

Hardware-aware Coding Function Design for Compressive Single-photon 3D Cameras

David Parra, Felipe Gutierrez-Barragan, Trevor Seets, and Andreas Velten

Abstract—Single-photon cameras are becoming increasingly popular in time-of-flight 3D imaging because they can time-tag individual photons with extreme resolution. However, their performance is susceptible to hardware limitations, such as system bandwidth, maximum laser power, sensor data rates, and in-sensor memory and compute resources. Compressive histograms were recently introduced as a solution to the challenge of data rates through an online in-sensor compression of photon timestamp data. Although compressive histograms work within limited in-sensor memory and computational resources, they underperform when subjected to real-world illumination hardware constraints. To address this, we present a constrained optimization approach for designing practical coding functions for compressive single-photon 3D imaging. Using gradient descent, we jointly optimize an illumination and coding matrix (i.e., the coding functions) that adheres to hardware constraints. We show through extensive simulations that our coding functions consistently outperform traditional coding designs under both bandwidth and peak power constraints. This advantage is particularly pronounced in systems constrained by peak power. Finally, we show that our approach adapts to arbitrary parameterized impulse responses by evaluating it on a real-world system with a non-ideal impulse response function.

Index Terms—single-photon avalanche diodes, 3d imaging, time-of-flight, optimization, hardware-aware, coding

1 INTRODUCTION

SINGLE-PHOTON cameras (SPCs) based on single-photon avalanche diodes (SPADs) are capable of measuring the time-of-arrival of individual photons [1, 2] with extreme time resolution, down to picoseconds [3]. This makes SPAD-based SPCs a popular sensor for time-resolved imaging where a scene is illuminated with a fast pulsed periodic light signal, and the SPAD measures the reflected light waveform. These imaging modalities include time-of-flight (ToF) 3D imaging [4–7], non-line-of-sight (NLoS) imaging [8, 9], and fluorescence lifetime (FLIM) imaging [10]. In particular, single-photon 3D ToF cameras are used in applications such as autonomous vehicles [11] and mobile devices [12], each with different constraints. As single-photon 3D cameras are adopted in more domains and devices, maximizing their accuracy while adhering to hardware constraints becomes critical.

The illumination and sensor hardware limitations of a single-photon 3D camera play an important role in determining their cost and practicality. For instance, although ultra-short pulsed lasers may be a desirable illumination source, they require sophisticated electronics with high bandwidth [13], which can be impractical and expensive [14, 15]. Furthermore, thermal, electrical, and material constraints on laser diodes limit the laser signal peak power [16, 17]. Finally, a precise 3D ToF camera must account for irregular pulse waveforms that may deviate from an assumed model (e.g., Gaussian pulses). Therefore, bandwidth, peak power, and illumination waveforms are key illumination hardware constraints that should be con-

sidered in designing a practical single-photon 3D camera.

On the sensor side, data rates and the amount of in-sensor compute and memory are constrained. SPAD pixels in single-photon 3D cameras can generate billions of photon timing events per second. Although in-sensor timestamp histogramming can make this large volume more manageable [12], the in-sensor memory required increases proportionally to the temporal and spatial resolution, making it unfeasible to store high-resolution histograms [19]. Moreover, regardless of the limited in-sensor memory, transferring high-resolution histograms from a megapixel SPAD array [20, 21] will lead to data rates of hundreds of Gbit/s [22] exceeding the limits of the commonly used MIPI CSI-2 implementations (~ 10 Gbit/s [23]). Hence, to meet data rate constraints, single-photon 3D cameras must perform in-sensor compression of the photon data stream, and this compression algorithm must also function within the sensor’s limited computational and memory resources.

The compressive histograms framework provides a solution for online in-sensor compression of photon timestamp data [18, 19, 22]. Compressive histograms encode each photon’s timestamp information through a coded projection with multiple “coding functions” [19]. Recent on-chip implementations demonstrate that compressive histograms can adhere to the highly constrained in-sensor compute and memory resources [22, 24]. However, current solutions do not consider illumination hardware constraints (Fig. 1(b)) in the coding function design, which can lead to sub-optimal performance in real-world systems, as shown in Fig. 1(c). This paper investigates the coding function design for compressive single-photon 3D cameras while accounting for hardware constraints.

To address this problem, we formulate it as a constrained optimization that jointly optimizes the illumination and the corresponding coding matrix, i.e., the coding functions

- D. Parra, T. Seets, and A. Velten are with the University of Wisconsin-Madison, Madison, WI, 53706. (email deparra@wisc.edu; velten@wisc.edu).
- F. Gutierrez-Barragan is an independent researcher.

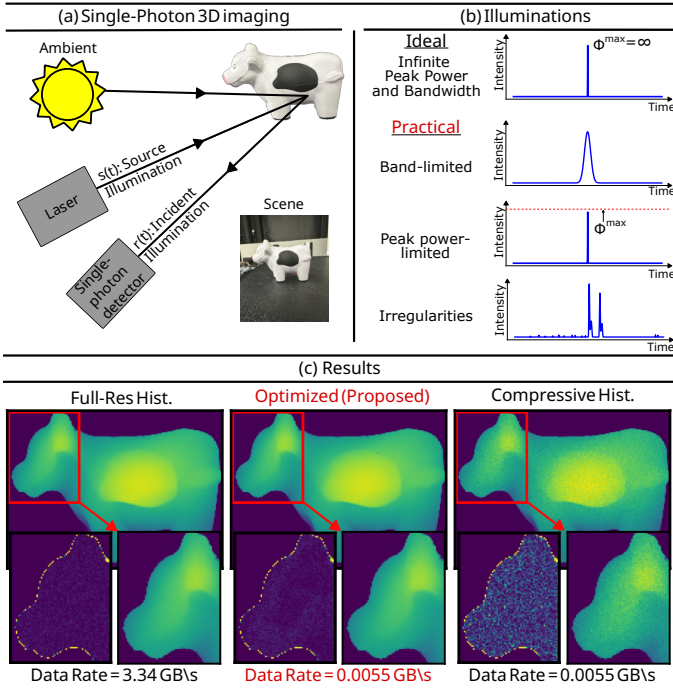


Fig. 1. **Hardware-aware coding function design for compressive single-photon 3D cameras.** (a) Overview of a single-photon 3D camera based on the time-of-flight principle, which illuminates a scene with a pulsed laser and detects the reflected signal at a single-photon detector. (b) The ideal source illumination will have infinite bandwidth and power, but practical hardware considerations demonstrate illumination bandwidth limitations, peak power limitations, and non-ideal impulse response functions. (c) Simulated depth map and depth error results using full-res. histogramming, compressive histogramming [18], and our optimized coding functions with finite bandwidth and finite peak power.

Given a specific hardware setup, our method identifies a set of high-performing coding functions that adhere to these constraints with minimal loss in depth accuracy. These coding functions can adapt to arbitrary parametrized impulse response functions (IRFs), including those affected by irregularities (Fig. 1(b)). Through simulations, we demonstrate that our method achieves superior depth accuracy compared to prior compressive histogram solutions [18, 19, 25] in hardware-constrained systems. Notably, peak power constraints lead to an optimized illumination with a symmetric, multi-peaked waveform that shows significant improvement over traditional methods (Fig. 1(c)). Finally, we show that our optimized coding function can adapt to real-world systems with a non-ideal impulse response function and evaluate its performance.

2 RELATED WORKS

Compression in single-photon imaging: One way to reduce data rates is compressive histogramming [18, 19, 22, 26], which is an online, in-sensor compression of histogram data. A simple approach to compressive histograms is coarse histograms [14, 27]. Although several studies have made an effort to improve coarse histograms [7, 25, 28], it has been shown to perform sub-optimally compared to approaches based on Fourier [18, 19, 29] and Gray [19] codes. Recent work has extended compressive histogramming by incorporating spatial and temporal information [22]. Although

current compressive histogram methods can account for bandwidth during the depth estimation step [19], they fail to use practical illumination constraints in the code design. We demonstrate that these codes become suboptimal when applied to real-world systems with band-limited and peak power-limited illuminations.

Additional strategies to reduce histogram data have been proposed, including equi-depth histograms [30, 31] and foveated SPADs [32]. Event-based detection has been explored for both passive [33, 34] and active single-photon imaging [35], which, as a byproduct, reduces data volume by capturing a sparse set of photon data. Furthermore, low-resolution “proximity” SPAD sensors, which have low data rates, combined with algorithmic upsampling techniques, are a promising low-cost 3D imaging solution [36–40], but at the cost of a computationally expensive 3D reconstruction step. In this work, we limit our analysis to incorporating hardware constraints within the compressive histogram framework.

Code optimization in active imaging: Code design for active 3D imaging has been explored in the areas of indirect ToF (i-ToF) [41–46] and structured light (SL) [47–50]. Previous work has focused on optimizing codes for real-world systems in both i-ToF [42] and SL [50]. However, to the best of our knowledge, hardware-aware code optimization for single-photon imaging has not yet been explored.

3 SINGLE-PHOTON TIME OF FLIGHT IMAGING

Single-photon ToF imaging involves a periodic pulsed laser (of period τ) that illuminates the scene, and a SPAD detector that captures the returned photons. The intensity of the laser at time t can be modeled as a continuous function $s(t) \geq 0$ and the radiance *incident* on the sensor as $r(t) \geq 0$. In practice, the illumination emitted by our system is constrained by its impulse response, defined by the function $h(t)$.

3.1 Single-photon Histogram Formation

A single-photon 3D imaging system will convert the incident waveform $r(t)$ into discrete time bins. We denote the number of time bins as N , the left edge of the i th time bin as b_i , and the size of a single time bin as Δ . The number of photons incident on the SPAD follows a Poisson distribution, and the mean is given by,

$$r_i = \Phi_i^{\text{sig}} \int_{b_i}^{b_i+\Delta} s(t - 2d/c) dt + \Phi_i^{\text{bkg}}, \quad (1)$$

where d is scene distance and c is the speed of light. Φ_i^{sig} is the mean number of signal photons for the i th bin, and Φ_i^{bkg} is the background and dark count photon count for the i th bin.

To estimate the incident waveform, a SPAD pixel will form a histogram of photon arrival times. The SPAD pixel starts acquiring photons immediately once triggered by a laser pulse. In each laser cycle, after the first incident photon is detected, the SPAD pixel enters a dead time during which it cannot detect any more photons. The probability q_i that at least one photon is incident during bin i is given by,

$$q_i = 1 - e^{-r_i}. \quad (2)$$

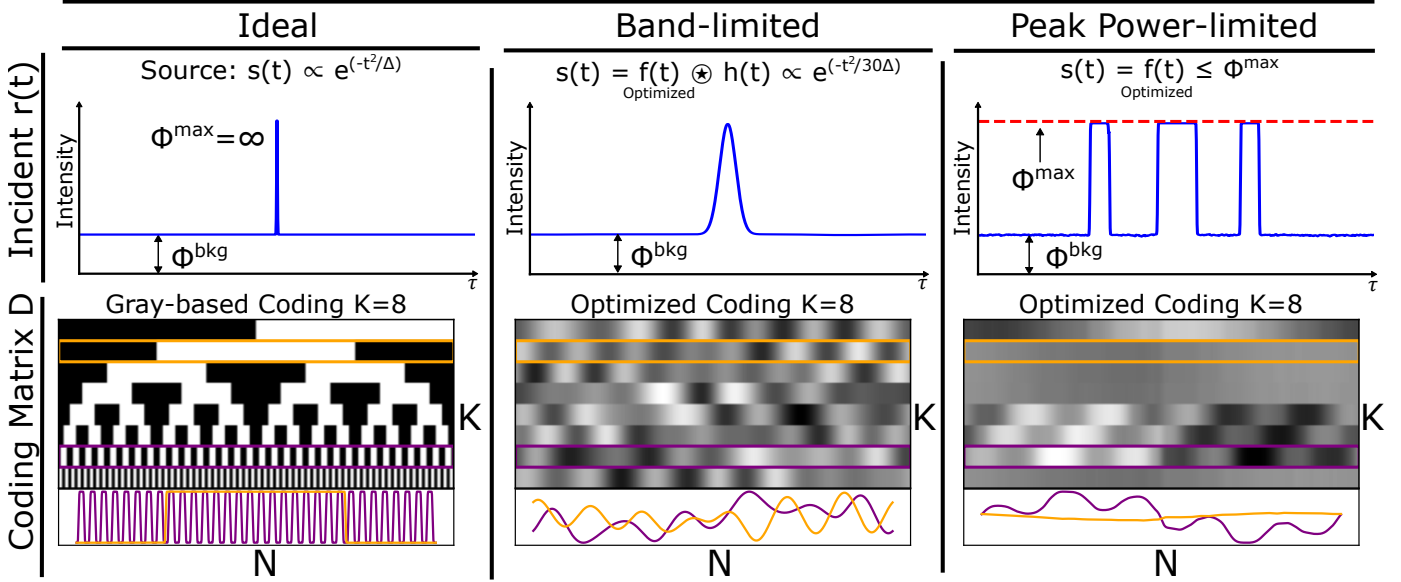


Fig. 2. **Overview of hardware-aware optimized coding functions.** Conventional compressive histograms (left) perform optimally with the ideal pulsed illumination. When hardware constraints, such as bandwidth (middle) and peak power (right), are introduced, compressive histograms become suboptimal, and the coding matrix and source illumination must be adjusted to prevent performance degradation. To address this, we propose a method to optimize both the illumination and coding matrix (or coding functions) based on bandwidth and/or peak power constraints. By adapting the coding functions, we demonstrate improved performance compared to compressive histograms.

A histogram of photon arrival times is constructed over L laser cycles. Let M_i be the number of photons detected in bin i . Thus, M_i follows a binomial distribution parameterized by the number of laser cycles L and probability q_i . Here, we assume that the SPAD is operated in asynchronous mode [51] or in the linear regime [5] such that pile-up distortion is minimized [5]. Thus, the detected histogram is an unbiased estimate of the incident waveform.

3.2 Compressive Histograms

As SPAD ToF cameras increase spatial and temporal resolution, data rates can become unmanageable, making it crucial to find strategies for reducing this. Compressive histograms address this by defining a $K \times N$ coding matrix D which will serve as a linear projection on a full-resolution histogram. Thus, the coded values are expressed as a matrix multiplication between the measured histogram M and coding matrix D , expressed as:

$$B_k = \sum_{i=0}^{N-1} D_{k,i} M_i. \quad (3)$$

The vector $B = (B_1, \dots, B_K)$ represents the coded values for a single pixel. To decode depths, we compute the zero-mean normalized cross-correlations (ZNCC), introduced in SL [50], between our coded values and coding matrix. The decoded depth with respect to time is given as:

$$\hat{t}_d \propto \arg \max_i \frac{D'_{:,i} - \text{mean}(D'_{:,i})}{\|D'_{:,i} - \text{mean}(D'_{:,i})\|} \cdot \frac{(B - \text{mean}(B))}{\|B - \text{mean}(B)\|}, \quad (4)$$

where $D'_{:,i}$ is the i th column of D correlated with $h(t)$. ZNCC decoding is a template-matching method that identifies the phase-shifted illumination that best matches the coded values B measured by our SPAD detector. To decode

depths for full-res. histograms we use matched filtering [52]. Since (3) is a linear operation, compressive histograms can employ an on-the-fly encoding of individual photon timestamps in-sensor, allowing the SPAD pixel to read out the vector B instead of the entire histogram.

The compressive histogram framework is the most effective with an ideal pulsed illumination, one that matches the SPAD sensor's temporal resolution and delivers ultra-high peak power. Although this framework has been evaluated using realistic instrument impulse responses, we demonstrate that under realistic hardware conditions, existing compressive histogram codes can perform sub-optimally. As shown in Fig. 3, when using the coding matrix from [19] with a band-limited illumination, several rows of the matrix D' are zeroed out, causing the corresponding entries in the coded vector B to be zero and contributing nothing to the ZNCC reconstruction. This highlights the need for hardware-aware code design for low-cost, compressive single-photon 3D cameras, which is the focus of this paper.

4 HARDWARE-AWARE CODING FUNCTION DESIGN

In this section, we outline the hardware constraints we considered and formulate the constrained optimization for practical single-photon code design.

4.1 Hardware Limitations

When developing a low-cost, constrained single-photon 3D camera, it is essential to address several key illumination hardware limitations that will directly impact system performance. While high-end systems may circumvent hardware limitations, low-cost implementations must consider the following constraints:

System Bandwidth: The bandwidth of our system is characterized by its impulse/frequency response. Ideally, a ToF




Fig. 3. **Gray codes subjected to strict bandwidth constraints.** Gray-based [19] compressive histograms cannot efficiently encode the illumination’s waveform when severely band-limited.

camera’s frequency response should be high enough so that the laser’s pulse width matches the timing resolution of our SPAD sensor [53], assuming sufficiently high photon counts [54]. However, this is not always practical in low-cost, constrained systems. Most LiDAR systems use semiconductor pulsed lasers [55], which are pulsed by modulating a power supply current. Higher modulation speeds become increasingly ineffective at larger currents [56, 57], ultimately restricting the bandwidth of the system (about 1-10 ns pulse widths for LiDAR [58]). Consequently, a lower frequency response in our illumination can negatively affect timing resolution [19] and subsequently depth recovery performance.

Peak laser power: In addition to matching sensor resolution, the ideal pulsed laser can produce ultra-high peak power. However, the peak intensity achievable in semiconductor lasers is tightly limited due to the highly nonlinear behavior caused by the high gain realized in a small device [59]. High depletion creates intensity-dependent beam distortion and self-focusing that destroys the diode [60]. Although higher power is achievable with larger diodes, this exacerbates the bandwidth problem. Moreover, due to eye-safety regulations, managing high peak power is crucial, as excessive peak power levels pose a significant risk of eye injury [61]. Reported peak optical powers range from ≤ 1 W in short-range flash-LiDARs [28] to hundreds or thousands of watts in long-range scanning systems [62]. However, reducing peak laser power to meet tighter system limits can cause a decrease in system performance, especially in low signal-to-noise ratio (SNR) regimes.

Impulse response irregularities: Beyond the limitations of bandwidth and peak power, real-world systems can introduce undesirable characteristics in the IRF, causing it to deviate from well-defined functions like a Gaussian pulse. These irregularities can arise from factors such as misaligned hardware, lens aberrations, and clock jitter. Consequently, these non-ideal impulse responses lead to sub-optimal performance of existing compressive histograms in practical single-photon 3D imaging.

4.2 Constrained Optimization:

We aim to jointly optimize the illumination and coding matrix under hardware constraints described in section 4.1. The system’s bandwidth limitation is characterized by its impulse response $h(t)$, which acts as a low-pass filter that smooths the input illumination and attenuates frequencies outside the supported bandwidth. Thus, $s(t)$ is the output illumination after the input illumination $f(t)$ is smoothed by $h(t)$. The goal is to design the input illumination $f(t)$.

For our simulation results, we model the IRF as a Gaussian pulse with width σ , i.e., $h(t) \propto e^{-t^2/\sigma}$. However, in real-world systems, $h(t)$ may deviate from a Gaussian shape. Nonetheless, our optimization framework is flexible and can accommodate arbitrary parameterized IRFs. For a pulsed illumination, the output illumination is simply the system’s IRF, i.e., $s(t) = h(t)$. The peak power of our system constrains the input illumination, ensuring that $f(t)$ is bounded by the maximum laser power. Let f, h, s be vectors representing the N -point discretization of $f(t), h(t), s(t)$, respectively. Therefore, finding f, D can be formulated as a constrained optimization problem, represented mathematically as:

$$\begin{aligned} \min_{f, D} \quad & l((t_{d_j})_{j=0}^{J-1}, (\hat{t}_{d_j})_{j=0}^{J-1}) \\ \text{subject to} \quad & (f \otimes h)_i = s_i \\ & f_i \leq \Phi^{\max}, i = 0, \dots, N-1 \end{aligned} \quad (5)$$

where l is the loss function between J true depths $(t_{d_j})_{j=0}^{J-1}$ and predicted depths $(\hat{t}_{d_j})_{j=0}^{J-1}$. Φ^{\max} is the peak photon count, expressed as $\Phi^{\max} = p^{\text{factor}} * \Phi^{\text{sig}}$, where $\Phi^{\text{sig}} = \sum_i^{N-1} \Phi_i^{\text{sig}}$ is the mean signal photon count. Bandwidth and IRF constraints are addressed in the second line of 5, whereas the third line contains the peak power limitations. Predicted depth accuracy is strongly influenced by the output illumination $s(t)$ and the coding matrix D . Given that the system response $h(t)$ is fixed, our goal is to determine the optimal input illumination $f(t)$ and coding matrix D that minimizes our loss function l .

5 IMPLEMENTATION AND TRAINING

To solve the constrained optimization problem, we construct a network that can be divided into two sections: one for the coding matrix and one for the illumination.

Coding matrix section: The coding matrix D in our network is defined as a 1D convolutional layer with kernel size N and shape $K \times N$, which is equivalent to a fully connected linear layer. This design ensures the coding matrix can be implemented on-chip. To decode depths, we implement a ZNCC layer as defined in (4). However, since the argmax operation in the ZNCC depth decoding is not differentiable, we use softmax.

Illumination section: The illumination section enforces all the hardware constraints described in section 4.1. Any bandwidth or IRF limitations are integrated through a 1D convolutional layer that performs the circular convolution between h and f . Peak power constraints are implemented through a clamping layer that clips f at the peak photon count Φ^{\max} . For peak power-limited and band-limited illuminations, we filter with h after clamping to remove high frequencies introduced by clipping. To facilitate training, we initialize our illumination to $f_i = 1.0, i = 0, \dots, N-1$.

Training with Poisson noise: Since we are optimizing for SPADs, the dominant noise source is Poisson. To simulate a measured histogram, we apply Poisson noise to the incident vector $((r_i)_{i=0}^{N-1})$ in the network. However, Poisson noise is not differentiable and cannot be implemented without breaking gradient flow. Therefore, to mimic Poisson noise, we approximate it with Gaussian noise where the mean




Fig. 4. **Simulated results with finite bandwidth and infinite peak power.** This figure shows the mean absolute error on the left and the root mean squared error on the right, obtained from Monte Carlo simulations using the different coding schemes, evaluated at an IRF width of $\sigma = 30\Delta$ and infinite peak power. The optimized coding functions achieve the closest accuracy to the FRH, regardless of SBR, when the mean photon count is above 150. For photon counts below 150, both the Gray and the optimized codings exhibit a large gap between RMSE and MAE due to the presence of outliers.

equals the variance. When photon counts are sufficiently high, this Gaussian noise converges to Poisson noise. To ensure our coding functions adapt to SPAD sensors, we test all of our results using Poisson noise.

The network uses gradient descent to jointly optimize both the illumination and the coding matrix. First, the illumination’s waveform is optimized according to the provided hardware constraints. Then the coding matrix is adjusted so that each row’s frequency content most effectively samples the optimized illumination. Both the illumination parameters and coding matrix weights are updated together at every backpropagation step.

Note that we generate distinct coding functions corresponding to a specific set of hardware constraints. Although these coding functions vary according to hardware constraints, they remain constant across different SNR levels. Consequently, optimal performance is achieved when the coding functions are applied under the specific hardware constraints for which they were optimized.

5.1 Training

The dataset used for training consisted of generated labels, where each label represented a ground truth depth and associated SNR level. For our experiments, coding functions were trained in a moderate SNR regime. We trained all models using the ADAM optimizer [63], and used a learning rate of 0.013 for bandwidth-limited results and 0.0018 for peak power-limited results. The learning rate decay was set to 0.3 – 0.4 for both the peak power-limited and bandwidth-limited results. It took between 10 and 30 epochs to converge to the local minimum, with fewer epochs required for the bandwidth-limited results. We trained using the L1 loss function with total variation (TV) regularization applied to the 1D convolution layer (coding matrix) weights to reduce noise. See the supplement for more training and implementation details.

6 EVALUATION OF OPTIMIZED CODES

The following section evaluates our optimized coding functions at various SNR levels with respect to hardware limitations discussed in section 4.1.

6.1 Baselines and Performance Metrics

We compare our method against the following baselines:

- **Truncated Fourier [18]:** This coding matrix D is the first K rows of the discrete Fourier transform matrix, not including the zeroth harmonic.
- **Continuous Gray [19]:** This coding matrix D is based on the Hamiltonian codes introduced in [41]. The rows are the K -bit Gray Codes [64] linearly interpolated to length N .
- **Full-resolution Histogramming (FRH):** Coding matrix D is the $N \times N$ identity matrix. Full-resolution histogramming provides no compression but offers the most accurate depth recovery.

For the baselines, we use a pulsed illumination. To determine the efficacy of a code, we computed its root mean squared error (RMSE) and mean absolute error (MAE) through Monte Carlo simulations over the depth range for different photon counts and signal-to-background ratio (SBR) levels. RMSE and MAE are easily comparable performance metrics and are well-defined indicators of a code’s performance [41, 51]. See the supplement for further details on performance metrics and Monte Carlo simulation.

6.2 Optimized Band-limited Coding Functions

Given finite bandwidth, but infinite peak power, the optimal band-limited illumination converges to the IRF, $h(t)$. This suggests that, with infinite peak power, the best-performing illumination is the narrowest pulse feasible by the hardware. Nevertheless, reduced bandwidth leads to an increase in pulse width, making system performance highly dependent on the choice of compressive histogram. Coding matrix rows that contain frequency components beyond the illumination’s bandwidth contribute nothing to reconstruction, as they are zeroed after filtering. High-frequency codes, such as Gray-based codes, will be inefficient at encoding an illumination with a low frequency response. Conversely, codes such as Truncated Fourier are effective for low-frequency waveforms, but may fail to capture higher frequency components as bandwidth increases. As a result, neither Truncated Fourier nor Gray codes are perfectly suited for varying impulse responses and may underperform with a band-limited illumination. On the other hand, our proposed codes are explicitly optimized to match the system’s impulse response. This leads to a coding matrix that maintains a balance between high and low frequencies and is robust across different bandwidths. See the supplement for visualizations of our optimized band-limited coding functions.

6.3 Band-limited System Performance

In this section, we evaluate the performance of the baselines described in Section 6.1 and our optimized coding functions under conditions where the bandwidth is finite but infinite peak laser power. Fig. 4 shows the resulting MAE and RMSE from our Monte Carlo simulations with an IRF width of $\sigma = 30\Delta$, where $K = 8$ for our compressive codes and $N = 1024$ for the FRH. As the illumination bandwidth decreases ($\sigma > \Delta$), the effective time resolution of our system decreases, leading to diminished FRH

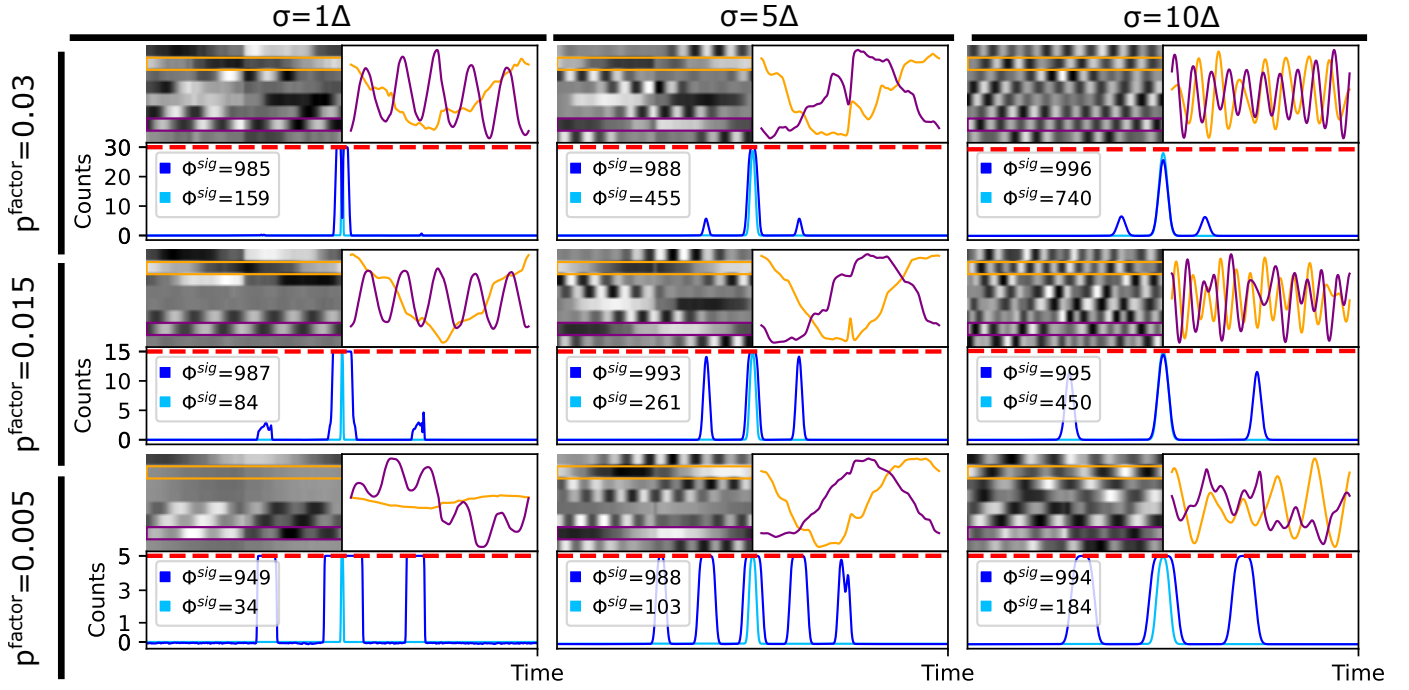


Fig. 5. **Hardware-constrained optimized coding functions.** This figure plots the optimized coding functions found for different bandwidths and peak power constraints. The columns indicate the width of the IRF $h(t)$ and the rows indicate the peak factor p^{factor} . For each p^{factor} and σ we display the coding matrix D on the top right, display the 2nd and 7th row of D on the top left, and display the pulsed illumination (light blue) and the optimized illumination (blue) incident on the sensor at the bottom. Here we optimize our codes with a constant mean photon count of $\Phi^{\text{sig}} = 1000$. Note that the middle of the illumination plot is zero.

performance compared to the ideal pulse ($\sigma = \Delta$). In this strict bandwidth regime, Gray-based codes experienced the greatest performance degradation, whereas Truncated Fourier exhibited minimal loss. Regardless, we identify an optimized coding function that approaches FRH performance more closely than either compressive histograms. As indicated by the RMSE and MAE, for photon counts above 150 and SBR above 0.15, the optimized codes yield the fewest outliers and produce less variance in depth estimation compared to Truncated Fourier and Gray codes. The higher RMSE observed for the optimized coding functions at photon counts ≤ 150 arises from outliers. We attribute these outliers to training on Gaussian noise (mean equal to variance) as an approximation of Poisson noise, as well as emphasizing medium-to-high photon counts. Nevertheless, the optimized coding functions continue to achieve near-zero errors for the majority of depths, as shown by the MAE Monte Carlo simulation (Fig. 4). See the supplement for additional band-limited results.

6.4 Optimized Peak Power-limited Coding Functions

Fig. 5 shows the resulting illuminations and coding matrices after optimizing for peak factors $p^{\text{factor}} = [0.005, 0.015, 0.030]$, IRF widths $\sigma = [1\Delta, 5\Delta, 10\Delta]$, and a constant signal photon count of $\Phi^{\text{sig}} = 1000$. Under conditions of finite bandwidth and peak power, we observe that the optimal illumination diverges from the pulsed waveform. Instead, the optimized illumination converges to a symmetric waveform with multiple peaks, almost resembling a square wave. This behavior becomes more pronounced as the peak photon count decreases. Notably,

several of the illuminations are binary and, thus, more efficient to implement in hardware. For conventional pulsed illumination, a peak power limitation can lead to a reduction in the total signal photons Φ^{sig} when clipping. To mitigate this photon loss, the optimized illumination adapts by re-distributing photons dependent on the severity of the peak factor p^{factor} . However, as we increase p^{factor} , the optimal illumination converges back to the IRF. Fig. 5 illustrates how the coding matrix adapts in response to illumination changes. Notably, for the optimized illumination characterized by an IRF width of $\sigma = 1\Delta$ and a peak factor of $p^{\text{factor}} = [0.015, 0.005]$, several of the rows within the coding matrix are nearly zero. This implies these near-zero rows do not encode any important frequency information about our illumination. Consequently, only a subset of rows in the coding matrix is necessary for accurate depth decoding, and the remainder can be potentially discarded. Note that the total signal photon count Φ^{sig} and peak photon count Φ^{max} are fixed during optimization, ensuring that the illumination does not violate these constraints.

6.5 Peak Power-Limited System Performance

In this section, we analyze each of the baselines described in section 6.1 and our optimized coding functions under finite bandwidth and finite peak power. In our experiments, we benchmark the coding functions against pulse-based methods where peak power is constrained in two ways: clipping the maximum pulse energy at Φ^{max} , and maintaining constant pulse energy by increasing pulse width.

Clipped energy results: Fig. 6 shows the RMSE and MAE for the different compressive codes ($K = 8$) and FRH




Fig. 6. **Simulated results with finite bandwidth and finite peak power when clipping pulse energy.** This figure shows the MAE on the left and the RMSE on the right, obtained from Monte Carlo simulations using the different coding schemes, evaluated at an IRF width of $\sigma = 5\Delta$ and peak factor of $p^{\text{factor}} = 0.05$ (clipped).

($N = 1024$) tested with a peak factor of $p^{\text{factor}} = 0.005$ and an IRF width of $\sigma = 5\Delta$ when clipping pulse energy at Φ^{max} . Clipping significantly degrades the performance of pulse-based methods due to reduced signal photons. In contrast, the optimized illumination adapts by maximizing the number of signal photons permitted by the system while not exceeding the peak photon count. Consequently, the optimized coding functions achieve the lowest RMSE and MAE across all photon counts among compressive histogram baselines, and even outperform FRH at counts below 1500. At low photon counts, pulse-based methods become unstable under clipping, while the optimized coding functions maintain performance despite the limited peak laser power.

Preserving energy results: Fig. 7 shows the RMSE and MAE for the different compressive codes ($K = 8$) tested with a peak factor of $p^{\text{factor}} = 0.005$ and an IRF width of $\sigma = 5\Delta$ when maintaining constant pulse energy. In this configuration, increasing the pulse width avoids clipping, resulting in significant improvement for pulse-based methods. However, this results in reduced temporal resolution and lower peak power in the pulse waveform. Consequently, preserving pulse energy does not overcome the limitations imposed by peak laser power, which becomes the main bottleneck for the pulsed baselines (at $p^{\text{factor}} = 0.005$). As a result, the optimized coding functions demonstrate the lowest RMSE and MAE among compressive histograms and FRH when signal photon counts are above 200. Notably, FRH never surpasses the optimized codes. For photon counts below 200, the optimized codes exhibit a wider gap in RMSE and MAE due to occasional outliers (from training on Gaussian noise), yet still achieve near-zero error for the majority of depths.

These findings indicate that clipping and preserving pulse energy yield suboptimal performance of pulse-based methods, even for FRH. To avoid performance loss in low-to-medium peak power regimes, it is essential to reconsider the choice of illumination and coding matrix based on hardware constraints. Please see the supplement for additional peak power-limited results.

7 COMPRESSIVE SINGLE-PHOTON 3D IMAGING

In this section, we evaluate the performance of our proposed coding functions on real-world data from a scanning-




Fig. 7. **Simulated results with finite bandwidth and finite peak power when preserving pulse energy.** This figure shows the MAE on the left and the RMSE on the right, obtained from Monte Carlo simulations using the different coding schemes, evaluated at an IRF width of $\sigma = 5\Delta$ and peak factor of $p^{\text{factor}} = 0.05$ (constant).

based hardware setup [51], and simulated flash-illumination scenes. Additionally, we analyze the memory efficiency of our optimized coding functions for an on-chip implementation.

7.1 Real-world Scanning-based System Results

To evaluate the ability of our optimized codes to adapt to real-world hardware, we use data acquired from a scanning-based single-photon system [51], which has been preprocessed as described in [19]. The post-processed histogram data have $\Delta = 8$ ps and $N = 2188$. The system IRF, shown in Fig. 8, exhibits a double peak, likely due to lens inter-reflections in the hardware setup. Fig. 8 shows the recovered 3D reconstruction of a porcelain face scan from compressive histograms and our optimized coding functions, where $K = 10$. When presented with this atypical IRF, compressive histograms suffer from performance degradation, even when the IRF is taken into account. In particular, truncated Fourier exhibits systematic depth errors, implying this code struggles to resolve the double peaks in the IRF. While Gray-based codes provide more stable results, they are prone to moderate depth errors across most pixels and large errors in certain regions (outliers). In contrast, our optimized coding demonstrates robustness to the irregular IRF and yields the closest reconstruction to the FRH. Additionally, the hardware scans in Fig. 8 include pixels with dense indirect reflections, which appear as extended tails in the histogram. As a result, all the compressive codes show higher RMSE values due to these pixels. Nonetheless, the optimized coding still recovers most depths with minimal error, while remaining robust to pixels with dense indirect reflections. Note that only the coding matrix was optimized for the hardware results. Please refer to the supplement for additional hardware results.

7.2 Simulated Flash-Illumination System Results

To evaluate the performance of our optimized coding functions in a flash-illumination-based system, we used rendered histogram data generated with MitsubaToF [65, 66], obtained from [67]. Here, we have $N = 2048$ for the band-limited results and $N = 1024$ for the peak power-limited results. Different illuminations are simulated by convolving the incident waveform with the rendered histograms and

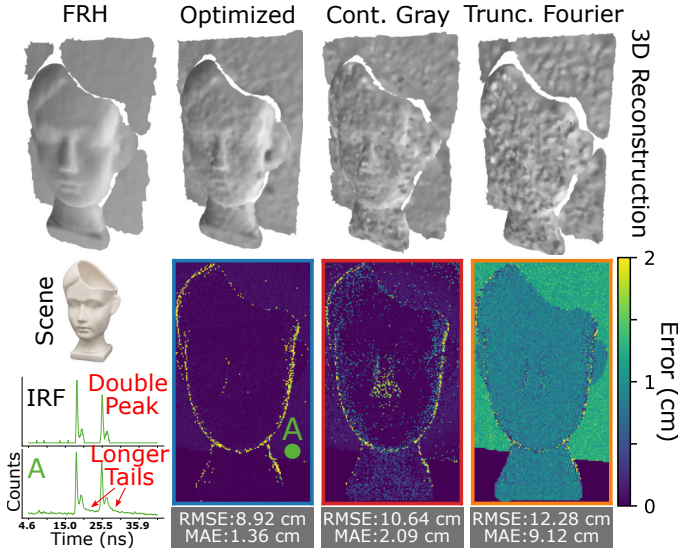


Fig. 8. **3D reconstruction visualization.** Real-world scans of a porcelain face obtained using Gray codes, truncated Fourier, and our optimized codes. RMSE and MAE are calculated from the FRH as the reference. Several background histograms (blue dot) show longer tails, which are likely due to indirect reflections.

then scaling and offsetting the results based on the desired photon count and SBR. The resulting scenes include both dense and sparse indirect reflections.

Band-limited coding functions: Fig. 9 plots the resulting depth maps and error for compressive histograms and the optimized coding functions with an IRF width $\sigma = 30\Delta$ and infinite peak power where $K = 12$ ($K = 11$ for Gray codes). Truncated Fourier produces systematic errors in the presence of indirect reflections, but remains relatively robust to large errors. Although Gray-based coding handles indirect reflections more effectively, this advantage diminishes under low-bandwidth constraints. The band-limited optimized coding functions are robust to incident waveforms that do not deviate far from the trained waveform, including both dense and sparse indirect reflections with a small secondary peak. However, they are susceptible to outliers caused by large, sparse indirect reflections. These outliers can be mitigated by increasing K as shown by the figure results. Overall, the optimized coding functions still outperform compressive histograms in terms of both RMSE and MAE, even under moderate indirect reflections.

Peak power-limited coding functions: Fig. 10 plots the resulting depth maps and error for compressive histograms and the optimized coding functions with IRF width $\sigma = 10\Delta$ and peak factor of $p^{\text{factor}} = 0.005$ when maintaining constant pulse energy. Similar to the band-limited results, the optimized coding functions are prone to outliers in the presence of sparse indirect reflections, which is further exacerbated by the multiple peaks in the peak power-limited illumination. However, in severely constrained systems, peak laser power becomes the limiting factor in coding scheme performance rather than indirect reflections. For the pulse-based methods, preserving energy by widening the pulse does not compensate for the lack of peak laser power. Consequently, our optimized coding functions achieve the lowest MAE among all baselines (when $p^{\text{factor}} = 0.005$),

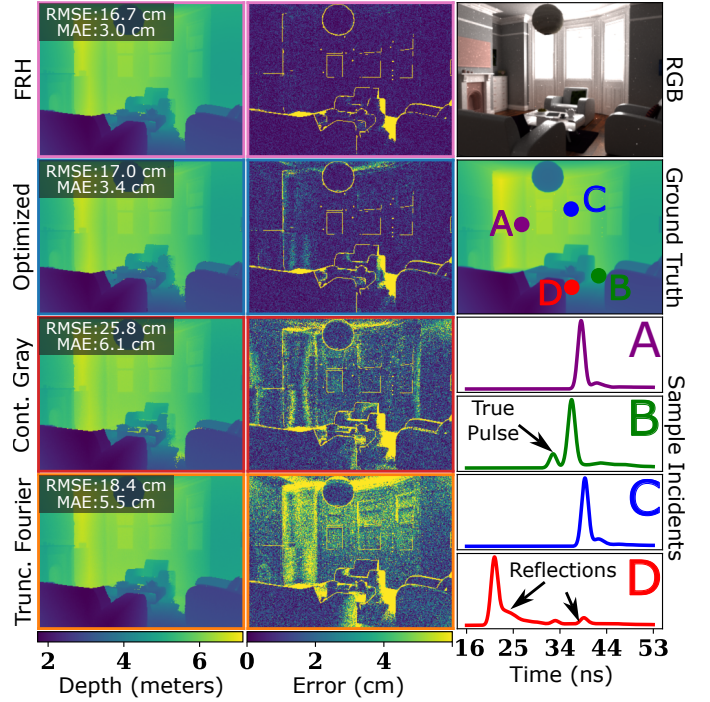


Fig. 9. **Simulated flash-illumination for band-limited coding functions.** Depth maps (first col) and error (second col) results for different coding schemes (rows) evaluated with IRF width $\sigma = 30\Delta$ and infinite peak power. The rightmost column plots the RGB image, ground truth depth map, and four incident functions taken from the scene in that order.

including FRH, despite challenges from indirect reflections. Please refer to the supplement for additional flash-illumination results and an analysis of the optimized coding functions in the presence of different indirect reflections.

7.3 Memory-Efficiency of Optimized Coding Matrices

One benefit of Fourier and Gray coding functions is that they have memory-efficient representations [22, 24]. A naive in-pixel approach would require storing the full 64-bit coding matrix, which, unless shared by multiple pixels, would lead to a sizeable in-sensor memory overhead. Therefore, an efficient representation of the optimized coding matrix is essential for on-chip implementation. To assess this, we evaluate the performance (Fig. 11) of the optimized coding functions when quantized to values between 1 and 64 bits and reconstructed using 10 to 90 Fourier coefficients. We find that both peak power-limited and band-limited coding matrices can be represented using as few as 4 bits or 40 Fourier coefficients with minimal performance loss. Gutierrez-Barragan et al. [22] recently showed that 8-bit Fourier codes, when implemented on the UltraPhase chip [24], consume less power than data-intensive codes like coarse histogramming [25]. Together, these findings suggest that our proposed coding functions have a feasible on-chip implementation under typical in-sensor memory constraints. Please see the supplement for memory analysis of additional coding functions.

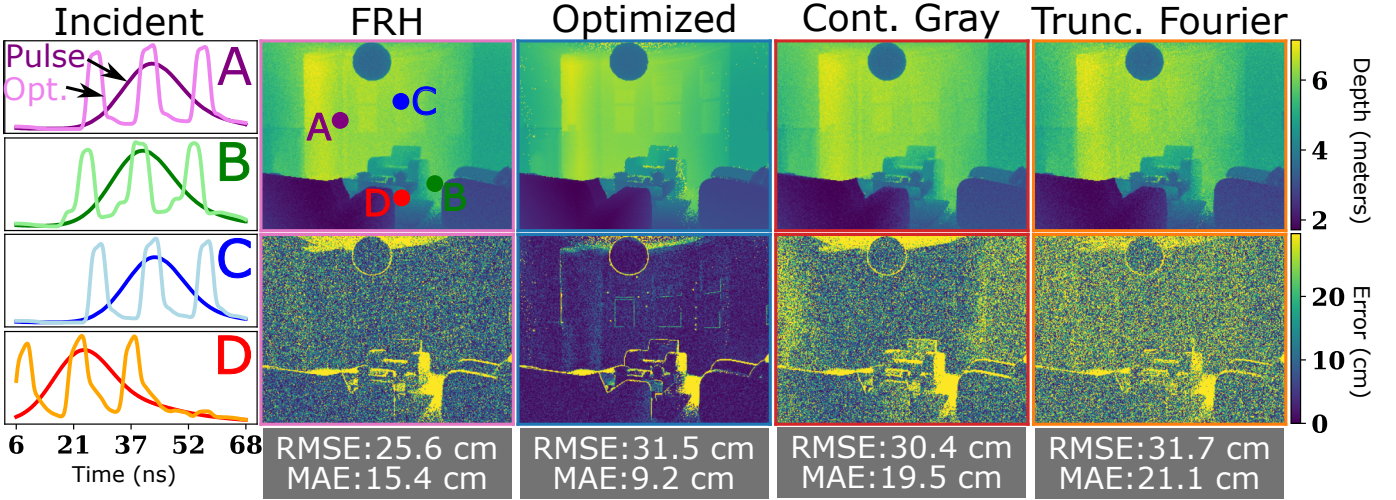


Fig. 10. **Simulated flash-illumination for peak power-limited coding functions.** Depth maps (first row) and error (second row) results for different coding schemes (cols) when preserving pulse energy evaluated at IRF width $\sigma = 10\Delta$ and $p^{\text{factor}} = 0.005$. The leftmost column plots the four incident functions taken from the scene for both the pulsed and optimized illuminations.

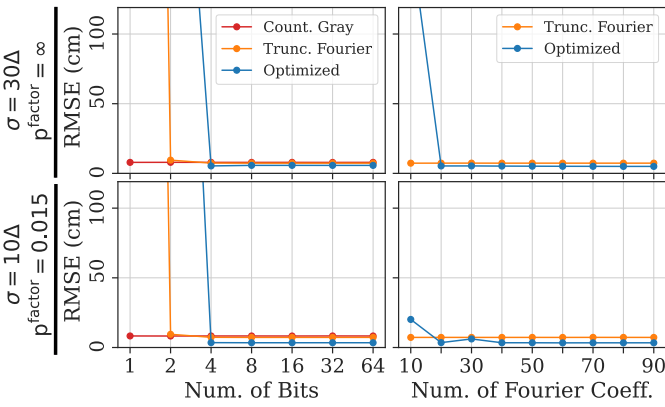


Fig. 11. **Memory efficiency analysis of coding functions.** This figure plots the RMSE of each compressive code as a function of bit depth (left col) and number of Fourier coefficients (right col) used to represent the coding matrix, evaluated at low SBR (0.1). The top row presents results for band-limited coding functions with infinite peak power, while the bottom row shows results for the peak power-limited coding functions.

8 DISCUSSION AND LIMITATIONS

Hardware constraints, such as laser power, sensor data rates, system bandwidth, and the illumination waveform, strongly influence the performance and practicality of single-photon 3D cameras. While compressive histograms were introduced to reduce data rates, they perform suboptimally under real-world hardware limitations. We propose a constrained optimization approach that identifies a set of high-performing coding functions while adhering to hardware constraints. We show that our method consistently outperforms compressive histograms across a wide range of scenarios. We hope this work takes a step toward enabling low-cost, low-power, and practical single-photon 3D imaging systems.

Poisson noise in training neural networks: Poisson noise is not differentiable. Therefore, to train our illumination, we approximated Poisson noise through Gaussian noise where the mean equals the variance. However, at low photon

counts, this approximation is weak. In the photon-starved regime, our optimized coding functions can be prone to large depth errors, likely due to this approximation. A possible workaround is to optimize the illumination with Gaussian noise and then optimize the coding matrix using Poisson noise to improve the code’s ability to adapt to single-photon sensors at low SNR.

Scene Dependent Considerations: We proposed a framework for designing coding functions optimized to a specific hardware impulse response. However, our optimization did not account for scene-dependent factors such as indirect reflections, surface albedo, or geometry. While the experimental and simulation results show that several coding functions are fairly robust to indirect reflections, this was not explicitly optimized. On the other hand, the performance of the peak power-limited coding functions may degrade due to the presence of multiple peaks in the illumination waveform. A possible solution would be to incorporate scene response into our network by adding a convolution layer.

Framework for practical time-of-flight imaging: While this paper focused on finding optimal coding functions for single photon ToF, this optimization framework is potentially viable for other imaging modalities that encode a transmitted signal, provided the decoding is differentiable. Examples include correlation-based ToF (C-ToF) 3D imaging, NLoS imaging, FLIM imaging, and gated imaging. Consider a C-ToF camera whose illumination is temporally modulated according to the function $m(t)$. A brightness value is computed as the correlation between the radiance incidence $r(t)$ and the sensor exposure function $d(t)$. Our proposed framework can be applied to finding high-performing modulation $m(t)$ and demodulation functions $d(t)$ that adhere to hardware constraints. Furthermore, our model has the potential to be extended to other ToF imaging methods such as NLoS and FLIM. However, we anticipate that applications needing full temporal information will achieve lower compression due to increased signal complexity and photon noise.

REFERENCES

- [1] S. Cova, M. Ghioni, A. Lacaita, C. Samori, and F. Zappa, "Avalanche photodiodes and quenching circuits for single-photon detection," *Applied Optics*, vol. 35, no. 12, pp. 1956–1976, 1996.
- [2] E. R. Fossum, J. Ma, S. Masoodian, L. Anzagira, and R. Zizza, "The quanta image sensor: Every photon counts," *Sensors*, vol. 16, no. 8, p. 1260, 2016.
- [3] S. Pellegrini, G. S. Buller, J. M. Smith, A. M. Wallace, and S. Cova, "Laser-based distance measurement using picosecond resolution time-correlated single-photon counting," *Measurement Science and Technology*, vol. 11, no. 6, p. 712, 2000.
- [4] J. Tachella, Y. Altmann, N. Mellado, A. McCarthy, R. Tobin, G. S. Buller, J.-Y. Tourneret, and S. McLaughlin, "Real-time 3d reconstruction from single-photon lidar data using plug-and-play point cloud denoisers," *Nature Communications*, vol. 10, no. 1, pp. 1–6, 2019.
- [5] A. Gupta, A. Ingle, A. Velten, and M. Gupta, "Photon-flooded single-photon 3d cameras," in *Proc. CVPR*, June 2019.
- [6] J. Rapp, Y. Ma, R. M. Dawson, and V. K. Goyal, "High-flux single-photon lidar," *Optica*, vol. 8, no. 1, pp. 30–39, 2021.
- [7] J. Rapp, J. Tachella, Y. Altmann, S. McLaughlin, and V. K. Goyal, "Advances in single-photon lidar for autonomous vehicles: Working principles, challenges, and recent advances," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 62–71, 2020.
- [8] M. Buttafava, J. Zeman, A. Tosi, K. Eliceiri, and A. Velten, "Non-line-of-sight imaging using a time-gated single photon avalanche diode," *Optics Express*, vol. 23, no. 16, pp. 20997–21011, 2015.
- [9] X. Liu, I. Guillén, M. La Manna, J. H. Nam, S. A. Reza, T. H. Le, A. Jarabo, D. Gutierrez, and A. Velten, "Non-line-of-sight imaging using phasor-field virtual wave optics," *Nature*, pp. 1–4, 2019.
- [10] J. Lee, J. V. Chacko, B. Dai, S. A. Reza, A. K. Sagar, K. W. Eliceiri, A. Velten, and M. Gupta, "Coding scheme optimization for fast fluorescence lifetime imaging," *ACM Trans. Graph.*, vol. 38, no. 3, jun 2019.
- [11] Ouster, "Os1 lidar sensor," <https://ouster.com/products/hardware/os1-lidar-sensor>, accessed: 7 April 2025.
- [12] C. Zhang, N. Zhang, Z. Ma, L. Wang, Y. Qin, J. Jia, and K. Zang, "A 240× 160 3d-stacked spad dtof image sensor with rolling shutter and in-pixel histogram for mobile devices," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 2, pp. 3–11, 2021.
- [13] J. Nissinen and J. Kostamovaara, "A high repetition rate cmos driver for high-energy sub-ns laser pulse generation in spad-based time-of-flight range finding," *IEEE Sensors Journal*, vol. 16, no. 6, pp. 1628–1633, 2016.
- [14] B. Lashkari and A. Mandelis, "Comparison between pulsed laser and frequency-domain photoacoustic modalities: Signal-to-noise ratio, contrast, resolution, and maximum depth detectivity," *Review of Scientific Instruments*, vol. 82, no. 9, p. 094903, 09 2011.
- [15] P. K. Upputuri and M. Pramanik, "Performance characterization of low-cost, high-speed, portable pulsed laser diode photoacoustic tomography (pld-pat) system," *Biomed. Opt. Express*, vol. 6, no. 10, pp. 4118–4129, Oct 2015.
- [16] B. Rosenstein, A. A. Ishaaya, D. Belker, and A. Shirakov, "Overcoming peak-power limitations of fiber lasers," *SPIE Newsroom*, 2014.
- [17] Y. Jaouën, G. Canat, S. Grot, and S. Bordais, "Power limitation induced by nonlinear effects in pulsed high-power fiber amplifiers," *Comptes Rendus Physique*, vol. 7, no. 2, pp. 163–169, 2006, high power fiber lasers and amplifiers.
- [18] M. P. Sheehan, J. Tachella, and M. E. Davies, "A sketching framework for reduced data transfer in photon counting lidar," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 989–1004, 2021.
- [19] F. Gutierrez-Barragan, A. Ingle, T. Seets, M. Gupta, and A. Velten, "Compressive single-photon 3d cameras," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17854–17864.
- [20] A. C. Ulku, C. Bruschini, I. M. Antolovic, E. Charbon, Y. Kuo, R. Ankri, S. Weiss, and X. Michalet, "A 512×512 spad image sensor with integrated gating for widefield flim," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 25, no. 1, p. 6801212, Jan-Feb 2019.
- [21] S. Hutchings, N. Johnston, I. Gyongy, T. Al Abbas, N. Dutton, M. Tyler, S. Chan, J. Leach, and R. Henderson, "A reconfigurable 3-d-stacked spad imager with in-pixel histogramming for flash lidar or high-speed time-of-flight imaging," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 11, pp. 2947–2956, Nov. 2019.
- [22] F. Gutierrez-Barragan, F. Mu, A. Ardelean, A. Ingle, C. Bruschini, E. Charbon, Y. Li, M. Gupta, and A. Velten, "Learned compressive representations for single-photon 3d imaging," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 10756–10766.
- [23] MIPI Alliance, "Evolving CSI-2 Specification – technology brief," https://www.mipi.org/sites/default/files/MIPI_CSI-2_Specification_Brief.pdf, mIPI Alliance, Accessed April 8, 2025.
- [24] A. Ardelean, "Computational imaging spad cameras," PhD thesis, École Polytechnique Fédérale de Lausanne (EPFL), 2023.
- [25] I. Gyongy, S. W. Hutchings, A. Halimi, M. Tyler, S. Chan, F. Zhu, S. McLaughlin, R. K. Henderson, and J. Leach, "High-speed 3d sensing via hybrid-mode imaging and guided upsampling," *Optica*, vol. 7, no. 10, pp. 1253–1260, 2020.
- [26] V. Poisson, W. Guicquero, G. Sicard *et al.*, "Luminance-depth reconstruction from compressed time-of-flight histograms," *IEEE Transactions on Computational Imaging*, vol. 8, pp. 148–161, 2022.
- [27] X. Ren, P. W. R. Connolly, A. Halimi, Y. Altmann, S. McLaughlin, I. Gyongy, R. K. Henderson, and G. S. Buller, "High-resolution depth profiling using a range-gated cmos spad quanta image sensor," *Opt. Express*, vol. 26, no. 5, pp. 5541–5557, Mar 2018.
- [28] C. Zhang, S. Lindner, I. M. Antolović, J. Mata Pavia, M. Wolf, and E. Charbon, "A 30-frames/s, 252 × 144 spad flash lidar with 1728 dual-clock 48.8-ps tdc, and pixel-wise integrated histogramming," *IEEE Journal of*

- Solid-State Circuits*, vol. 54, no. 4, pp. 1137–1151, 2019.
- [29] M. P. Sheehan, Y. Sun, and P. T. Boulinguez, “Surface detection for sketched single photon lidar,” *EUSIPCO*, 2021.
- [30] A. Ingle and D. Maier, “Count-free single-photon 3d imaging with race logic,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [31] K. Sadekar, D. Maier, and A. Ingle, “Single-photon 3d imaging with equi-depth photon histograms,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024, to appear.
- [32] J. Folden, A. Ingle, and S. J. Koppal, “Foveaspad: Exploiting depth priors for adaptive and efficient single-photon 3d imaging,” *IEEE Transactions on Computational Imaging*, vol. 10, pp. 1722–1735, 2024.
- [33] V. Sundar, A. Ardelean, T. Swedish, C. Bruschini, E. Charbon, and M. Gupta, “Sodacam: Software-defined cameras via single-photon imaging,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 8165–8176.
- [34] V. Sundar, M. Dutson, A. Ardelean, C. Bruschini, E. Charbon, and M. Gupta, “Generalized event cameras,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 25007–25017.
- [35] D. Yao, G. Mora-Martín, I. Gyongy, S. Scholes, J. Leach, S. McLaughlin, and Y. Altmann, “Bayesian neuromorphic imaging for single-photon lidar,” *Opt. Express*, vol. 32, no. 14, pp. 25 147–25 164, Jul 2024.
- [36] F. Mu, C. Sifferman, S. Jungerman, Y. Li, M. Han, M. Gleicher, M. Gupta, and Y. Li, “Towards 3d vision with low-cost single-photon cameras,” 2024.
- [37] Y. Li, X. Liu, W. Dong, H. Zhou, H. Bao, G. Zhang, Y. Zhang, and Z. Cui, “Deltar: Depth estimation from a light-weight tof sensor and rgb image,” 2022.
- [38] A. Ruget, L. Wilson, J. Leach, R. Tobin, A. McCarthy, G. S. Buller, S. McLaughlin, and A. Halimi, “Plug-and-play algorithm for 3d guided video super-resolution of single-photon lidar data,” *Opt. Express*, vol. 33, no. 5, pp. 11 775–11 792, Mar 2025.
- [39] M. Sun, S. Zhuo, and P. Y. Chiang, “Multi-scale histogram-based probabilistic deep neural network for super-resolution 3d lidar imaging,” *Sensors*, vol. 23, no. 1, p. 420, 2022.
- [40] G. Mora-Martín, S. Scholes, A. Ruget, R. Henderson, J. Leach, and I. Gyongy, “Video super-resolution for single-photon lidar,” *Opt. Express*, vol. 31, no. 5, pp. 7060–7072, Feb 2023.
- [41] M. Gupta, A. Velten, S. K. Nayar, and E. Breitbach, “What are optimal coding functions for time-of-flight imaging?” *ACM Transactions on Graphics (TOG)*, vol. 37, no. 2, p. 13, 2018.
- [42] F. Gutierrez-Barragan, S. A. Reza, A. Velten, and M. Gupta, “Practical coding function design for time-of-flight imaging,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1566–1574.
- [43] J. Li, T. Yue, S. Zhao, and X. Hu, “Fisher information guidance for learned time-of-flight imaging,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 16 334–16 343.
- [44] M. Gupta, S. K. Nayar, M. B. Hullin, and J. Martin, “Phasor imaging: A generalization of correlation-based time-of-flight imaging,” *ACM Transactions on Graphics (ToG)*, vol. 34, no. 5, pp. 156:1–156:18, 2015.
- [45] A. Kadambi, R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar, “Coded time of flight cameras: Sparse deconvolution to address multipath interference and recover time profiles,” *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, pp. 167:1–167:10, 2013.
- [46] R. Tadano, A. K. Pediredla, and A. Veeraraghavan, “Depth selective camera: A direct, on-chip, programmable technique for depth selectivity in photography,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 3595–3603.
- [47] W. Chen, P. Mirdehghan, S. Fidler, and K. N. Kutulakos, “Auto-tuning structured light by optical stochastic gradient descent,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 5970–5980.
- [48] M. Gupta, A. Agrawal, A. Veeraraghavan, and S. G. Narasimhan, “Structured light 3d scanning in the presence of global illumination,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 713–720.
- [49] M. Gupta and N. Nakhate, “A geometric perspective on structured light coding,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 87–102.
- [50] P. Mirdehghan, W. Chen, and K. N. Kutulakos, “Optimal structured light à la carte,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6248–6257.
- [51] A. Gupta, A. Ingle, and M. Gupta, “Asynchronous single-photon 3d imaging,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 7909–7918.
- [52] G. Turin, “An introduction to matched filters,” *IRE Transactions on Information Theory*, vol. 6, no. 3, pp. 311–329, 1960.
- [53] N. Hagen, M. Kupinski, and E. L. Dereniak, “Gaussian profile estimation in one dimension,” *Applied Optics*, vol. 46, no. 22, pp. 5374–5383, Aug 2007.
- [54] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. C. Wong, and J. H. Shapiro, “Photon-efficient imaging with a single-photon camera,” *Nature Communications*, vol. 7, p. 12046, 2016. [Online]. Available: <https://doi.org/10.1038/ncomms12046>
- [55] D. Li, C. Zhang, P. Zhou *et al.*, “Evolution of laser technology for automotive lidar, an industrial viewpoint,” *Nature Communications*, vol. 15, no. 1, p. 1234, 2024.
- [56] S. Wen, M. Wang, J. Xie, and D. Wu, “Large current nanosecond pulse generating circuit for driving semiconductor laser diode,” *Microwave and Optical Technology Letters*, vol. 61, no. 4, pp. 867–872, 2019.
- [57] Z. Chen, Y. Yan, J. Shu, K. Liu, and J. Qiu, “Modulated high power and narrow pulse width laser drive circuit for lidar system,” *Electronics*, vol. 10, no. 7, 2021.
- [58] D. Zhai, Z. Li, W. Zhang, M. Fei, S. Lu, X. Chen, H. Pan, and G. Wu, “High-precision uav-borne single-photon lidar by adaptive averaging,” *Opt.*

Express, vol. 33, no. 6, pp. 13 660–13 672, Mar 2025. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-33-6-13660>

- [59] J. Huang and L. W. Casperson, “Gain and saturation in semiconductor lasers,” *Optical and quantum electronics*, vol. 25, pp. 369–390, 1993.
- [60] T. Kitatani, M. Kondow, S. Nakatsuka, Y. Yazawa, and M. Okai, “Room-temperature lasing operation of gainnas-gaas single-quantum-well laser diodes,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 3, no. 2, pp. 206–209, 1997.
- [61] *ANSI Z136.1: American National Standard for Safe Use of Lasers*, American National Standards Institute (ANSI) Std., 2020, aNSI Z136.1-2020.
- [62] S. Royo and M. Ballesta-Garcia, “An overview of lidar imaging systems for autonomous vehicles,” *Applied Sciences*, vol. 9, no. 19, 2019. [Online]. Available: <https://www.mdpi.com/2076-3417/9/19/4093>
- [63] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [64] F. Gray, “Pulse code communication,” US Patent 2 632 058, 1953.
- [65] W. Jakob, “Mitsuba Renderer,” <http://www.mitsuba-renderer.org>, 2010, online; accessed 22 June 2025.
- [66] A. K. Pediredla, A. C. Sankaranarayanan, M. Buttafava, A. Tosi, and A. Veeraraghavan, “Signal processing based pile-up compensation for gated single-photon avalanche diodes,” *arXiv preprint arXiv:1806.07437*, 2018.
- [67] F. Gutierrez-Barragan, H. Chen, M. Gupta, A. Velten, and J. Gu, “itof2dtof: A robust and flexible representation for data-driven time-of-flight imaging,” *IEEE Transactions on Computational Imaging*, vol. 7, pp. 1205–1214, 2021.



Trevor Seets received a Ph.D. (2025) in Electrical Engineering from University of Wisconsin-Madison, Madison, WI in 2025. Currently, Trevor works as a Senior Software Engineer at Ubiccept. His research interests include machine learning, statistical signal processing, and computational imaging.



David Parra is a third-year PhD student in Computer Science at the University of Wisconsin-Madison, advised by Prof. Andreas Velten. His research interests are in computational imaging, computer vision, and deep learning. He earned his Bachelor’s degree from the University of Arizona (2021).



Felipe Gutierrez-Barragan is a Senior Software Engineer at Ubiccept, developing efficient image processing algorithms optimized for computer vision systems and emerging sensors such as high-speed HDR SPAD cameras. He received his B.S (2016), M.S. (2019), and Ph.D. (2022) in Computer Sciences from the University of Wisconsin-Madison. His Ph.D. research focused on practical modifications to spad-based and indirect time-of-flight 3D cameras, aiming to reduce their power consumption and data

bandwidth while preserving or enhancing their precision. His research interests are in computational imaging, computer vision, and machine learning.



Andreas Velten is Associate Professor at the Department of Biostatistics and Medical Informatics and the Department of Electrical and Computer Engineering at the University of Wisconsin-Madison and directs the Computational Optics Group. He obtained his PhD in Physics at the University of New Mexico in Albuquerque and was a postdoctoral associate of the Camera Culture Group at the MIT Media Lab. He has been included in the MIT TR35 list of the world’s top innovators under the age of 35 and

is a senior member of the National Academy of Inventors and SPIE as well as a fellow of Optica. He is co-Founder of Onlume, a company that develops surgical imaging systems, and Ubiccept, a company developing single photon imaging solutions.