439-011-0266-9.

[22] S. Vinogradov. Analytical models of probability distribution and excess noise factor of solid state photomultiplier signals with crosstalk. *Nuclear Instruments and Methods in Physics Research Section A*, 695:247–251, 2012. doi: 10.1016/j.nima.2011.11.086.

[23] J. Rosado, V. M. Aranda, F. Blanco, and F. Arqueros. Modeling crosstalk and afterpulsing in silicon photomultipliers. *Nuclear Instruments and Methods in Physics Research Section A*, 787: 153–156, 2015. doi: 10.1016/j.nima.2014.11.080.

[24] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, 1982. doi: 10.1109/TMI.1982.4307558.

[25] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963. doi: 10.1137/0111030.

[26] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. doi: 10.1145/358669.358692.

[27] Stefan Gundacker, Rosana Martinez Turtos, Nicolaus Kratochwil, Rosalinde Hendrika Pots, Marco Paganoni, Paul Lecoq, and Etiennette Auffray. Experimental time resolution limits of modern sipms and tof-pet detectors exploring different scintillators and cherenkov emission. *Physics in Medicine & Biology*, 65(2):025001, 2020. doi: 10.1088/1361-6560/ab63b4.

[28] K. Abe and others (T2K Collaboration). The t2k experiment. *Nuclear Instruments and Methods in Physics Research Section A*, 659(1):106–135, 2011. doi: 10.1016/j.nima.2011.06.067.

[29] Stefan E. Brunner and Dennis R. Schaart. Bgo as a hybrid scintillator / cherenkov radiator for cost-effective time-of-flight pet. *Physics in Medicine & Biology*, 62(11):4421–4439, 2017. doi: 10.1088/1361-6560/aa6a49.

[30] Sun Il Kwon, Alberto Gola, Alessandro Ferri, Claudio Piemonte, and Simon R. Cherry. Dual-ended readout of bismuth germanate to improve timing resolution in time-of-flight pet. *Physics in Medicine & Biology*, 64(10):105007, 2019. doi: 10.1088/1361-6560/ab18da.

[31] Paul Lecoq, Christian Morel, John O. Prior, Dimitris Visvikis, Stefan Gundacker, Etiennette Auffray, Peter Král, et al. Roadmap toward the 10 ps time-of-flight pet challenge. *Physics in Medicine & Biology*, 65(21):21RM01, 2020. doi: 10.1088/1361-6560/ab9500.

[32] Dennis R. Schaart. Physics and technology of time-of-flight pet detectors. *Physics in Medicine & Biology*, 66(9):09TR01, 2021. doi: 10.1088/1361-6560/abee56.

[33] Craig S. Levin and Edward J. Hoffman. Calculation of positron range and its effect on the fundamental limit of pet system spatial resolution. *Physics in Medicine and Biology*, 44(3): 781–799, 1999. doi: 10.1088/0031-9155/44/3/019.

[34] William W. Moses. Fundamental limits of spatial resolution in pet. *Nuclear Instruments and Methods in Physics Research Section A*, 648:S236–S240, 2011. doi: 10.1016/j.nima.2010.11.092.

[35] K. Eguchi et al. First results from kamland: Evidence for reactor antineutrino disappearance. *Physical Review Letters*, 90(2):021802, 2003. doi: 10.1103/PhysRevLett.90.021802.

[36] G. Alimonti et al. The borexino detector at the laboratori nazionali del gran sasso. *Nuclear Instruments and Methods in Physics Research Section A*, 600(3):568–593, 2009. doi: 10.1016/j.nima.2008.11.076.

[37] A. Abusleme and others (JUNO Collaboration). Juno physics and detector. *Progress in Particle and Nuclear Physics*, 123:103927, 2022. doi: 10.1016/j.ppnp.2021.103927.

[38] A. Abusleme, others (JUNO, and Daya Bay Collaborations). Optimization of the juno liquid scintillator composition using a daya bay antineutrino detector. *Nuclear Instruments and Methods in Physics Research Section A*, 988:164823, 2021. doi: 10.1016/j.nima.2020.164823.

[39] M. Askins and others (Theia Collaboration). Theia: an advanced optical neutrino detector. *The European Physical Journal C*, 80(5):416, 2020. doi: 10.1140/epjc/s10052-020-7977-8.

[40] Y. Fukuda et al. The super-kamiokande detector. *Nuclear Instruments and Methods in Physics Research Section A*, 501(2–3):418–462, 2003. doi: 10.1016/S0168-9002(03)00425-X.

[41] Javier Duarte et al. Fast inference of deep neural networks in fpgas for particle physics. *Journal of Instrumentation*, 13(7):P07027, 2018. doi: 10.1088/1748-0221/13/07/P07027.

[42] S. Agostinelli et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A*, 506(3):250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8.