Supplementary Document for "Asynchronous Single-Photon 3D Imaging"

Anant Gupta, Atul Ingle, Mohit Gupta.

{anant,ingle,mohitg}@cs.wisc.edu

Supplementary Note 1. Asynchronous Image Formation Model and MLE Waveform Estimator

In this supplementary note we derive the Poisson-Multinomial histogram model of Eq. (S3) in the main paper. We then derive the generalized Coates's estimator (Eq. (4)), which is the maximum likelihood estimator (MLE) of the photon flux waveform in asynchronous acquisition. Scene depth is computed by locating the peak of the estimated photon flux waveform.

Joint Distribution of Measured Histogram

Here we derive the joint distribution of the histogram $(N_1, \ldots, N_B, N_{B+1})$ measured in the asynchronous acquisition mode. Recall that the $B+1^{\text{st}}$ bin is added for mathematical convenience to record SPAD cycles with no detected photons. For the l^{th} SPAD cycle we define a one-hot random vector $(O_{l,1}, O_{l,2}, \ldots, O_{l,B}, O_{l,B+1})$ that stores the bin index where the photon was recorded. Since the SPAD detects at most one photon per laser cycle, $(O_{l,i})_{i=1}^{B+1}$ contains zeroes everywhere except at the bin index corresponding to the photon detection. Its joint distribution is given by a categorical distribution [Suppl. Ref. 1]:

$$(O_{l,i})_{i=1}^{B+1} \sim (B+1)\text{-Categorical}\big((p_{l,i})_{i=1}^{B+1}\big). \tag{S1}$$

The final histogram of photon counts is obtained by summing these one-hot vectors over all laser cycles:

$$N_i = \sum_{l=1}^{L} O_{l,i}\,. \tag{S2}$$

Since N_i is a sum of L different $(B+1)$-Categorical random variables, the joint distribution is given by a Poisson-Multinomial Distribution [Suppl. Ref. 2]:

$$(N_i)_{i=1}^{B+1} \sim (L, B+1)\text{-PMD}\Big((p_{l,i})_{1\le l\le L, 1\le i\le B+1}\Big). \tag{S3}$$

The expected number of photon counts $\mathbb{E}[N_i]$ in the i^{th} bin is $\sum_{l=1}^{L} p_{l,i}$. Note that in the synchronous case, this reduces to a multinomial distribution because the one-hot random vector defined here no longer depends on the SPAD cycle index l.

Derivation of the Generalized Coates's Estimator for Asynchronous SPAD LiDAR (Eq. (4))

In this section, we derive the MLE $(\widehat{r}_i)_{i=1}^{B}$ of the true waveform $(r_i)_{i=1}^{B}$ for the asynchronous acquisition model and show that it is equal to the generalized Coates's estimator described in the main text. We assume that for each SPAD cycle $1 \le l \le L$, the TCSPC system stores one-hot random vectors $(O_{l,i})_{1\le i\le B+1}$.

For future reference, we define $J_{l,i}$ to be the set of bin indices preceding i in the l^{th} cycle, in a modulo-B sense[8]:

$$J_{l,i} = \begin{cases} \{s_l+1, \ldots, \underbrace{B, 1}_{\text{wrap around}}, \ldots, i-1\}, & \text{for } i \le s_l \\ \{s_l+1, \ldots, i-1\}, & \text{for } i > s_l\,. \end{cases} \tag{S4}$$

In the lth laser cycle, the joint distribution of $(O_{l,i})_{i=1}^{B+1}$ is given by the categorical distribution in Eq. (S1). Therefore, the likelihood function of the photon incidence probabilities $(q_i)_{i=1}^{B}$ is given by:

[8] For example, suppose $B=8$ and $s_l=3$. Then, $J_{l,7}=\{4,5,6\}$, and $J_{l,2}=\{4,5,6,7,8,1\}$.

$$
\begin{aligned}
\mathcal{L}(q_1, q_2, ..., q_B) &= \mathbb{P}\left((O_{1,i})_{i=1}^{B+1}, (O_{2,i})_{i=1}^{B+1}, ..., (O_{L,i})_{i=1}^{B+1} | q_1, q_2, ..., q_B\right) \\
&\overset{(a)}{=} \prod_{l=1}^{L} \mathbb{P}\left((O_{l,i})_{i=1}^{B+1} | q_1, q_2, ..., q_B\right) \\
&\overset{(b)}{=} \prod_{l=1}^{L} \prod_{i=1}^{B+1} p_{l,i}^{O_{l,i}} \\
&\overset{(c)}{=} \prod_{l=1}^{L} \prod_{i=1}^{B+1} \left(q_i \prod_{j \in J_{l,i}} (1 - q_j)\right)^{O_{l,i}} \\
&\overset{(d)}{=} \prod_{i=1}^{B+1} q_i^{\sum_{l=1}^{L} O_{l,i}} \left[\prod_{l=1}^{L} \prod_{i=1}^{B+1} \left(\prod_{j=1}^{B+1} (1 - q_j)^{\mathbb{1}(j \in J_{l,i})}\right)^{O_{l,i}}\right] \\
&\overset{(e)}{=} \prod_{i=1}^{B+1} q_i^{N_i} \prod_{j=1}^{B+1} \prod_{l=1}^{L} (1 - q_j)^{\sum_{i=1}^{B+1} \mathbb{1}(j \in J_{l,i}) O_{l,i}} \\
&\overset{(f)}{=} \prod_{i=1}^{B+1} q_i^{N_i} \prod_{j=1}^{B+1} \prod_{l=1}^{L} (1 - q_j)^{D_{l,j} - O_{l,j}} \\
&\overset{(g)}{=} \prod_{i=1}^{B+1} q_i^{N_i} (1 - q_i)^{D_i - N_i}
\end{aligned} \tag{S5}
$$

where (a) holds because measurements in different cycles are conditionally independent given the shift sequence; (b) follows from the definition of categorical distribution and the fact that for any fixed l, $O_{l,i} = 1$ for exactly one $1 \leq i \leq B+1$; (c) follows from Eq. (3); (d) uses the notation $\mathbb{1}$ for the indicator function; (e) uses the definition of N_i from Eq. (S2); (f) follows from the lemma proved below; (g) follows from Eq. (S2) and Def. 1 and rearrangement of the terms in the preceding product.

Since the likelihood is factorizable in q_i, we can calculate the MLE element-wise as:

$$
\begin{aligned}
\widehat{q_i} &= \arg\max_{q_i} \; q_i^{N_i} (1 - q_i)^{D_i - N_i} \\
&= \frac{N_i}{D_i}.
\end{aligned}
$$

Since $q_i = 1 - e^{-r_i}$, by the functional invariance property of the MLE [Suppl. Ref. 4], the MLE for the photon flux waveform r_i is given by the generalized Coates's estimator of Eq. (4). We estimate scene depth by locating the peak of the estimated waveform: $\widehat{\tau} = \arg\max_i \widehat{r}_i$.

Finally, we prove the following lemma that was used in step (f) in the derivation above.

Lemma (Proof of step (f)). For $1 \leq j \leq B+1$ and $1 \leq l \leq L$

$$
\sum_{i=1}^{B+1} \mathbb{1}(j \in J_{l,i}) O_{l,i} = D_{l,j} - O_{l,j}
$$

Proof. Let i^* denote the bin index where the photon was detected in the l^{th} SPAD cycle, i.e., $O_{l,k} = 1$ iff $k = i^*$ and 0 otherwise. Then:

$$
\sum_{i=1}^{B+1} \mathbb{1}(j \in J_{l,i}) O_{l,i} = \mathbb{1}(j \in J_{l,i^*}) O_{l,i^*}.
$$

If $j = i^*$, then by definition $j \notin J_{l,i^*}$. Therefore, LHS = 0. Also, $D_{l,j} = O_{l,j} = 1$ in this case, so RHS = 0. If $j \neq i^*$, there are two cases. Case 1: $j \in J_{l,i^*}$. Then $D_{l,j} = 1$ and $O_{l,j} = 0$. Therefore, LHS = RHS = 1. Case 2: $j \notin J_{l,i^*}$. Then $D_{l,j} = O_{l,j} = 0$. Therefore, LHS = RHS = 0. □

Supplementary Note 2. Proofs of Results 1 and 2

In this section we provide detailed mathematical proofs for two key theoretical results in the main paper. Recall that we use an upper bound on the probability of depth error (ℓ_0 error) as a surrogate for RMSE. Result 1 establishes the importance of a constant expected denominator sequence. It shows that a shifting strategy that minimizes an upper bound on the ℓ_0 error must have a denominator sequence that is constant (on average) for all histogram bins. Result 2 shows that the uniform shifting strategy (which allocates approximately equal number of shifts to all histogram bins over whole depth range) achieves a constant expected denominator sequence.

Proof of Result 1

In this section, we will derive an upper bound on ℓ_0 depth error which corresponds to the probability that the depth estimate using the generalized Coates's estimator derived in Supplementary Note 1 is different from the true depth bin.

An upper bound on ℓ_0 error: To ensure that the depth estimate $\widehat{\tau}$ is correct, the bin corresponding to the true depth should have the highest counts after the generalized Coates's correction (Eq. (4)) is applied. Therefore, for a given true depth τ, we want to minimize the probability of depth error:

$$\begin{aligned}\mathbb{P}(\widehat{\tau} \neq \tau) &= \mathbb{P}\left(\bigcup_{i\neq\tau}(\widehat{q}_i > \widehat{q}_\tau)\right) \\ &\leq \sum_{i\neq\tau}\mathbb{P}\left(\widehat{q}_i > \widehat{q}_\tau\right) \\ &= \sum_{i\neq\tau}\mathbb{P}\left(\widehat{q}_i - \widehat{q}_\tau > 0\right)\end{aligned}$$

where the first inequality follows from the union bound. Note that $\widehat{q}_i - \widehat{q}_\tau$ has a mean $q_i - q_\tau$ and variance $\sigma_i^2 + \sigma_\tau^2$ (assuming uncorrelated). The variance is given by [Suppl. Ref. 3] $\sigma_i^2 = \frac{q_i(1-q_i)}{\mathbb{E}[D_i]}$.

For large L, by the central limit theorem, we have:

$$\widehat{q}_i - \widehat{q}_\tau \sim \mathcal{N}(q_i - q_\tau, \sigma_i^2 + \sigma_\tau^2).$$

Using the Chernoff bound for Gaussian random variables, we get:

$$\begin{aligned}\mathbb{P}(\widehat{q}_i > \widehat{q}_\tau) &\leq \exp\left(-\frac{(q_i - q_\tau)^2}{2(\sigma_i^2 + \sigma_\tau^2)}\right) \\ &= \frac{1}{2}\exp\left(-\frac{1}{2}\frac{(q_i - q_\tau)^2}{\frac{q_i(1-q_i)}{\mathbb{E}[D_i]} + \frac{q_\tau(1-q_\tau)}{\mathbb{E}[D_\tau]}}\right)\end{aligned}$$

Assuming a uniform prior on τ over the entire depth range, we get the following upper bound on the average probability of error:

$$\begin{aligned}\frac{1}{B}\sum_{\tau=1}^{B}\mathbb{P}(\widehat{\tau} \neq \tau) &\leq \frac{1}{B}\sum_{\tau=1}^{B}\sum_{i\neq\tau}\frac{1}{2}\exp\left(-\frac{1}{2}\frac{(q_i - q_\tau)^2}{\frac{q_i(1-q_i)}{\mathbb{E}[D_i]} + \frac{q_\tau(1-q_\tau)}{\mathbb{E}[D_\tau]}}\right) \\ &\approx \frac{1}{B}\sum_{i,\tau=1}^{B}\frac{1}{2}\exp\left(-\frac{1}{2}\frac{(q_i - q_\tau)^2}{\frac{q_i(1-q_i)}{\mathbb{E}[D_i]} + \frac{q_\tau(1-q_\tau)}{\mathbb{E}[D_\tau]}}\right)\end{aligned} \tag{S6}$$

We can minimize the probability of error indirectly by minimizing this upper bound. The upper bound involves exponential quantities which will be dominated by the least negative exponent, which in turn is dominated by the index i with the largest value of $1/\mathbb{E}[D_i]$. Therefore, the denominator sequence that minimizes this upper bound must maximize $\min_i \mathbb{E}[D_i]$. Given that the total expected denominator is constant under a fixed number of cycles (see proof below), this is equivalent to making the denominator sequence uniform.

The above analysis assumes that photon detections across cycles are independent, and therefore holds for all acquisition schemes with a gating mechanism that is fixed in advance, including synchronous and asynchronous acquisition (deterministic). Note that it does not hold for photon-driven shifting, where a photon detection in one cycle can affect that in another.

Proof of constant total expected denominator: Let $\Xi \equiv \sum_{i=1}^{B} \mathbb{E}[D_i]$ be the total expected denominator. Assuming a low SBR scenario where the background flux dominates the signal flux, we take $r_i \approx \Phi_{\text{bkg}}$ for all i, i.e., an almost uniform incident waveform. To calculate the total denominator, we sum up the contributions of each cycle. If the SPAD active time is m, each cycle ends with either a photon detection in one of the m time bins, or with no photon detections. A cycle with a photon detection in the i^{th} bin ($1 \leq i \leq m$) contributes i units to the total denominator, since each bin before and including the detection bin was active. By the same argument, a cycle with no photon detections contributes m units. Therefore, the total expected denominator is given by:

$$\begin{aligned}
\Xi &= \sum_{l=1}^{L} \left[\sum_{i=1}^{m} \mathbb{P}(O_{l,s_l \oplus i} = 1) \cdot i + \mathbb{P}(O_{l,s_l \oplus i} = 0 \text{ for } 1 \leq i \leq m) \cdot m \right] \\
&= \sum_{l=1}^{L} \left[\sum_{i=1}^{m} p_{l,s_l \oplus i} \cdot i + \left(1 - \sum_{j=1}^{m} p_{l,s_l \oplus j} \right) \cdot m \right] \\
&\overset{(a)}{=} \sum_{l=1}^{L} \left[\sum_{i=1}^{m} (1 - e^{-\Phi_{\text{bkg}}}) e^{-(i-1)\Phi_{\text{bkg}}} \cdot i + e^{-m\Phi_{\text{bkg}}} \cdot m \right] \\
&= \frac{L(1 - e^{-m\Phi_{\text{bkg}}})}{1 - e^{-\Phi_{\text{bkg}}}} ,
\end{aligned} \tag{S7}$$

where $\oplus$ denotes addition modulo-B, and (a) follows from the fact that for a uniform waveform, $p_{l,s_l \oplus i} = (1 - e^{-\Phi_{\text{bkg}}})(e^{-\Phi_{\text{bkg}}})^{|J_{l,s_l \oplus i}|}$. Moreover, $J_{l,s_l \oplus i} = i - 1$. Therefore, Ξ remains constant for a given L, regardless of the shift sequence used.

Proof of Result 2

For $1 \leq i \leq B$, under uniform shifting:

$$\begin{aligned}
\mathbb{E}[D_i] &\overset{(a)}{=} \sum_{l=1}^{L} \mathbb{E}[D_{l,i}] \\
&= \sum_{l=1}^{L} \prod_{j \in J_{l,i}} (1 - q_j) \\
&\overset{(b)}{\approx} \sum_{l=1}^{L} \prod_{j \in J_{l,i}} e^{-\Phi_{\text{bkg}}} \\
&= \sum_{l=1}^{L} e^{-\Phi_{\text{bkg}} |J_{l,i}|} \\
&= \sum_{l=1}^{L} e^{-\Phi_{\text{bkg}} \{(i - s_l) \bmod B\}} \\
&\overset{(c)}{=} \frac{L}{B} \left(\sum_{i=0}^{B-1} e^{-i\Phi_{\text{bkg}}} \right)
\end{aligned}$$

where (a) follows from Def. 1; (b) relies on the assumption that the pileup is mainly due to ambient light Φ_{bkg} and not the source light Φ_{sig}; and (c) assumes that L is a multiple of B. The last step follows because under uniform shifting, the set of shifts $\{s_l\}_{l=1}^{L}$ spans the discrete space $\{0, 1, ..., B-1\}$ uniformly. Moreover, the shifts remain uniform when transformed by a constant modulo addition. Therefore, $\mathbb{E}[D_i]$ is the sum of identical terms, and hence same for all i.

If L is not a multiple of B, the shift sequence cannot achieve every possible shift from 0 to $B-1$ an equal number of times and some bins might use more shifts than others, especially when $L < B$.However, the expected denominator sequence is still approximately uniform, with the approximation becoming more accurate as L increases. This can be intuitively seen from the gradual elimination of the "saw teeth" in Fig. 3 as the number of shifts increases.

Supplementary Note 3. Achieving Uniform Shifts through Laser and SPAD Cycle Mismatch

The implementation shown in Fig. 3 relies on the SPAD cycle period being different from the laser cycle. By default, the duration for which the SPAD gate is kept on (called the *active time window*, shown as white boxes in Fig. 3), is set equal to the laser cycle period $B\Delta$. After each active time window, we force the SPAD gate to be turned off (gray boxes in Fig. 3) for a duration equal to t_{d} (irrespective of whether a photon was detected). As a result, the SPAD cycle length is equal to $B\Delta + t_{\mathrm{d}}$. This ensures that the dead time from one SPAD cycle does not extend into the next when a photon is detected close to the end of the current cycle. The total number of SPAD cycles over a fixed acquisition time, T, is therefore limited to $L = \lfloor {}^{T}/_{(B\Delta+t_{\mathrm{d}})} \rfloor$.

Since the length of the laser cycle is $B\Delta$ and the SPAD cycle is $B\Delta + t_{\mathrm{d}}$, we automatically achieve the shift sequence $(0, t_{\mathrm{d}}, 2t_{\mathrm{d}}, \ldots)$. Moreover, we can achieve any arbitrary shift sequence in practice by introducing additional artificial delays $(\varepsilon_l)_{l=1}^{L}$ (using, say, a programmable delayer) at the end of each SPAD cycle that extend the inactive times beyond t_{d}. In general, the additional delay $\sum_l \varepsilon_l$ will require either increasing T or decreasing L, both of which are undesirable. In the next section, we show that it is possible to choose ε_l's such that the total extra delay $\sum_l \varepsilon_l \leq 1.5B\Delta$, by making the SPAD cycle length co-prime with the laser. Therefore, uniform shifting can be implemented in practice at a negligible additional cost in terms of total acquisition time.

Quantifying acquisition time cost of uniform shifting

In this section, we bound the additional delay incurred due to uniform shifting. We consider the general case where the SPAD active time window can be different from the laser cycle period. We also present a practical method for implementing uniform shifting that relies on artificially introducing a mismatch in the laser and SPAD repetition frequencies.

Let t_{d} be the dead time of the SPAD, T be the total acquisition time, B be the number of bins in the histogram and Δ be the size of each histogram bin. Let $N = \lfloor {}^{T}/_{\Delta} \rfloor$ be the total acquisition time and $n_d = \lfloor {}^{t_{\mathrm{d}}}/_{\Delta} \rfloor$ be the dead time in units of the histogram bin size. Suppose the SPAD is kept active for m bins in each cycle. Our goal is to find a small positive shift ϵ such that $m + n_d + \epsilon$ becomes asynchronous to B, in the sense that uniform shifts modulo-B are achieved in N bins, and in equal proportions for all shifts.

Let $L = \lfloor {}^{N}/_{m+n_d} \rfloor$ be the number of whole cycles that would be obtained if no shifts were used. For simplicity, assume B is a multiple of L. We divide the range B into L equally spaced intervals and design a shift sequence $(\epsilon_i)_{i=1}^{L}$ such that each of the L cycles is aligned with the start of exactly one of these intervals.

When $\epsilon = 0$, the dead time provides a shift amount equal to ${}^{(m+n_d)L}/_{B} \bmod L$, in units of the interval size ${}^{B}/_{L}$. By rounding it up to the nearest integer, we get the effective shift $s = {}^{(m+n_d+\epsilon)L}/_{B} = \lceil {}^{(m+n_d)L}/_{B} \rceil$. We consider two cases:

Case 1: Suppose s and L are co-prime. In this case, the i^{th} cycle has a unique shift $(i-1)s \bmod L$ for $1 \leq i \leq L$.

Case 2: Suppose s and L are not co-prime and their greatest common divisor is g. We can express $s = g \cdot k$ and $L = g \cdot l$ for some co-prime integers k and l. At the end of every l SPAD cycles, each of the l shifts $\{g, 2g, 3g, \ldots, ..., l \cdot g\}$ is attained once. Since we have g of these "groups" of l SPAD cycles, we can add an offset of i to the shift amounts of the i^{th} group for $i = 0, 1, ..., g-1$. This will ensure a unique shift for each of the L cycles.

Additional delay due to shifting: The additional delay due to shifting is $L\epsilon$ in the Case 1 and $L\epsilon + (g-1){}^{B}/_{L}$ in Case 2. Since ${}^{L\epsilon}/_{B} < 1$ and $g \leq L/2$, the total additional delay incurred due to shifting time is negligible and is at most $1.5B$ bins, i.e., at most one-and-a-half extra laser cycle periods.

Achieving uniform shift sequence through frequency mismatch: Note that in both these cases, uniform shifting can be achieved conveniently by operating the SPAD cycle frequency asynchronous to the laser frequency: ${}^{L}/_{B\epsilon(L+s)}$ in the Case 1 and ${}^{L}/_{B\epsilon(L+s+{}^{1}/_{l})}$ in Case 2.

Supplementary Note 4. Derivation of m^{opt}

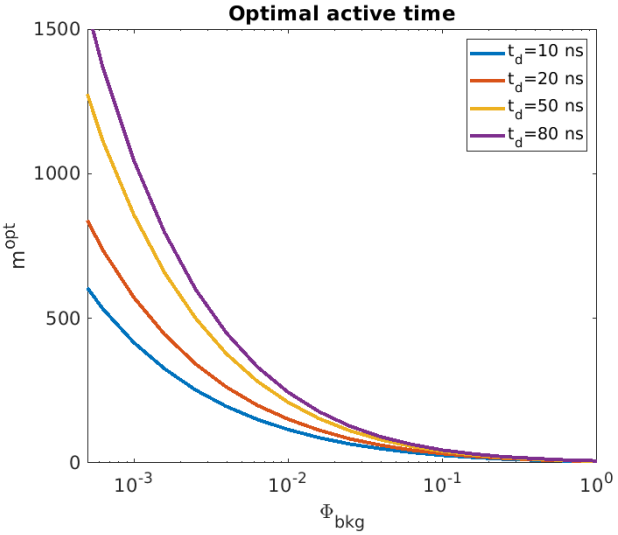

Supplementary Figure 1. **Effect of ambient flux and dead time on m^{opt} .** The optimal active time is a decreasing function of ambient flux Φ_{bkg} and an increasing function of dead time t_{d}.

In this section, we derive the optimal active time for uniform shifting. We assume that the expected denominator sequence is constant (achieved, say, using the method of laser and SPAD cycle mismatch from the previous section). Recall that we use an upper bound on the ℓ_0 depth error as a surrogate for RMSE. From the proof of Result 1 we know that this is equivalent to using the smallest denominator sequence value as a surrogate for RMSE. However, since the denominator sequence is constant on average, we can use the total expected denominator Ξ from Eq. (S7) as a surrogate for depth accuracy—a lower total expected denominator would correspond to higher depth errors and vice versa. We now derive the value of m that maximizes Ξ.

Since the length of each SPAD cycle is $m\Delta + t_{\text{d}}$, the number of cycles in a fixed acquisition time T is given by

$$L = \frac{T}{m\Delta + t_{\text{d}}}. \tag{S8}$$

From Eqs. (S7) and (S8), we get:

$$\Xi(m) = \frac{T}{m\Delta + t_{\text{d}}} \frac{1 - e^{-m\Phi_{\text{bkg}}}}{1 - e^{-\Phi_{\text{bkg}}}}.$$

where we have explicitly included the dependence of Ξ on the active time m. Intuitively, the first term represents the average number of SPAD measurement cycles that can fit in time T and the second term is the expected denominator value in each cycle (assumed to be uniform over all histogram bins in each cycle). The optimal active time m^{opt} is the one that maximizes the total denominator. Solving for $d\Xi/dm = 0$, yields [Suppl. Ref. 5]:

$$m^{\text{opt}} = -\frac{1}{\Phi_{\text{bkg}}} \text{LambertW}(-e^{-t_{\text{d}}\Phi_{\text{bkg}}/\Delta - 1}) - \frac{t_{\text{d}}}{\Delta} - \frac{1}{\Phi_{\text{bkg}}}. \tag{S9}$$

Interpreting m^{opt}: Suppl. Fig. 1 shows the behavior of m^{opt} for different values of dead time t_{d} and varying ambient flux Φ_{bkg}. Observe that the optimal active time decreases with increasing Φ_{bkg}. Under high ambient photon flux, the average time until the first photon detection after the SPAD active time window begins is small. As a result, keeping the SPAD active for a longer duration is inefficient because it unnecessarily increases the length of the measurement cycle, and reduces the number of measurements that can be acquired over a fixed acquisition time. Conversely, at lower flux levels, the optimal active time increases to match the increased average inter-photon arrival time.

For a fixed background flux level, m^{opt} is higher for longer dead times. When the dead time is long, increasing the active time duration increases the probability of detecting a photon while not increasing the measurement cycle length by much.

A uniform shift sequence with SPAD active time equal to m^{opt} can be implemented using the method described in Supplementary Note 3.

Supplementary Note 5. Photon-Driven Shifting: Histogramming and Error Analysis

In this section we provide a proof of Result 3 which states that photon-driven shifting achieves a uniform shift sequence for sufficiently large acquisition times.

This section also proposes an algorithm for computing the generalized Coates's estimator for photon-driven shifting. Unlike uniform shifting where the shift sequence is deterministic and can be pre-computed, the shift sequence in photon-driven shifting is random and depends on the actual photon detection times. We show how the arrival timestamps can be used to compute D_i and N_i needed for the generalized Coates's estimator. It is also possible to estimate the waveform in free-running acquisition using a Markov chain-based model of the arrival times [Suppl. Ref. 6]. While the Markov chain-based estimator is based on solving a non-convex optimization problem, the proposed generalized Coates's estimator has a closed-form expression. Both achieve equivalent performance in terms of depth recovery accuracy.

Proof of Result 3

Let $(S_l)_{l=1}^{L}$ denote the stochastic shift sequence, where $0 \leq S_l \leq B-1$ and $\ominus$ denote subtraction modulo-B with a wrap around when the index falls below 1. This shift sequence forms a Markov chain with state space $[0, B-1]$ and transition density given by [Suppl. Ref. 6]:

$$f_{S_{i+1}|S_i}(s_{i+1}|s_i) = \begin{cases} \left(\frac{1-e^{-\Phi_{\text{bkg}}}}{1-e^{-B\Phi_{\text{bkg}}}}\right) e^{-(s_{i+1}\ominus s_i + B - t_d)\Phi_{\text{bkg}}} & \text{if } s_{i+1} \ominus s_i < t_{\text{d}} \\ \left(\frac{1-e^{-\Phi_{\text{bkg}}}}{1-e^{-B\Phi_{\text{bkg}}}}\right) e^{-(s_{i+1}\ominus s_i - t_d)\Phi_{\text{bkg}}} & \text{otherwise.} \end{cases}$$

The uniform distribution $f(s) = \frac{1}{B}$ is a stationary distribution for this irreducible aperiodic Markov chain, and hence it converges to its stationary distribution [Suppl. Ref. 7] as $L \to \infty$. Let $f^k(s)$ denote the distribution of the k^{th} shift. We have:

$$\lim_{k\to\infty} f^k(s) = f(s)$$

Therefore, as $L \to \infty$, the empirical distribution of $(s_1, s_2, \ldots, s_L) \to f(s)$ and all shifts are achieved with equal probability making the shift sequence uniform. This also leads to a constant expected denominator sequence as shown in the simulations below.

Simulations: Suppl. Fig. 2 shows the expected denominator sequence for different total acquisition times at three different ambient flux levels. There are two main observations here. First, for short acquisition times there is a depth-dependent bias which disappears as T increases. Second, for fixed T, the depth dependent bias is higher for higher flux levels. This is because at high ambient light, the SPAD detects a photon almost deterministically after each dead time window elapses which causes the Markov chain $(S_l)_{l=1}^{L}$ to have a longer mixing time than at lower ambient flux levels.

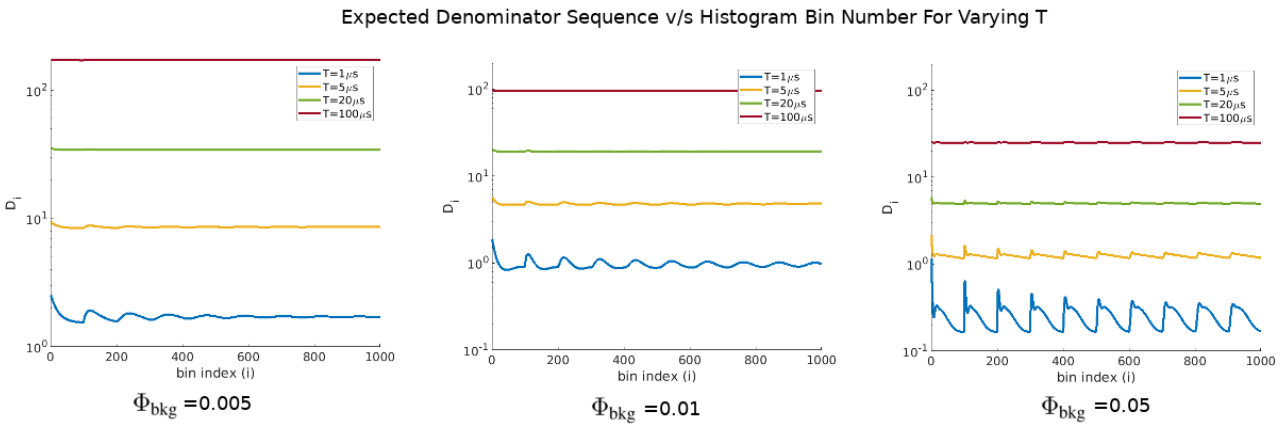

Supplementary Figure 2. **Effect of flux and acquisition time on denominator sequence.** The expected denominator sequence in the photon-driven mode has a position dependent bias which disappears as the total acquisition time increases.

Histogram and denominator sequence computation

In this section, we provide details about the algorithm for computing N_i and D_i from the shift sequence and the sequence of photon arrival times. This leads to a computationally tractable method for computing the generalized Coates's estimate for the flux waveform and hence estimating scene depths.

Let $(u_1, u_2, \ldots, u_L)$ denote the photon arrival times (in terms of bin index) in each SPAD cycle measured with respect to the most recent laser cycle. Note that $1 \leq u_i \leq B$.

The histogram of photon counts is given by:

$$N_i = \sum_{l=1}^{L} \mathbb{1}(u_l = i)$$

To compute the denominator sequence, we loop through the list of photon arrival times u_l. For each photon detection time we increment D_j for every bin index $j \in J_{l,u_l}$.

In the photon-driven shifting mode, the denominator sequence can also be computed in closed form. The shift sequence $(s_l)_{l=1}^{L}$ is determined by the photon arrival times u_i as $s_{l+1} = u_l \oplus n_d$ where n_d is the dead time in units of bins. As before, N_i is given by the histogram of $(u_l)_{l=1}^{L}$. D_i can be computed in closed form in terms of the histogram counts: For each bin index i, there are $\frac{T}{B\Delta}$ depth bins in total which can potentially detect photons. However, a photon detection in any depth bin prohibits the n_d bins that follow it from detecting photons. Therefore, the value of D_i at the i^{th} bin is given by subtracting these bins from the total:

$$D_i = \frac{T}{B\Delta} - \sum_{j=1}^{t_d/\Delta} N_{i \ominus j}.$$

As before for the case of deterministic shifting, the likelihood function of the photon incidence probabilities $(q_i)_{i=1}^{B}$ is given by:

$$\mathcal{L}(q_1, q_2, ..., q_B) = \prod_{i=1}^{B+1} q_i^{N_i} (1 - q_i)^{D_i - N_i}\,.$$

Therefore, the generalized Coates's estimator for photon-driven shifting is given as the following closed-form expression:

$$\begin{aligned} \widehat{\tau} &= \arg\max_i \frac{N_i}{D_i} \\ &= \arg\max_i \frac{N_i}{\frac{T}{B\Delta} - \sum_{j=1}^{t_d/\Delta} N_{i \ominus j}}\,. \end{aligned} \tag{S10}$$

Markov chain model-based estimator [Suppl. Ref. 6]

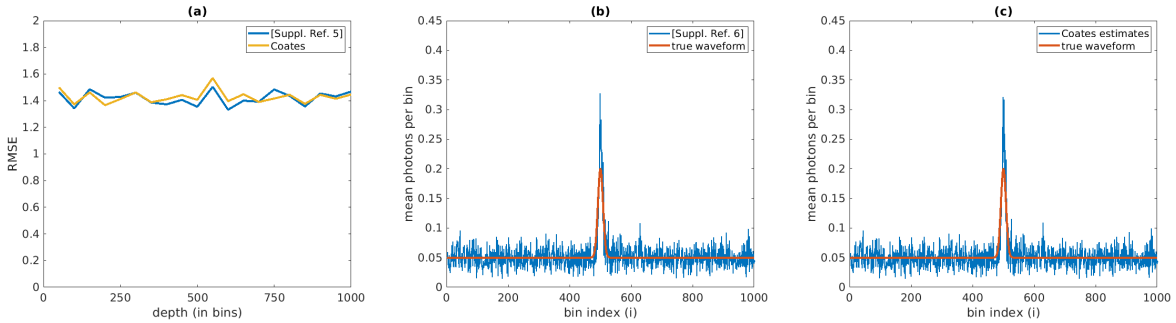

Supplementary Figure 3. **Comparison between [Suppl. Ref. 6] and the proposed generalized Coates's estimator.** (a) shows the RMSE curve for the two methods. Both methods have the same performance across depth. (b) and (c) show that the estimated waveforms using the two methods are similar.

It is also possible to use a Markov chain for modeling the photon arrival times in free-running acquisition. The resulting estimator [Suppl. Ref. 6] is based on solving a non-convex optimization problem, whereas the proposed generalized Coates's

estimator has a closed-form expression, as shown above. Both achieve equivalent performance in terms of depth recovery accuracy, as shown in Suppl. Fig. 3. In the regime of large acquisition time ($T = 25\,\mu$s), and with $\Phi_{\text{bkg}} = 0.05, \Phi_{\text{sig}} = 0.15$, the performance of both the estimators in terms of depth error is almost equivalent.

Performance Gain as a Function of Depth

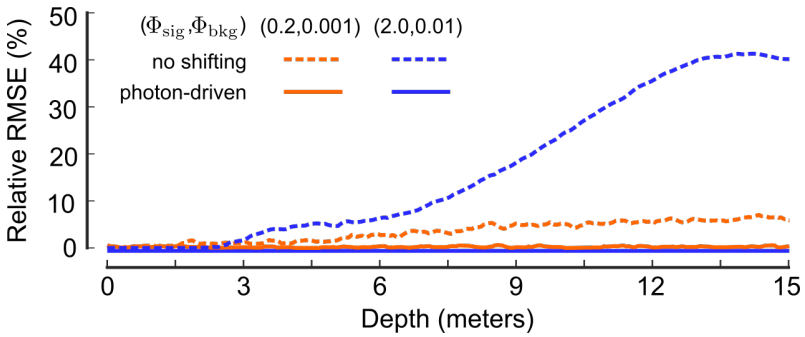

Supplementary Figure 4. Asynchronous acquisition schemes minimize overall depth error, and as a byproduct also provide uniform error with depth. The performance gain is larger at longer distances.

Suppl. Fig. 4 shows error vs. depth for photon-driven and conventional synchronous acquisition for different source and background flux levels. There are two key observations:

1. The performance gain from photon-driven acquisition is greater at longer distances. (Although not shown here, similar conclusion can be drawn for deterministic shifting as well.)

2. RMSE is approximately depth-invariant for photon-driven shifting.

Achieving uniform error over all depths is not our primary goal, but a byproduct of our optimal acquisition schemes. We define optimality in terms of a surrogate of ℓ_2 depth error metric which measures overall depth error. Result 1 shows that the optimal acquisition scheme that minimizes our error metric also has a constant denominator sequence for all depths.

Supplementary Note 6. Asynchronous Shifting: Practitioner's View

In this supplementary note we provide practical guidelines for design of asynchronous acquisition strategies. We provide a comparison between photon-driven shifting and uniform shifting with optimal active time m^{opt} and show that there are certain regimes where photon-driven shifting has slightly worse error performance than uniform shifting. We also analyze the effect of number of SPAD cycles on depth error performance of photon-driven acquistion.

Derivation of Expected Denominator Sequence for Photon-Driven Acquisition

We consider a low SBR scenario with $r_i \approx \Phi_{\text{bkg}}$ for all i. Let $(D_i)_{i=1}^{B}$ denote the denominator sequence obtained over L SPAD cycles and a fixed acquisition time T. Since the length of each SPAD cycle is random, the number of cycles in time T is also random. By the equivalence of the length of the active period of a cycle and its contribution to the total denominator, we have:

$$D_1 + D_2 + ... + D_B = M_1 + M_2 + ... + M_L$$

where M_i is a random variable denoting the length of the i^{th} measurement window. Note that the M_i's are i.i.d. with a discrete exponential distribution:

$$P(M_i = k) = e^{-(k-1)\Phi_{\text{bkg}}}\left(1 - e^{-\Phi_{\text{bkg}}}\right)$$

for $k = 1, 2, 3, \ldots$. Therefore, $\mathbb{E}[M_i] = \frac{1}{1-e^{-\Phi_{\text{bkg}}}}$. Taking expectation on both sides of the above equation, we get:

$$\mathbb{E}[D_1 + D_2 + ... + D_B] = \mathbb{E}[M_1 + M_2 + ... + M_L] \tag{S11}$$

$$\approx \mathbb{E}[L]\mathbb{E}[M_i] = \frac{\mathbb{E}[L]}{1 - e^{-\Phi_{\text{bkg}}}}\,, \tag{S12}$$

where, we have assumed that M_i and L are approximately independent. We also have:

$$(M_1 + t_{\text{d}}) + (M_2 + t_{\text{d}}) + ... + (M_L + t_{\text{d}}) = T\,.$$

Taking the expectation on both sides and using Eq. (S12), we get:

$$\frac{\mathbb{E}[L]}{1 - e^{-\Phi_{\text{bkg}}}} + \mathbb{E}[L]t_{\text{d}} = T$$

which implies that

$$\mathbb{E}[L] = \frac{T}{\frac{1}{1-e^{-\Phi_{\text{bkg}}}} + t_{\text{d}}}\,.$$

Combining with Eq. (S11), we get:

$$\mathbb{E}[D_1 + D_2 + ... + D_B] = \frac{T}{1 + (1 - e^{-\Phi_{\text{bkg}}})t_{\text{d}}}\,.$$

From Result 3, we know that $\mathbb{E}[D_i]$ is constant $\forall\, i$ as $L \to \infty$. Under this assumption, we have for $1 \le i \le B$:

$$\mathbb{E}[D_i]|_{\text{photon-driven}} = \frac{T}{B(1 + (1 - e^{-\Phi_{\text{bkg}}})t_{\text{d}})}. \tag{S13}$$

Comparing expected denominator sequences for uniform shifting and photon-driven shifting: Recall from Supplementary Note 4 the expected denominator sequence for uniform shifting is given by:

$$\mathbb{E}[D_i]|_{\text{uniform}} = \frac{T(1 - e^{-m\Phi_{\text{bkg}}})}{B(m\Delta + t_{\text{d}})(1 - e^{-\Phi_{\text{bkg}}})}. \tag{S14}$$

Observe that $\mathbb{E}[D_i]|_{\text{photon-driven}} > \mathbb{E}[D_i]|_{\text{uniform}}$ for all $m \ge 1$, Φ_{bkg} and t_{d}. To a first order approximation, the expected denominator sequence alone determines the depth estimation performance of the various techniques that we discussed. This suggests that photon-driven shifting should be always better than uniform shifting. However, in some cases, depending on the dead time and incident flux levels, deterministic shifting can outperform photon-driven acquisition. One such scenario is demonstrated in Suppl. Fig. 5 using both simulation and experimental results. At certain values of dead time, photon-driven shifting fails to achieve uniform shifts in the fixed acquistion time.

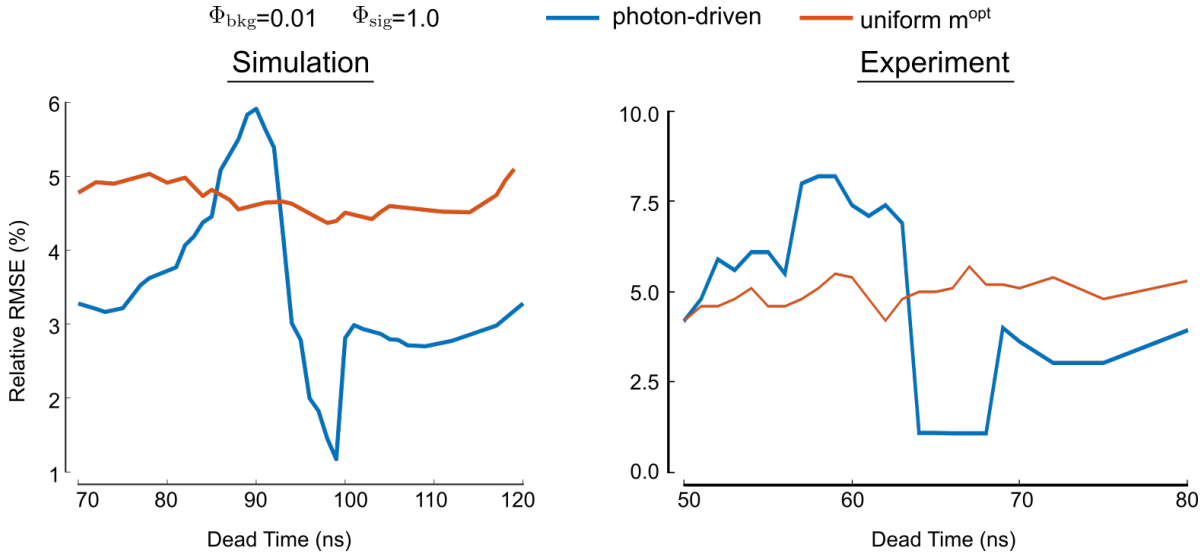

Supplementary Figure 5. In some special scenarios, depending on dead time and incident flux, uniform shifting may outperform photon-driven shifting.

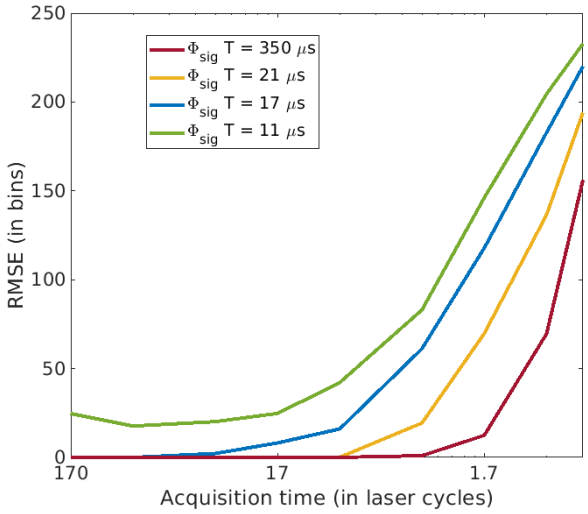

Supplementary Figure 6. **Effect of reducing laser cycles on photon-driven shifting with optimal attenuation.** Each curve has a fixed $\Phi_{\text{sig}}T$ product. As T decreases and Φ_{sig} increases, the RMSE rises. The effect of decreasing T can be overcome to some extent by increasing Φ_{sig}.

Effect of reducing laser cycles on photon-driven shifting Intuitively, one would expect asynchronous acquisition techniques to require a large number of cycles, in order to achieve uniform shifting and remove pileup as motivated in the main text. However, it turns out that our combined method of photon-driven shifting and optimal attenuation can work with very few laser cycles, provided that the signal flux is high. This is true even for relatively high ambient flux levels (we consider $\Phi_{\text{bkg}} = 0.05$). Intuitively, even though uniform shifting cannot occur with laser cycles much less than the number of bins B, the combination of high effective Φ_{sig} and low effective Φ_{bkg} obviates the need for shifting.

In Suppl. Fig. 6, we evaluate depth error performance as number of laser cycles is decreased, while keeping the $\Phi_{\text{sig}}T$ product constant. For all Φ_{sig} levels, RMSE curves have flat valleys, with errors starting to rise beyond a certain T and reaching the maximum possible error eventually. It is interesting to see the transition from the valley to max error for different Φ_{sig} levels. As the Φ_{sig} level is increased, the transition happens at a lower acquisition time, until T becomes less than about 2 laser cycles. Beyond this, the error cannot be reduced no matter how high Φ_{sig} is made. In some sense, this is the limiting point of our method.

Supplementary Note 7. Combining Shifting Strategies with Optimal Flux Attenuation

In this section we derive an expression for optimal flux attenuation for photon-driven acquisition. We also provide some information theoretic arguments for why it is important to combine asynchronous acquisition schemes with flux attenuation to achieve optimal depth error performance.

Optimal Attenuation Fraction for Photon-Driven Shifting

We will use similar techniques as the derivation of Result 1 from Supplementary Note 2 and introduce an additional variable Υ for flux attenuation fraction. In particular the expected denominator sequence in Eq. (S13) becomes:

$$\mathbb{E}[D_i] = \frac{T}{B(1+(1-e^{-\Upsilon\Phi_{\text{bkg}}})t_{\text{d}})}. \tag{S15}$$

From Eq. (S6), we have:

$$\begin{aligned}
\frac{1}{B}\sum_{\tau=1}^{B}\mathbb{P}(\widehat{\tau}\neq\tau) &\leq \frac{1}{B}\sum_{i,\tau=1}^{B}\frac{1}{2}\exp\left(-\frac{1}{2}\frac{(q_\tau-q_i)^2}{\frac{q_i(1-q_i)}{\mathbb{E}[D_i]}+\frac{q_\tau(1-q_\tau)}{\mathbb{E}[D_\tau]}}\right)\\
&\approx \frac{1}{B}\sum_{i,\tau=1}^{B}\frac{1}{2}\exp\left(-\frac{1}{2}\mathbb{E}[D_i](q_\tau-q_i)\right)\\
&= \frac{B}{2}\exp\left(-\frac{1}{2}\mathbb{E}[D_i](q_\tau-q_i)\right)
\end{aligned}$$

where in the second step, we have assumed that the denominator sequence is uniform, $q_i, q_\tau \ll 1$ and $q_\tau - q_i \approx q_\tau + q_i$. Substituting Eq. (S15), the optimal attenuation fraction is given by:

$$\begin{aligned}
\Upsilon^{\text{opt}} &= \arg\max_{\Upsilon}\left[\mathbb{E}[D_i](q_\tau-q_i)\right]\\
&= \arg\max_{\Upsilon}\frac{Te^{-\Upsilon\Phi_{\text{bkg}}}(1-e^{-\Upsilon\Phi_{\text{sig}}})}{B(1+(1-e^{-\Upsilon\Phi_{\text{bkg}}})t_d)}.
\end{aligned}$$

The optimal attenuation fraction Υ^{opt} depends on various system parameters (number of bins, acquisition time and the dead time), which are known. Υ^{opt} also depends on the signal and background flux levels, which can be estimated a priori using a small number of laser cycles. Using the estimated values, the optimal flux attenuation can be determined, which can be used to further reduce the pile-up, and thus improve the depth estimation performance.

Fig. 6 in the main paper shows simulation results of depth RMSE for the conventional synchronous mode and photon-driven shifting over a range of attenuation factors and two different dead times. The locations of the minima agree with the theoretical expression derived above. There are two key observations that can be made from the figure. First, the optimal attenuation fraction with shifting is much higher than that for conventional synchronous acquisition. Second, combining attenuation with photon-driven shifting can provide a large gain in depth error performance, reducing the RMSE to almost zero for some combinations of SBR and dead time.

Information Theoretic Argument for Combining Shifting with Flux Attenuation

We formalize the notion of discriminative power of the histogram data for distinguishing between incident flux waveform levels using a concept from statistics called *Fisher information* [Suppl. Ref. 4, pp. 35]. This concept was also used in [Suppl. Ref. 6]. Fisher information measures of the rate of change of the likelihood function with respect to the unknown parameters r_i. For any attenuation coefficient Υ, we can compute:

$$\mathcal{I}(r_i;\Upsilon) = \mathbb{E}\left[\left(\frac{\partial}{\partial r_i}\log\mathcal{L}(N_i;D_i,r_i,\Upsilon)\right)^2\right]$$

Substituting the expression for likelihood from Eq. (S5) and simplifying yields:

$$\mathcal{I}(r_i;\Upsilon) = \frac{D_i\Upsilon^2}{e^{\Upsilon r_i}-1} \tag{S16}$$

Asynchronous acquisition methods can only change D_i which appears in the numerator of the Fisher information. Attenuation can further increase the Fisher information and lead to better flux utilization than would be possible with shifting alone.

Supplementary Note 8. Details of Simulations and Experiments

Calculating RMSE

We calculate RMSE in a modulo-B sense in our simulations and experiments. For example, this means that the first and the last histogram bins are considered to be only 1 bin apart. Let M be the number of Monte Carlo runs, B be the number of histogram bins, τ_i be the true depth bin index and $\widehat{\tau}_i$ be the estimated depth bin index for the i^{th} Monte Carlo run. We calculate RMSE using the following formula:

$$\text{RMSE} = \left[\frac{1}{M} \sum_{i=1}^{M} \left(\frac{B}{2} - \left(\widehat{\tau}_i - \tau_i + \frac{B}{2} \right) \bmod B \right)^2 \right]^{\frac{1}{2}}.$$

Monte Carlo Simulations

In Figs. 4 and 7, the dead time was $10\,\text{ns}$, and the exposure time was $2.5\,\mu\text{s}$ In Fig. 6, the exposure times were $2.5\,\mu\text{s}$ and $2.9\,\mu\text{s}$ respectively for the two dead times, to acquire an equal number of SPAD cycles in both cases. At low incident flux and longer dead times, an active time duration longer than B may be needed to increase the probability that the SPAD captures a photon.

Experimental Setup

Suppl. Fig. 7 shows various components of our experimental setup. The SPAD is operated by a programmable control unit and photon timestamps are captured by a TCSPC module (not shown).

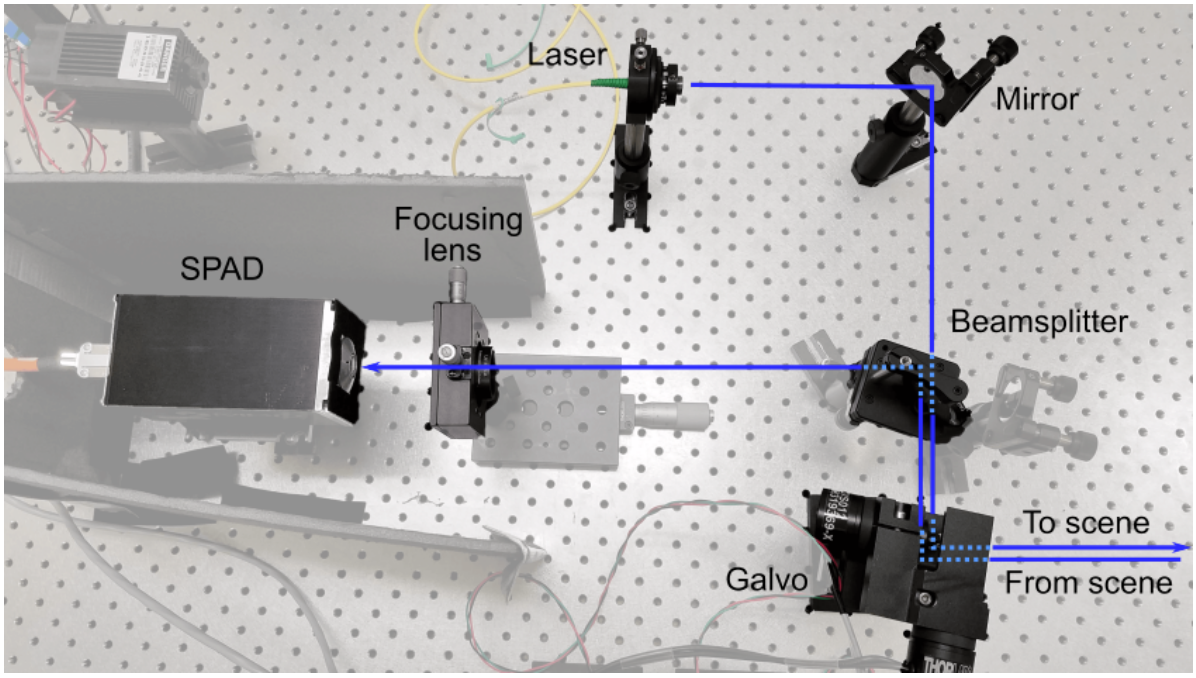

Supplementary Figure 7. **Experimental setup** The setup consists of a collimated pulsed laser and a single-pixel SPAD detector optically collocated using a beamsplitter. A pair of galvo mirrors scans the image plane and the light returning from the scene is focused on the detector using a focusing lens. The detector and focusing lens are enclosed in a light-tight box (not shown here) so that the all the light arriving at the detector pixel must pass through the lens.

Single-Pixel Experiment Results

Suppl. Fig. 8 shows depth RMSE at four different ambient flux levels for a range of SBR values. A combination of photon-driven shifting with optimal attenuation provides the best performance of all techniques. For high ambient flux levels, even at high SBR, the conventional 5% rule-of-thumb and shifting alone fail to provide acceptable depth reconstruction performance.

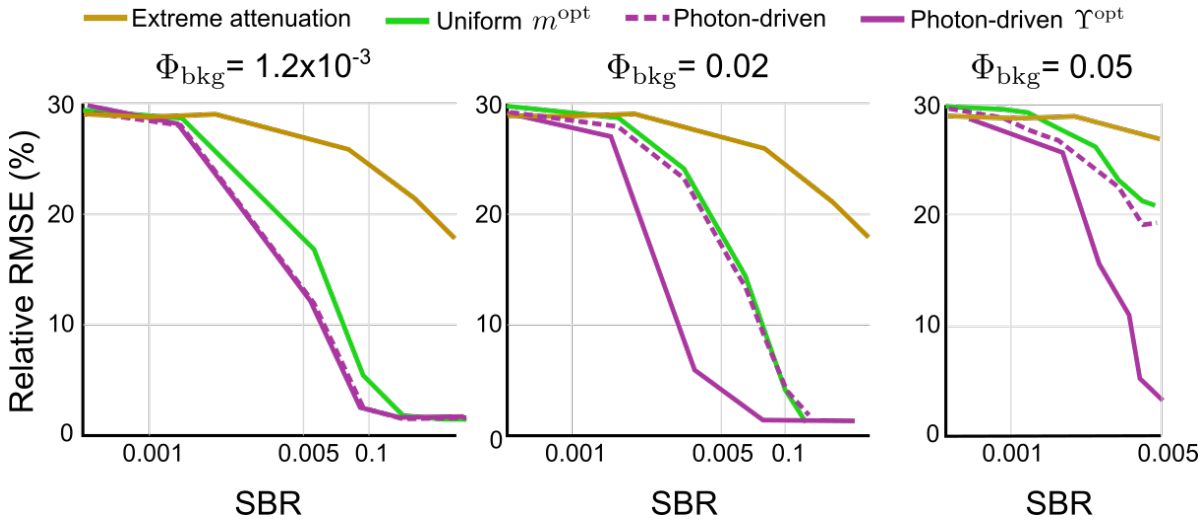

Supplementary Figure 8. **RMSE in single-pixel experiments.** Asynchronous schemes outperform the state-of-the-art method of synchronous acquisition with extreme attenuation at all flux levels. A combination of photon-driven shifting with optimal attenuation provides an order of magnitude lower RMSE at high ambient flux levels.

This suggests that the optimal acquisition strategy for SPAD-based LiDAR must use a combination of both shifting and attenuation.

Additional Result Showing Effect of Number of SPAD Cycles

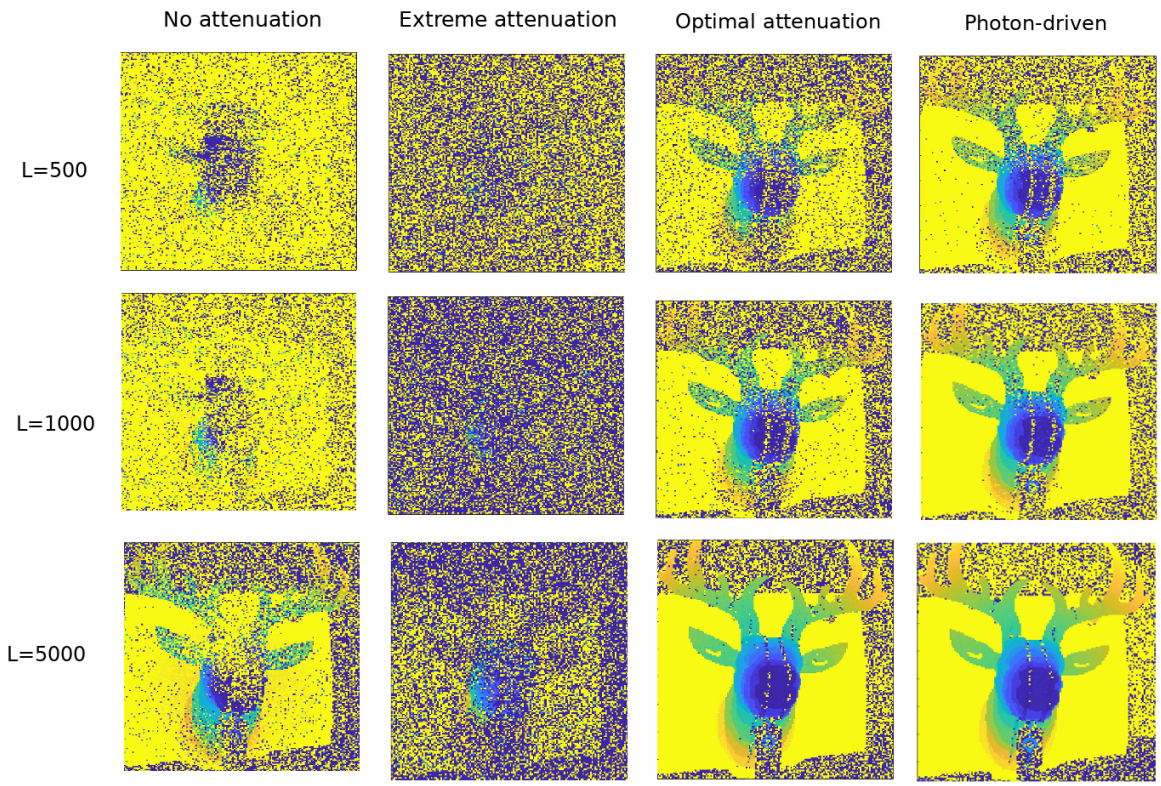

Supplementary Figure 9. **Depth reconstructions for varying number of SPAD cycles for "Reindeer" scene.** Observe that the reconstruction accuracy improves as the number of cycles increases. Photon-driven shifting provides better reconstruction performance than all synchronous acquisition schemes. The source flux $\Phi_{\text{sig}} = 0.5$ and ambient flux $\Phi_{\text{bkg}} = 0.01$ for this experiment which corresponds to an SBR of 50.

Supplementary References

[1] K. Murphy, Machine Learning: A Probabilistic Perspective. Cambridge, MA: MIT Press, 2012, pp. 35.

[2] C. Daskalakis, G. Kamath, and C. Tzamos. On the structure, covering, and learning of poisson multinomial distributions. arXiv preprint, 2015.

[3] A. Gupta, A. Ingle, A. Velten, and M. Gupta. Photon flooded single-photon 3d cameras. arXiv preprint arXiv:1903.08347, 2019.

[4] S. Kay, Fundamentals of Statistical Signal Processing: Estimation Theory. Upper Saddle River, NJ: Prentice Hall, 1993, pp. 173–174.

[5] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the LambertW function. Advances in Computational Mathematics, 5(1):329359, Dec 1996.

[6] J. Rapp, Y. Ma, R. Dawson, and V. K. Goyal. Dead time compensation for high-flux ranging. arXiv preprint arXiv:1810.11145, 2018.

[7] G. Grimmett, D. Stirzaker, Probability and Random Processes. Oxford, UK: Oxford University Press, 2001, pp. 227.