

# Supplementary Document for “Hardware-aware Coding Function Design for Compressive Single-photon 3D Cameras”

David Parra<sup>\*1</sup>, Felipe Gutierrez-Barragan<sup>2</sup>, Trevor Seets<sup>1</sup>, and Andreas Velten<sup>†1</sup>

<sup>1</sup>University of Wisconsin-Madison

<sup>2</sup>Independent Researcher

{deparra, seets}@wisc.edu, felipegb94@gmail.com, velten@wisc.edu

## 1 Single-Photon Hardware Constrained Coding Functions

In this section, we visualize the hardware-constrained compressive histograms and optimized coding functions and provide additional intuition for hardware-constrained code design.

### 1.1 Band-limited Coding Functions

Here we analyze the bandwidth of single-photon coding schemes. Recall that the impulse response function of our system is defined by  $h(t)$  and its discretization  $h = (h_i)_{i=0}^{N-1}$ , where  $N$  is the number of time bins. Also, recall  $D'$  is the filtered coding matrix, which is the circular correlation between  $D$  and the impulse response function (IRF)  $h(t)$ . The discrete Fourier series of the IRF is given as:

$$h_i = \sum_{k=0}^b H_k e^{\frac{j2\pi ki}{N}}, \quad (1)$$

where  $j = \sqrt{-1}$ ,  $i = 0, \dots, N-1$ , and  $H_k$  are the Fourier coefficients. For a coding matrix  $D$  to be effective in encoding the pulsed waveform, none of its rows should contain frequencies higher than the system bandwidth  $b$ . If a row contains frequencies larger than  $b$ , then the inner product between that row and  $h$  is zero [1]. This is represented in Supp. Figs. 1, 2, and 3, which plot the input illumination  $f(t)$  and decoding matrix  $D'$  for coding schemes with finite bandwidth, but infinite peak power. For the results, we model the IRF as a Gaussian pulse of width  $\sigma$ , ie  $h(t) \propto e^{-t^2/\Delta}$ . As bandwidth decreases, the decoding matrix  $D'$  becomes increasingly smooth. This effect is most pronounced in the Continuous Gray-based compressive histogram, which exhibits zeroed-out rows for IRF widths  $\sigma = [20\Delta, 30\Delta]$ . While Continuous Gray performs efficiently with an ideal narrow pulse ( $\sigma = 1\Delta$ ), its performance deteriorates at lower bandwidths. In contrast, Truncated Fourier codes remain unchanged after filtering and are less sensitive to bandwidth limitations. However, for lower  $K$ , they may fail to capture higher-frequency features of the waveform. In contrast, our proposed coding is hardware-aware and optimizes specifically for the system’s impulse response. This leads to a coding matrix that better matches the waveform’s frequency response than traditional compressive histograms.

### 1.2 Peak Power-limited Coding Functions

In our experiments, we compare the optimized coding functions to pulse-based methods under two implementations of peak power limitations: (1) clipping pulse energy and (2) preserving

---

<sup>\*</sup>deparra@wisc.edu

<sup>†</sup>velten@wisc.edu

energy by widening the pulse width. Clipping leads to significant photon loss, particularly under low peak power and high bandwidth conditions, and results in the greatest performance degradation for both compressive histograms and FRH. Preserving pulse energy by increasing the pulse width is more efficient at low peak power but suffers from reduced timing resolution and lower peak power for larger pulse widths. Conversely, our optimized illumination maximizes both average power ( $\Phi^{\text{sig}}$ ) and peak power ( $\Phi^{\text{max}}$ ) without violating hardware constraints. As a result, the optimized coding functions can outperform even FRH (for low-to-medium peak laser power) while achieving compression ratios of  $\geq 100\times$ . This suggests that preserving and clipping pulse energy are suboptimal, even for FRH, and that the illumination waveform must change to avoid performance loss. These findings show the importance of optimizing illumination and coding under strict hardware constraints. In the absence of such constraints, however, existing compressive histogram methods remain near-optimal.

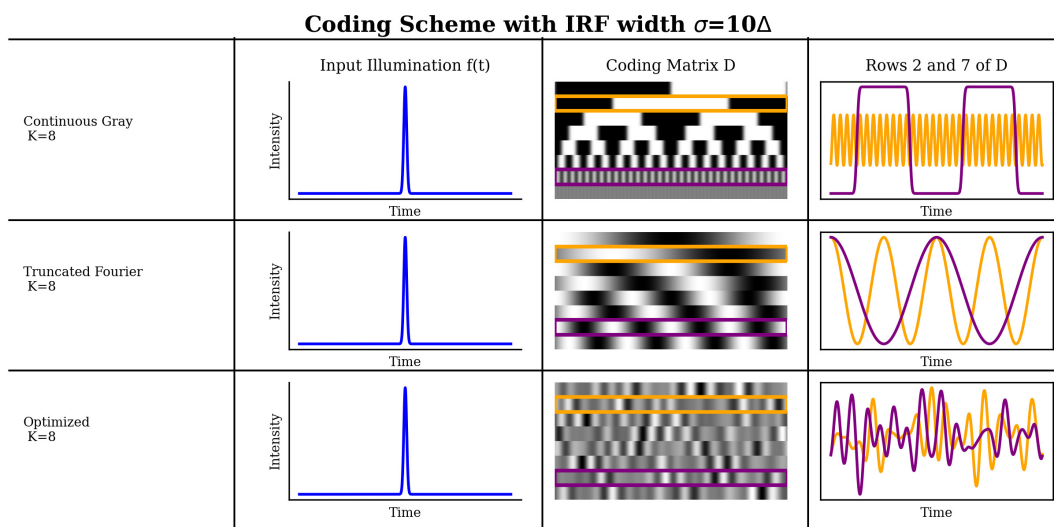


Figure 1: This figure plots the compressive histograms and the optimized coding schemes with band-limited illumination with IRF width  $\sigma = 10\Delta$

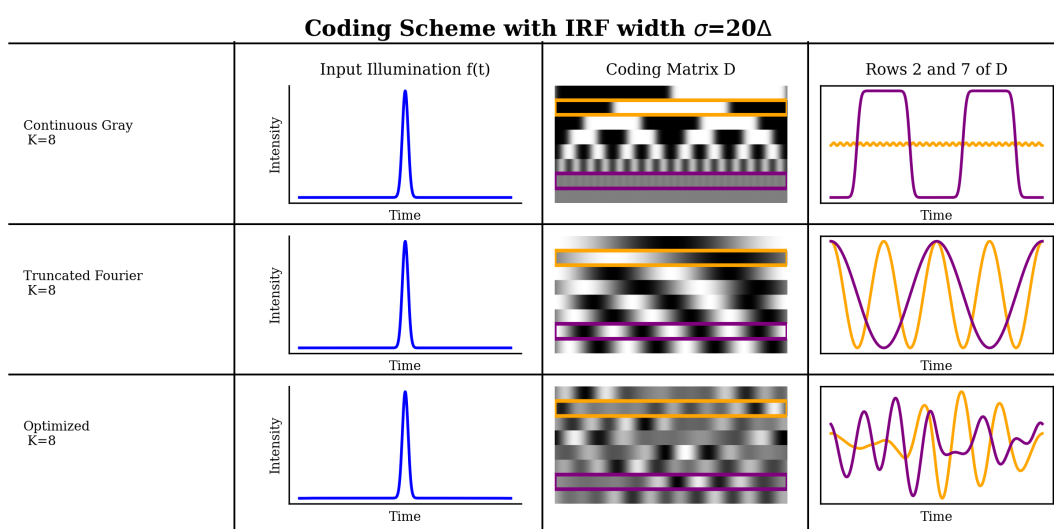


Figure 2: This figure plots the compressive histograms and the optimized coding schemes with band-limited illumination with IRF width  $\sigma = 20\Delta$

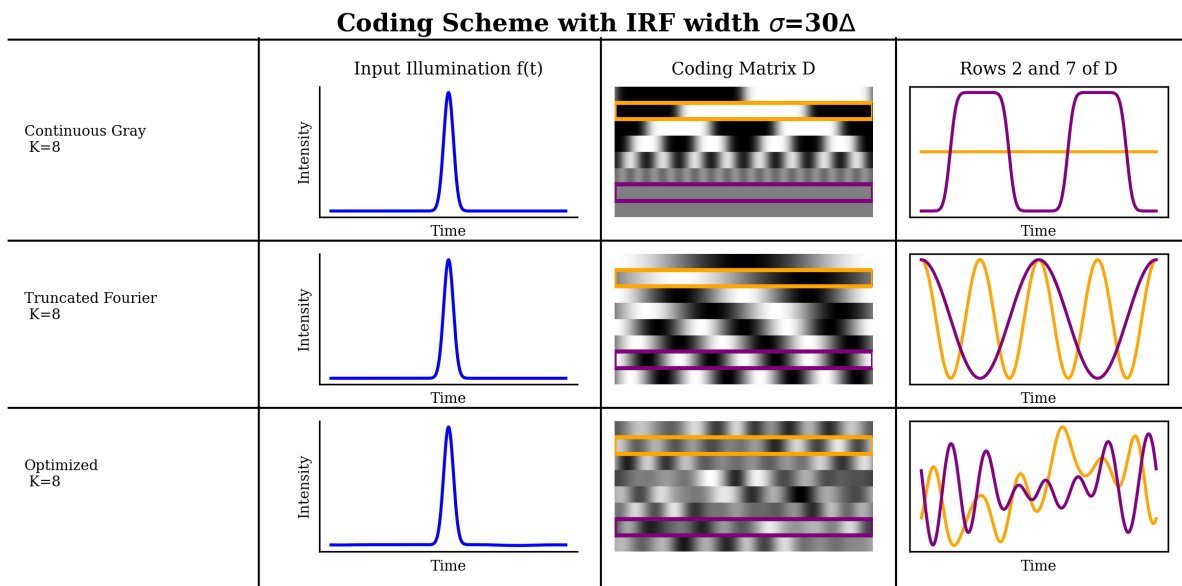


Figure 3: This figure plots the compressive histograms and the optimized coding schemes with band-limited illumination with IRF width  $\sigma = 30\Delta$

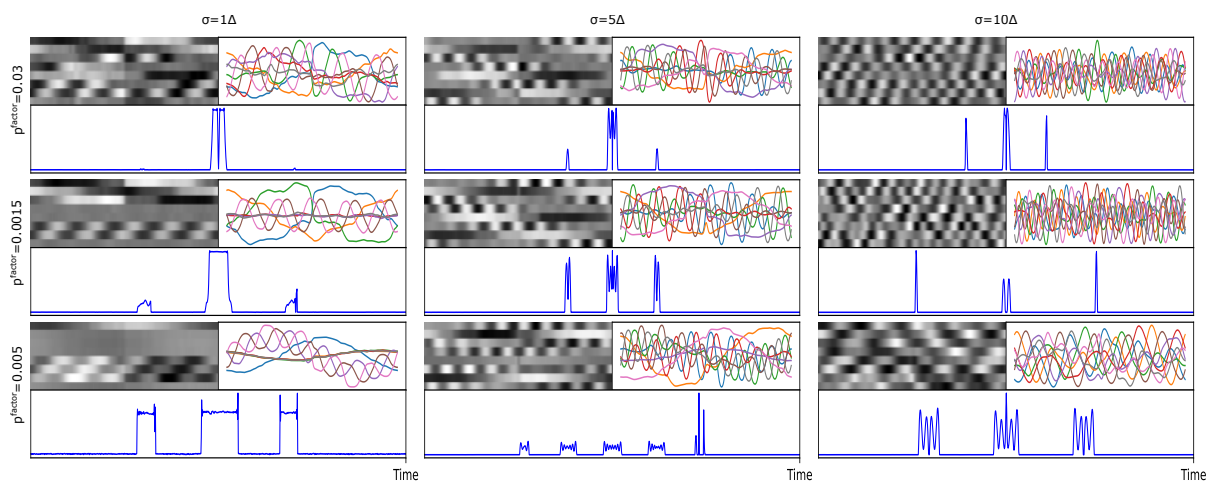


Figure 4: This figure plots the compressive histograms and the optimized coding schemes with peak-power limited illumination without filtering with the IRF.

## 2 Performance Evaluation

This section will detail the Monte Carlo Simulations’ performed and the performance metrics used to quantify a coding scheme’s efficacy. Supp. Sec. 9 presents additional results of Monte Carlo experiments, including error surface plots and contour plots.

### 2.1 Performance Metrics

A coding scheme should perform well across a wide range of photon counts, signal-to-background ratios (SBR), and depths. To evaluate the effectiveness of a coding scheme, we compute its Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), which have been shown to be reliable indicators of performance [2–4]. The MAE is defined as:

$$MAE = \frac{1}{n} \sum_{i=0}^n |t_d - \hat{t}_d| \quad (2)$$

MAE measures the average absolute error, providing insight into a code’s ability to reduce overall error and minimize small variations. The RMSE is defined as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=0}^n (t_d - \hat{t}_d)^2} \quad (3)$$

RMSE indicates how well a coding scheme minimizes large depth errors.

### 2.2 Monte Carlo Simulations

To perform our Monte Carlo simulations, we first fix a photon count  $\Phi^{\text{sig}} = \sum_{i=0}^{N-1} \Phi_i^{\text{sig}}$  and SBR (SBR =  $\frac{\Phi^{\text{sig}}}{N\Phi^{\text{bkg}}}$ ) level, for a given histogram length  $N$  and IRF  $h(t)$ . For each SNR level, the input illumination  $f(t)$  is filtered by  $h(t)$  to produce the output illumination  $s(t) = f(t) \otimes h(t)$ . We generate  $P$  equally spaced temporal shifts over the  $N$  time bins, and scale and offset  $s(t)$  according to the specified photon count and SBR. This gives us  $P$  ground truth depths  $t_d$ . For each  $s(t)$ , we simulate a histogram by drawing Poisson samples for each time bin. We then compute the predicted depth  $\hat{t}_d$  using ZNCC reconstruction with the coding matrix  $D$ . This process is repeated  $S$  times. Finally, we calculate MAE and RMSE using Equations 2 and 3, where the total number of samples is  $n = P \times S$ . We use  $N = 1024$ ,  $S = 5000$ , and  $D = 56$  for the main paper’s Monte Carlo simulations.

### 3 Analysis of Coding Functions with Indirect Reflections

In this section, we evaluate the performance of compressive histograms and our coding functions when subjected to sparse and dense indirect reflections. To simulate sparse indirect reflections, we added a second pulse (reflection) to the incident waveform with varying offsets and amplitudes. For dense indirect reflections, we added an exponential decay function to the tail of the histogram with varying levels of decay and amplitude.

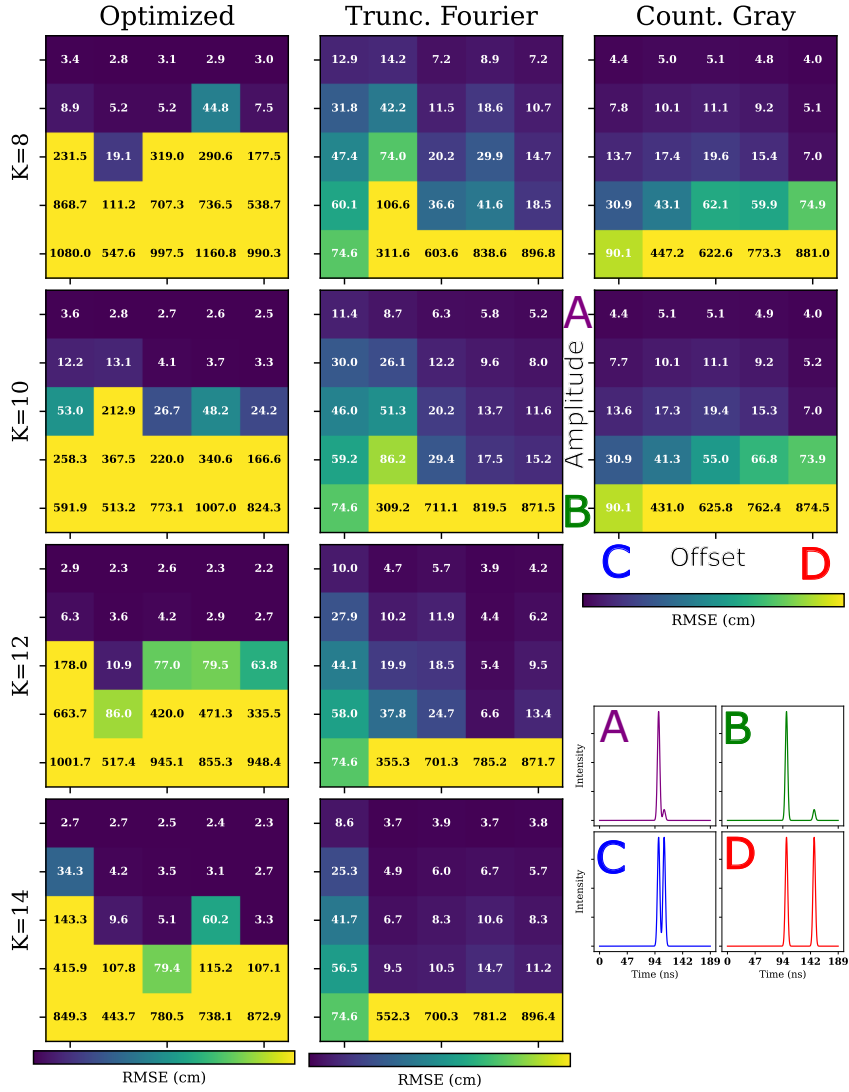
#### 3.1 Sparse Indirect Reflections

**Band-limited results:** Supp. Fig. 5 plots RMSE heatmaps for each of the band-limited compressive schemes (columns) and different  $K$  values (rows), evaluated under sparse indirect reflections with varying offsets and amplitudes of the second peak. Fourier-based codes suffer from the most performance degradation in pixels where the indirect reflection is close to the true pulse. Similar to the experimental results in the paper, truncated Fourier struggles to differentiate between the two peaks. On the other hand, it remains robust provided the second pulse is far enough from the true pulse. Gray-based codes are robust to varying amplitudes and offsets in the second pulse. However, in band-limited systems, this advantage is diminished due to the smoothing of high-frequency rows in Gray-based codes and is not improved by increasing  $K$ . The band-limited optimized coding functions are robust to small amplitudes and offsets in the second pulse. For larger, sparse indirect reflections in the incident waveform, these coding functions can produce large errors or outliers. We found that by increasing  $K$ , the band-limited coding functions saw the largest performance boost compared to other compressive codes and became increasingly resilient to sparse indirect reflections. We attribute the outliers to the optimization process, which assumes only direct reflections during training and struggles to fit incident functions that deviate significantly from the training data.

**Peak power-limited results:** Supp. Fig. 6 plots RMSE heatmaps for the peak power-limited ( $p^{\text{factor}} = 0.015$  for compressive schemes (cols) for different  $K$  values (rows) evaluated under the presence of sparse indirect reflection with varying offsets and amplitudes of the secondary peak. In these experiments, pulse energy is preserved by increasing pulse width. Therefore, pulse-based methods are not a function of IRF width, and we present the optimized coding functions with IRF widths  $\sigma = [5\Delta, 10\Delta]$ . Similar to the band-limited case, truncated Fourier performs well when the secondary peak is sufficiently far from the true peak, but its performance degrades as the peaks move closer together. Additionally, increasing  $K$  for truncated Fourier provides measurable improvements, albeit minor. Even with the increased pulse width due to preserving energy, Gray codes make the fewest errors with sparse indirect reflections, although Gray codes are fundamentally limited by bandwidth and see no improvement beyond  $K = 8$ . Due to the multiple peaks and only training on direct reflections, the optimized coding functions struggle with sparse indirect reflections. However, increasing  $K$  yields a significant performance improvement compared to compressive histograms. At  $K = 12$ , the optimized coding functions demonstrate moderate resilience to indirect reflections despite the presence of multiple peaks in the illumination. Opposite to truncated Fourier, we found that the optimized coding functions perform better when the reflection is close to the true depth and worse when it is far away. Among the configurations tested, the coding functions with IRF width  $\sigma = 5\Delta$  achieved the best overall performance.

For low peak power (Supp. Fig. 7 and  $p^{\text{factor}} = 0.005$ ), all coding methods produce large errors under the presence of sparse indirect reflections and fail to reliably decode depths. In this regime, limited peak laser power becomes a primary bottleneck, further exacerbated by sparse indirect reflections. Gray codes perform best under these constraints, despite their low bandwidth, though overall errors remain high. Notably, at low peak power, increasing  $K$  results in worse performance for truncated Fourier coding. This decline is due to the extremely low-frequency content of the widened pulse, which makes the additional Fourier components

ineffective. While increasing  $K$  improves performance for the optimized codes, they struggle most with sparse indirect reflections. This is due to the illumination having multiple peaks of the same amplitude (for  $p^{\text{factor}} = 0.005$ ), and the fact that the coding functions were trained exclusively on direct reflections during optimization. A possible solution is adding different scene responses in the model, such as those implemented in this analysis. This way, the model can adapt to incident functions that might be encountered in real-world scenarios. Regardless, all methods fail (including FRH) when the indirect reflection peak is higher than the true peak. Note that the illumination is held fixed for  $K = [10, 12]$  to evaluate its performance under indirect reflections with larger  $K$ .



**Figure 5: Evaluation of band-limited coding schemes in presence of sparse indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12, 14]$  (rows) when subjected to sparse indirect reflections. For each heatmap, the x-axis represents the temporal offset of the indirect reflection (second pulse) relative to the true pulse, while the y-axis denotes its amplitude. The bottom right figure plots example incident functions with sparse indirect reflections.

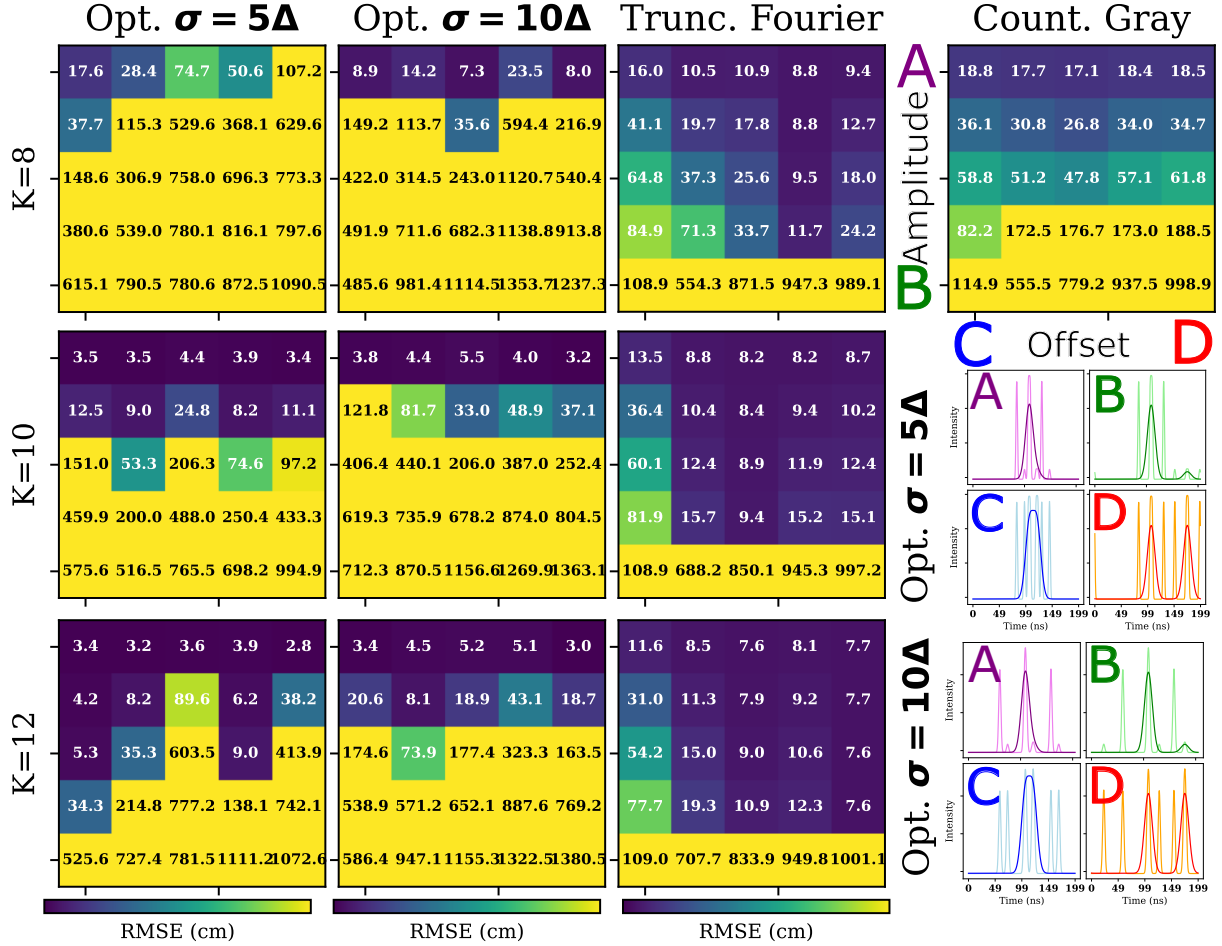


Figure 6: **Evaluation of peak power-limited (moderate peak power) coding schemes in presence of sparse indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12]$  (rows) when subjected to sparse indirect reflections and peak factor  $p^{\text{factor}} = 0.015$ . For each heatmap, the x-axis represents the temporal offset of the indirect reflection (second pulse) relative to the true pulse, while the y-axis denotes its amplitude. The two bottom right figure plots example incident functions with sparse indirect reflections for the optimized coding functions for IRF widths  $\sigma = [5\Delta, 10\Delta]$ .

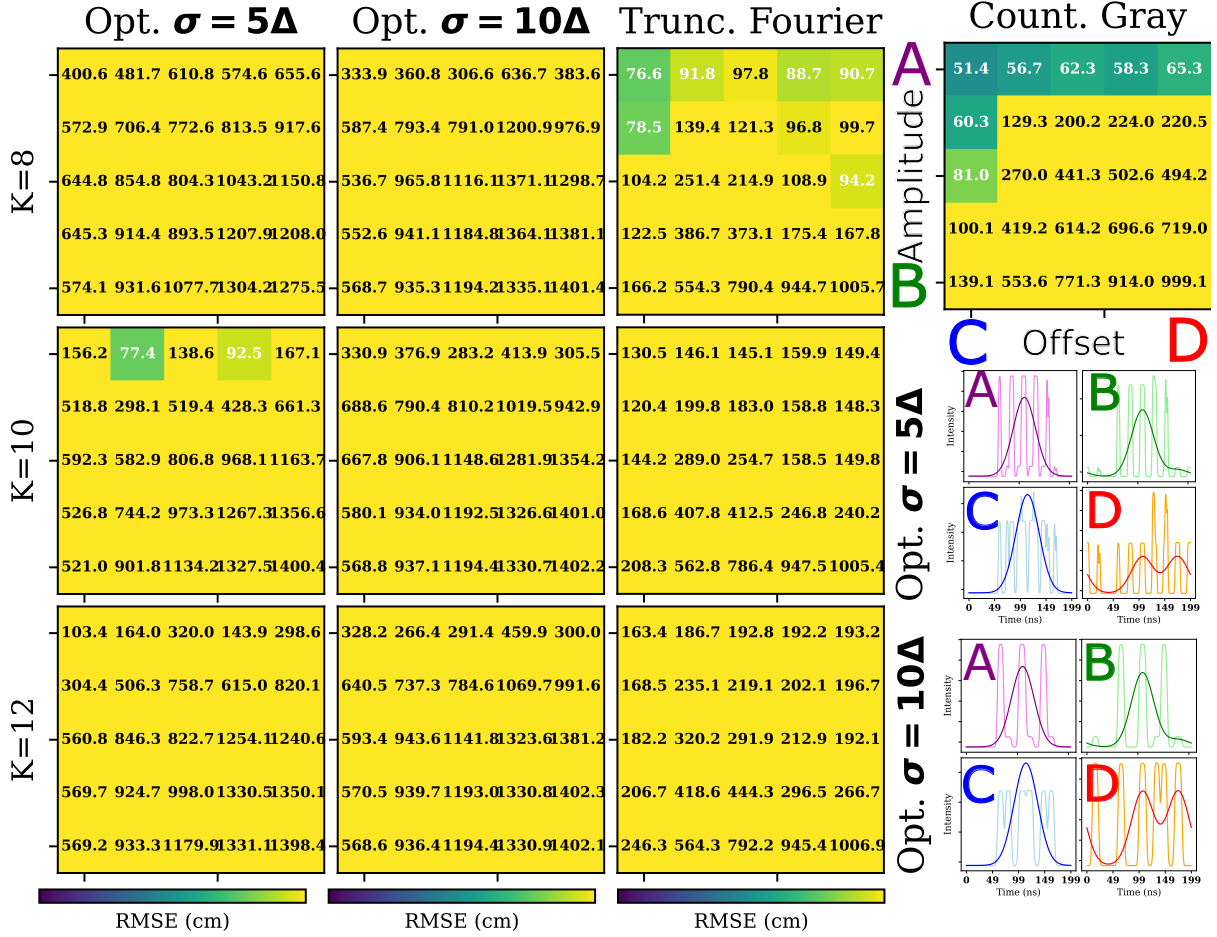


Figure 7: **Evaluation of peak power-limited (low peak power) coding schemes in presence of sparse indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12]$  (rows) when subjected to sparse indirect reflections and peak factor  $p^{\text{factor}} = 0.005$ . For each heatmap, the x-axis represents the temporal offset of the indirect reflection (second pulse) relative to the true pulse, while the y-axis denotes its amplitude. The two bottom right figure plots example incident functions with sparse indirect reflections for the optimized coding functions for IRF widths  $\sigma = [5\Delta, 10\Delta]$ .

### 3.2 Dense Indirect Reflections

Supp. Fig. 5 plots RMSE heatmaps for each of the band-limited compressive schemes (cols) for different  $K$  values (rows) evaluated under dense indirect reflection with varying decays and amplitudes of the exponential function (histogram tail). We find that both compressive histograms and the optimized coding functions are generally robust to dense indirect reflections, except for a few edge cases. This robustness is due to the presence of sufficiently high-frequency components in the coding matrix[4]. All coding schemes exhibit the most performance loss when the histogram tail has a large decay, though they perform well in other scenarios. Gray-based codes are effective in scenes with dense indirect reflections, despite the low bandwidth constraints. However, the band-limited optimized coding functions outperform both compressive histograms under dense indirect reflection.

Additionally, Supp. Fig. 6 and Fig. 10 present peak power-limited results for peak factors  $p^{\text{factor}} = 0.015$  and  $p^{\text{factor}} = 0.005$ , respectively. As observed in the sparse indirect reflection cases, at  $p^{\text{factor}} = 0.005$ , coding performance is primarily limited by the peak laser power rather than by the indirect reflections themselves. However, all coding schemes performed better under dense indirect reflections than sparse indirect reflections. For high and low peak power, Gray-based codes saw better performance than truncated Fourier, even when preserving pulse power, although they see no improvement beyond  $K = 8$ . Increasing  $K$  improves the performance of the optimized coding functions, while it leads to reduced performance for truncated Fourier. Notably, coding functions with IRF width  $\sigma = 5\Delta$  and  $K = [10, 12]$  perform significantly better under dense indirect reflections compared to those with IRF width  $\sigma = 10\Delta$ .

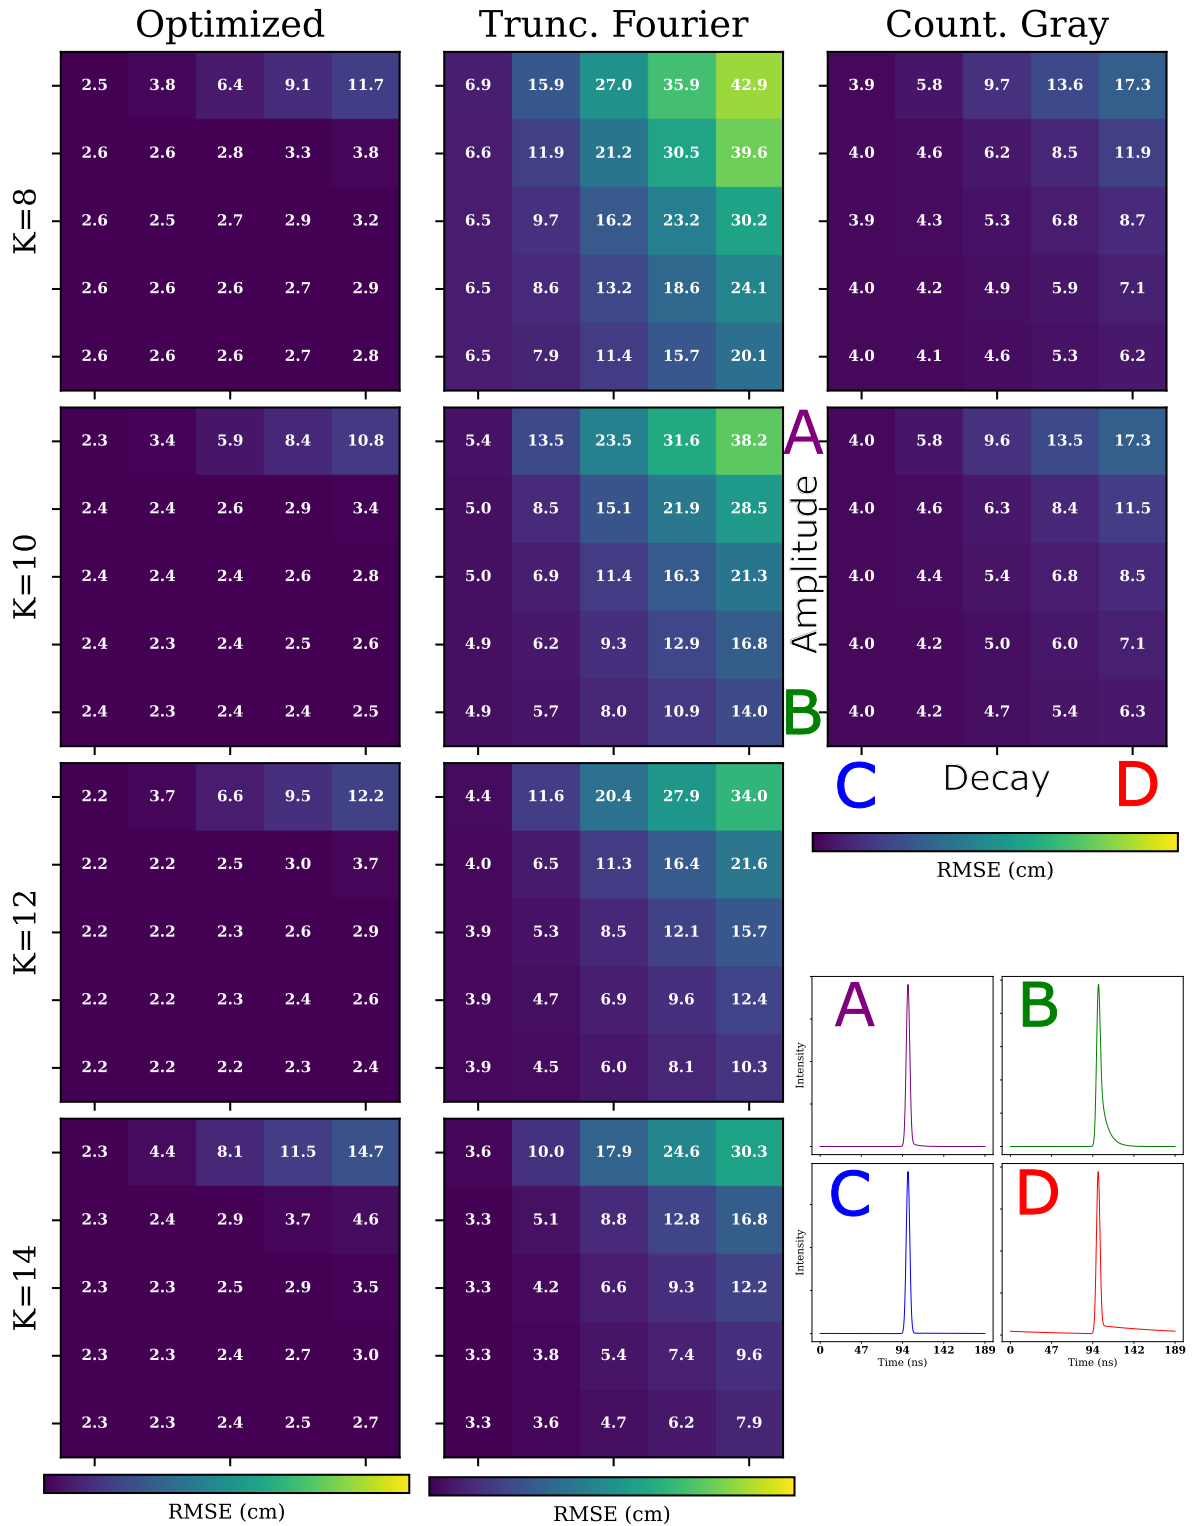


Figure 8: **Evaluation of band-limited coding schemes in presence of dense indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12, 14]$  (rows) when subjected to dense indirect reflections. For each heatmap, the x-axis represents the decay of the indirect reflection (exponential function), while the y-axis denotes its amplitude. The bottom right figure plots example incident functions with dense indirect reflections.

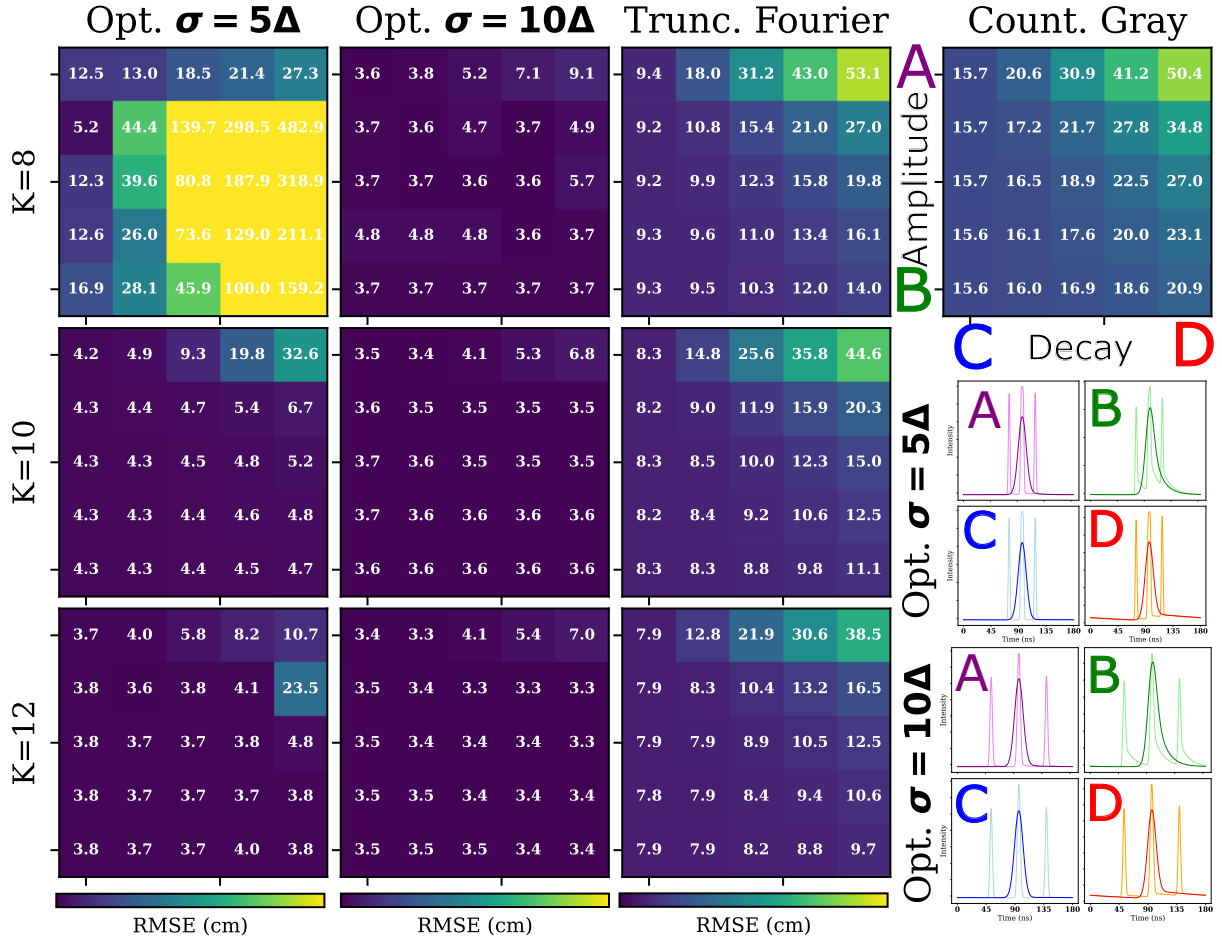


Figure 9: **Evaluation of peak power-limited (moderate peak power) coding schemes in presence of dense indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12]$  (rows) when subjected to dense indirect reflections and peak factor  $p^{\text{factor}} = 0.015$ . For each heatmap, the x-axis represents the decay of the indirect reflection (exponential function), while the y-axis denotes its amplitude. The two bottom right figure plots example incident functions with dense indirect reflections for the optimized coding functions for IRF widths  $\sigma = [5\Delta, 10\Delta]$ .

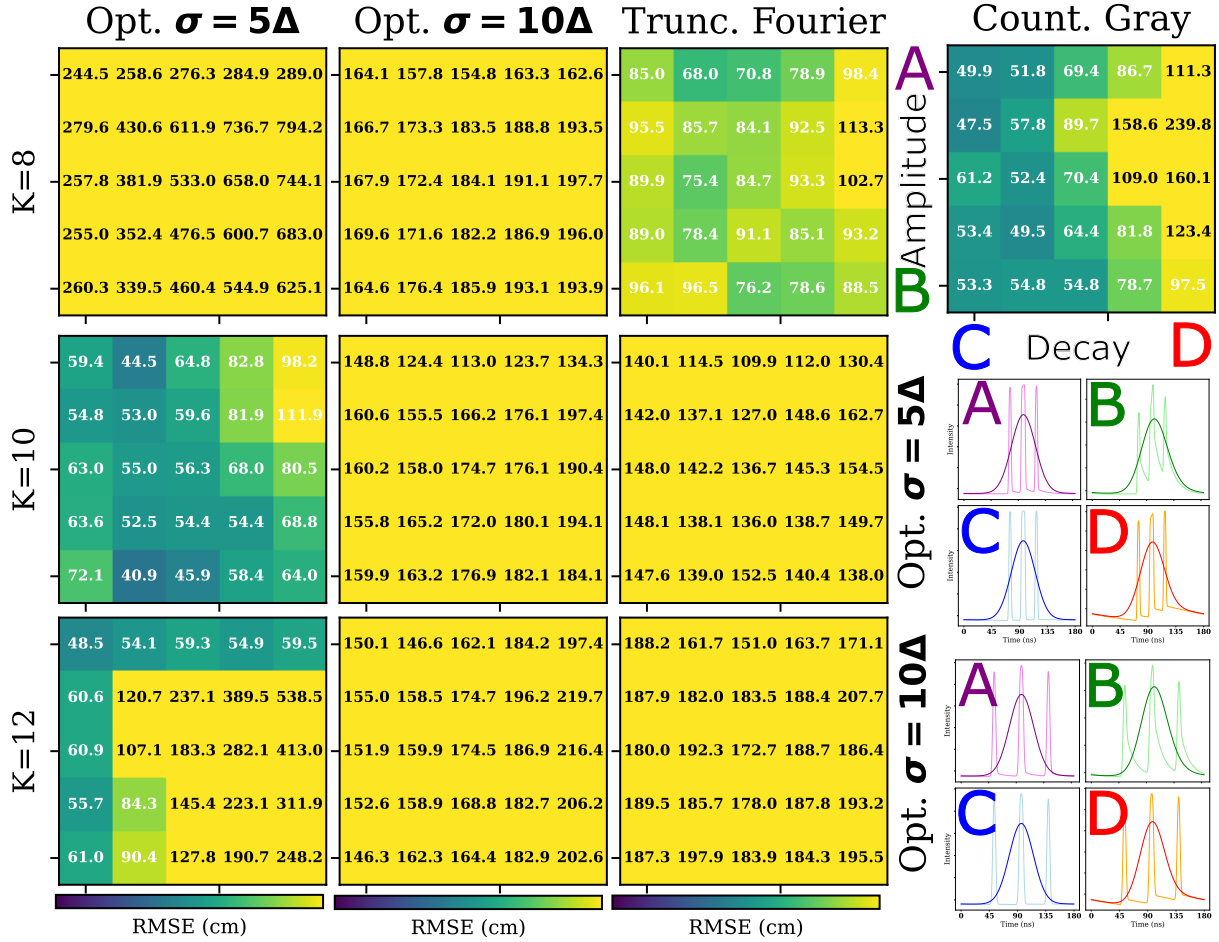


Figure 10: **Evaluation of peak power-limited (low peak power) coding schemes in presence of dense indirect reflections:** This figure presents RMSE heatmaps for compressive histograms and the proposed coding functions (cols), evaluated at  $K = [8, 10, 12]$  (rows) when subjected to dense indirect reflections and peak factor  $p^{\text{factor}} = 0.005$ . For each heatmap, the x-axis represents the decay of the indirect reflection (exponential function), while the y-axis denotes its amplitude. The two bottom right figure plots example incident functions with dense indirect reflections for the optimized coding functions for IRF widths  $\sigma = [5\Delta, 10\Delta]$ .

## 4 Additional Real-World Results

### 4.1 Experimental Hardware Scanning-based System

In this section, we give additional results using a real-world scanning-based hardware system obtained from [5]. Supp. Fig. 11 plots the resulting depth map of a Deer face evaluated using compressive histograms and our proposed coding functions ( $K = 10$ ). This scan includes pixels with dense indirect reflections, as indicated by the long tail in the measured histogram. For these pixels, all the coding schemes exhibit higher RMSE values. However, we find that even with indirect reflections, the optimized coding functions achieve the closest accuracy to FRH.

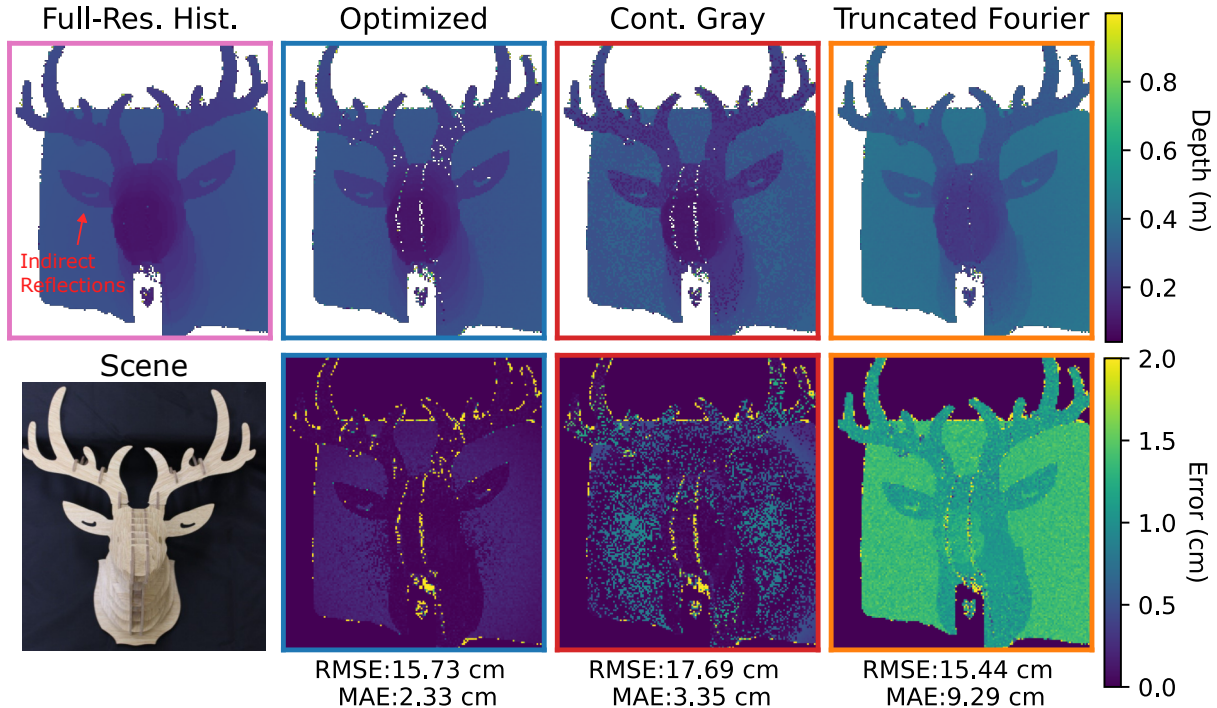


Figure 11: **Depth map results from hardware system.** Real-world scans of a Deer face with the different compressive codes. RMSE and MAE are calculated from the Full-res. histogram as baseline. White pixels indicate outliers in depth map results.

### 4.2 Simulated Flash-illumination System

In this section, we give additional flash-illumination simulation results not shown in the main paper obtained from [6]. 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 then scaling and offsetting the results based on the desired photon count and SBR. The following scenes include both dense and sparse indirect reflections.

### 4.3 Band-limited flash-illumination results

Supp. Fig. 12 and Fig. 13 present the depth maps and corresponding error for the various coding schemes, evaluated with an IRF width of  $\sigma = 30\Delta$  and infinite peak power, across different values of  $K$ . Each scene was simulated with 1000 mean photons and an SBR of 0.1. Truncated Fourier exhibits moderate performance improvements with increasing  $K$ . However, it continues to suffer from systematic errors in both sparse and dense indirect reflections when the reflection is close to the true peak. Gray-based coding shows no improvement beyond  $K = 8$ , as

its higher-frequency rows are smoothed out by the IRF and do not contribute to depth decoding. In contrast, the optimized coding functions demonstrate the greatest performance gains with increasing  $K$ . At higher  $K$ , they achieve near-zero error across most pixels, with only occasional outliers caused by indirect reflections.

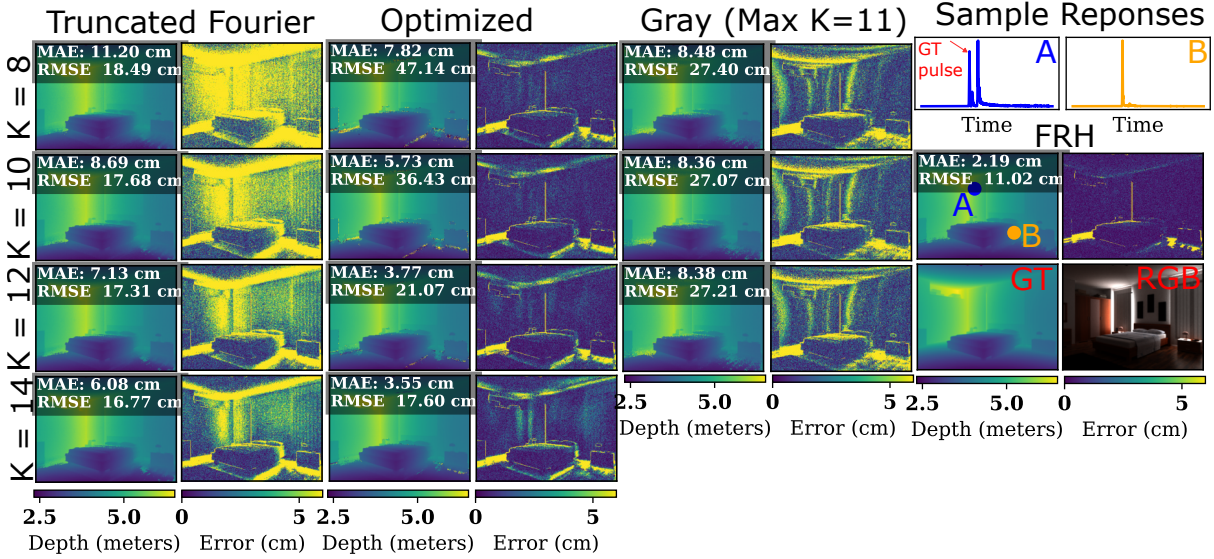


Figure 12: **Flash-illumination simulation result with band-limited coding functions:** This figure shows the resulting depth maps and error maps for each compressive code (cols) across different encoding sizes  $K = [8, 10, 12, 14]$  (rows). For each code and  $K$ , we display the depth map, corresponding error map, and report the RMSE and MAE. The final row includes visualizations of two pixels with indirect reflections, the ground truth depth and RGB images, and the FRH.

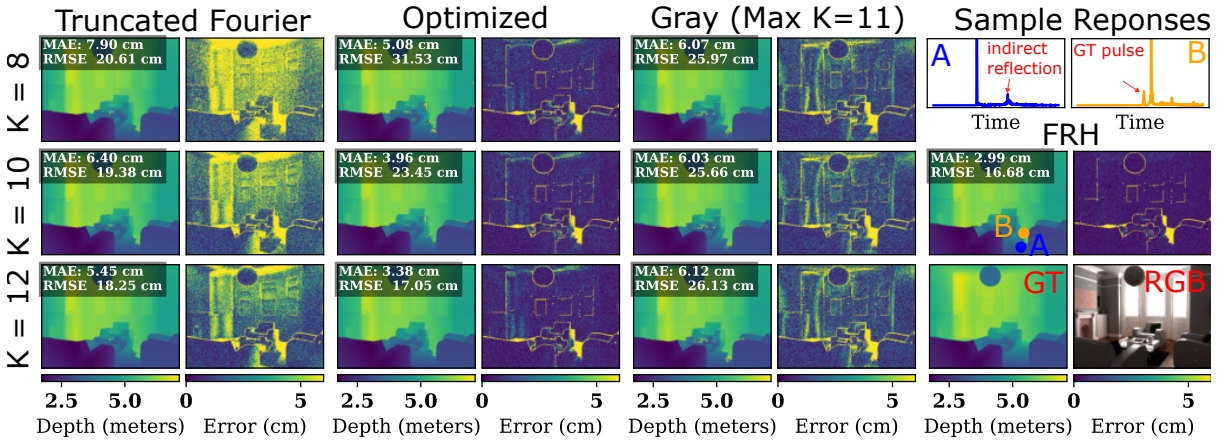


Figure 13: **Flash-illumination simulation result with band-limited coding functions:** This figure shows the resulting depth maps and error maps for each compressive code (cols) across different encoding sizes  $K = [8, 10, 12]$  (rows).

#### 4.4 Peak Power-limited Flash-illumination Results

Supp. Fig. 14 plots depth maps and corresponding error for the optimized coding functions evaluated at IRF widths  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  and peak factors  $p^{\text{factor}} = [0.030, 0.015, 0.005]$ , where  $K = 8$ . Each scene was simulated with 1000 mean photons and an SBR of 0.1. The peak power-limited, infinite bandwidth case ( $\sigma = 1\Delta$ ) performs significantly worse under both sparse and

dense indirect reflections. This arises from a combination of no smoothing of high-frequency components in the scene response and the low-frequency content of the coding matrices. In contrast, the band-limited coding functions (and peak power-limited) suppress these high frequencies and thus perform significantly better.

Supp. Fig. 15 and Fig. 16 present depth maps and corresponding errors for different coding schemes evaluated at  $K = [8, 12]$  and peak factors  $p^{\text{factor}} = [0.005, 0.015]$ . In these experiments, pulse energy is preserved by increasing pulse width. Therefore, bandwidth is no longer a limiting hardware constraint for pulse-based methods, and we present the optimized coding functions with IRF widths  $\sigma = [5\Delta, 10\Delta]$ . At low peak laser power ( $p^{\text{factor}} = 0.005$ ), performance is primarily limited by peak power rather than the scene response. In this regime, the optimized coding functions outperform all baselines (including compressive histograms and FRH), achieving the lowest MAE by a significant margin. Pulse-based methods exhibit poor performance, even when preserving pulse energy, due to the reduced temporal resolution and diminished peak power caused by increasing the pulse width. Among all schemes, increasing  $K$  benefits the optimized coding functions the most. In contrast, compressive histograms gain nothing from higher  $K$  due to the low-frequency content of the wide pulse. In particular, truncated Fourier shows performance loss at  $K = 12$ , as the additional frequency components of the coding matrix contribute no useful information for low-frequency pulse waveforms.

As peak laser power increases, the scene response becomes a more limiting factor ( $p^{\text{factor}} = 0.015$ ). Similar to the band-limited case, peak power-limited coding functions struggle with sparse indirect reflections that contain strong secondary peaks, often resulting in outliers, more so than pulse-based methods. Increasing  $K$  helps mitigate this effect, but only with small, sparse/dense indirect reflections. In such cases, truncated Fourier is an efficient coding scheme under higher peak power and when sparse indirect reflections are present. In particular, in the bedroom scene (Fig. 16) with many indirect reflections, truncated Fourier performs the best. These results highlight the importance of selecting appropriate coding schemes based on both hardware constraints and scene responses. Note that the illumination is held fixed for  $K = 12$  to evaluate its performance under indirect reflections with larger  $K$ .

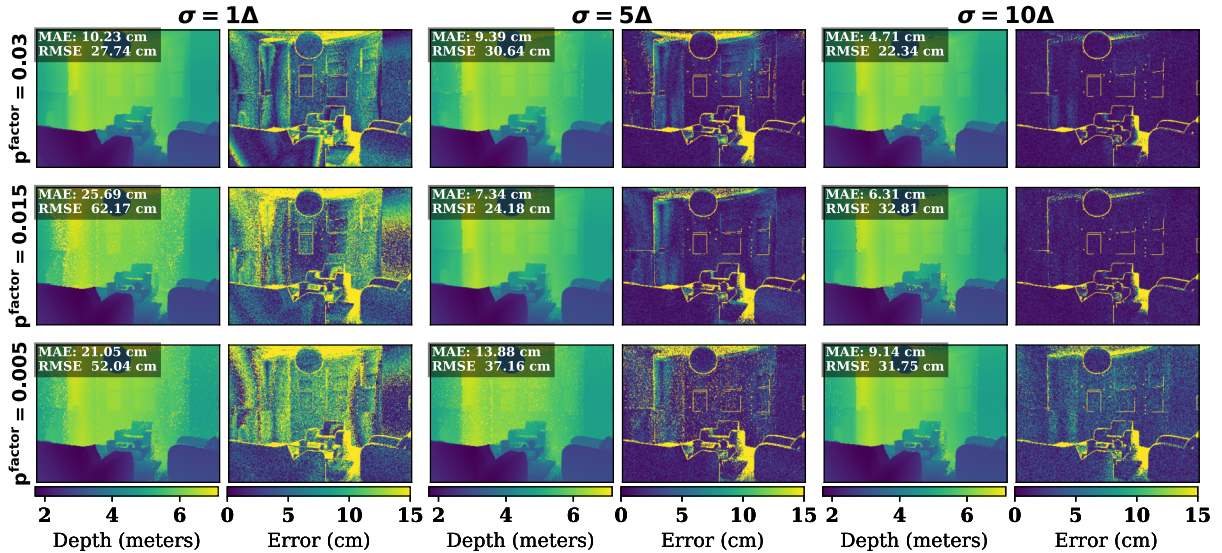


Figure 14: **Peak power-limited coding functions with flash-illumination system:** This figure plots the depth map and error for the optimized coding functions for IRF widths  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  and peak factors  $p^{\text{factor}} = [0.030, 0.015, 0.005]$ , where  $K = 8$ .

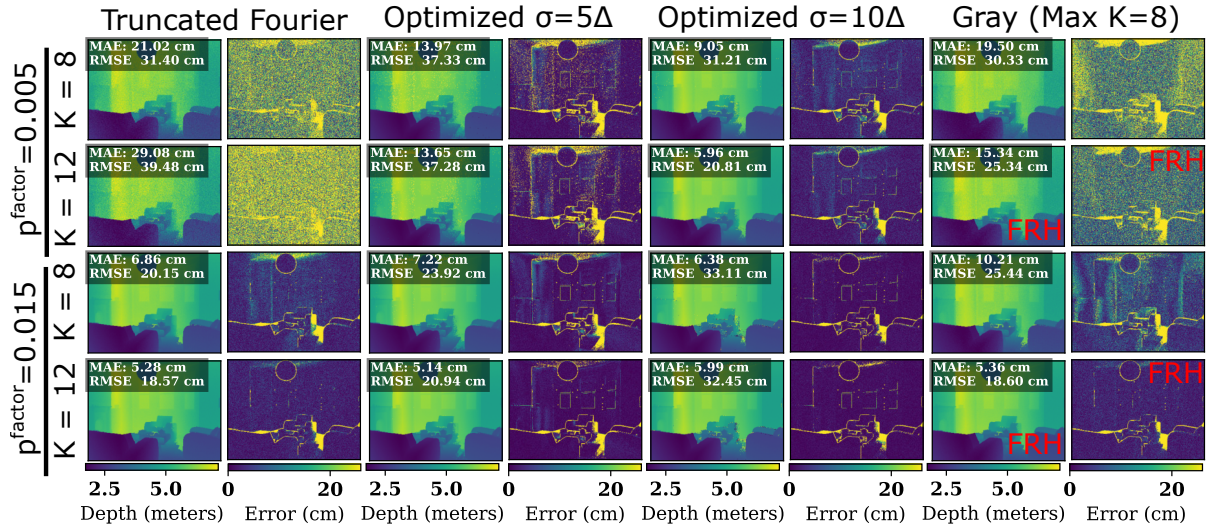


Figure 15: **Peak power-limited flash-illumination system with living room scene** : This figure shows the depth maps and corresponding errors for each coding scheme (cols) under peak factors of  $p^{\text{factor}} = [0.005, 0.015]$  when preserving pulse energy. Results are shown for  $K = [8, 12]$  (rows), with Gray codes fixed at  $K = 8$ . The bottom-left panel shows the FRH depth map and error.

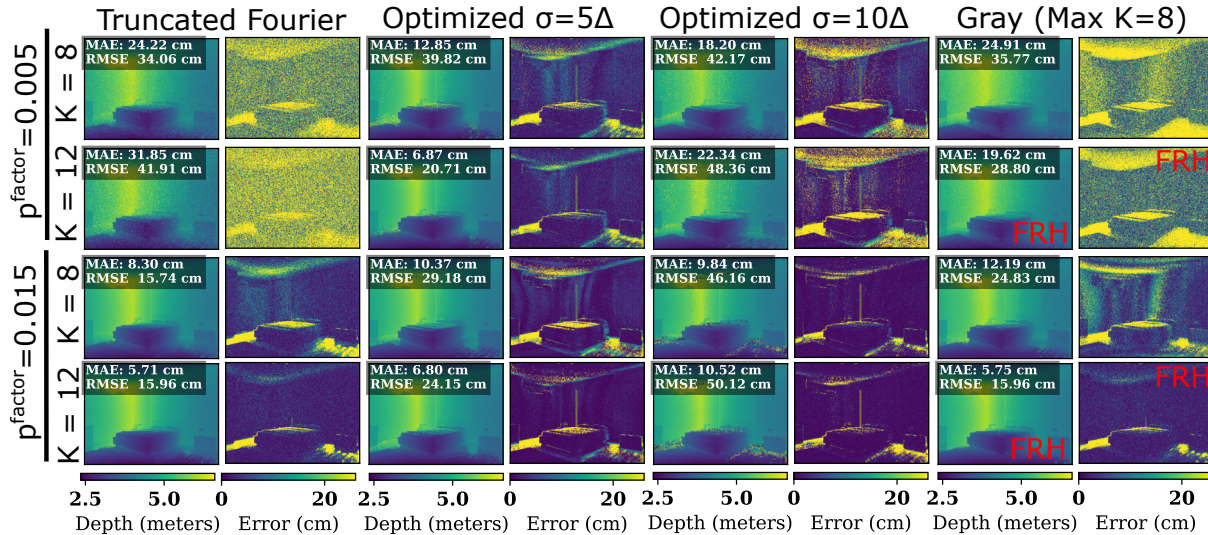


Figure 16: **Peak power-limited flash-illumination system with bedroom scene** : This figure shows the depth maps and corresponding errors for each coding scheme (cols) under peak factors of  $p^{\text{factor}} = [0.005, 0.015]$  when preserving pulse energy. Results are shown for  $K = [8, 12]$  (rows), with Gray codes fixed at  $K = 8$ . The bottom-left panel shows the FRH depth map and error.

## 5 Memory Analysis of Coding Functions

In this section, we analyze the memory overhead associated with implementing the optimized coding functions on-chip. Ideally, a coding matrix should have minimal bit-depth (1–8 bits), such as Gray codes, and support a memory-efficient representation, such as Fourier-based codes [7, 8]. A straightforward approach to in-pixel implementation involves storing the full coding matrix directly on-chip. However, unless shared across neighboring pixels, this approach incurs significant in-sensor memory overhead. Thus, a memory-efficient representation of our proposed coding functions is required for deployment. To evaluate this, we present the RMSE of the optimized coding functions when quantized to between 1 and 64 bits and reconstructed using 10 to 90 Fourier coefficients.

**Bit-depth Analysis:** Supp. Fig. 17 shows the resulting RMSE of the different coding schemes when quantized to 1 to 64 bits, evaluated at IRF widths  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  and infinite peak power. Additionally, Supp. Fig. 18 shows the same results but for IRF widths  $\sigma = [10\Delta, 20\Delta, 30\Delta]$  and peak factors  $p^{\text{factor}} = [0.030, 0.015, 0.005]$  (clipped pulse energy results). We find that the optimized coding functions show negligible performance loss when quantized to 4 bits for all band-limited and peak power-limited coding matrices, regardless of SBR. Furthermore, at high SBR and low bandwidth, this requirement is reduced to 2 bits with no performance loss.

**Fourier Analysis:** Supp. Fig. 19 shows the resulting RMSE of the different coding schemes when reconstructing  $D$  with 10 to 90 Fourier coefficients, evaluated at IRF widths  $\sigma = [10\Delta, 20\Delta, 30\Delta]$  and infinite peak power. Additionally, Supp. Fig. 20 shows the same results but for IRF widths  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  and peak factors  $p^{\text{factor}} = [0.030, 0.015, 0.005]$  (clipped pulse energy results). We find that most band-limited and peak power-limited coding matrices can be represented with fewer than 30 Fourier coefficients with negligible performance loss compared to a full representation. The peak power-limited coding functions further reduce the number of Fourier coefficients required for accurate representation, needing fewer than 15 coefficients for most configurations. For the high-bandwidth coding functions, up to 90 Fourier coefficients may be required at low SBR.

Our results demonstrate that the optimized coding functions can be quantized to as few as 4-bits and represented with 30 Fourier coefficients while still outperforming compressive histograms. Gutierrez et al. [8] recently showed an on-chip implementation of 8-bit Fourier codes on the UltraPhase Chip [7]. Taken together, these findings suggest that our coding functions have a memory-efficient implementation for online photon data compression.




Figure 17: **Bit depth analysis of band-limited coding functions:** This figure shows the RMSE as a function of bit depth for each compressive code. Each column corresponds to a different bandwidth,  $\sigma = [10\Delta, 20\Delta, 30\Delta]$ , and for each bandwidth, RMSE is reported under both high and low SBR conditions. Our optimized coding functions exhibit negligible performance degradation at bit depths of 4 bits or higher, regardless of SBR.




Figure 18: **Bit depth analysis of peak power-limited coding functions:** This figure shows the RMSE as a function of bit depth for each compressive code at peak powers of  $p^{\text{factor}} = [0.05, 0.015, 0.030]$  and bandwidths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$ . For each bandwidth and peak power, RMSE is reported under both high and low SBR conditions.




Figure 19: **Fourier analysis of band-limited coding functions:** This figure shows the RMSE as a function of the number of Fourier coefficients to reconstruct the coding matrix. Each column corresponds to a different bandwidth,  $\sigma = [10\Delta, 20\Delta, 30\Delta]$ , and for each bandwidth, RMSE is reported under both high and low SBR conditions. At high SBR, our optimized coding functions can be presented with fewer than 30 Fourier coefficients with negligible performance degradation. At low SBR, it may take up to 100 Fourier coefficients for the  $\sigma = 10\Delta$  coding functions to be efficiently represented.

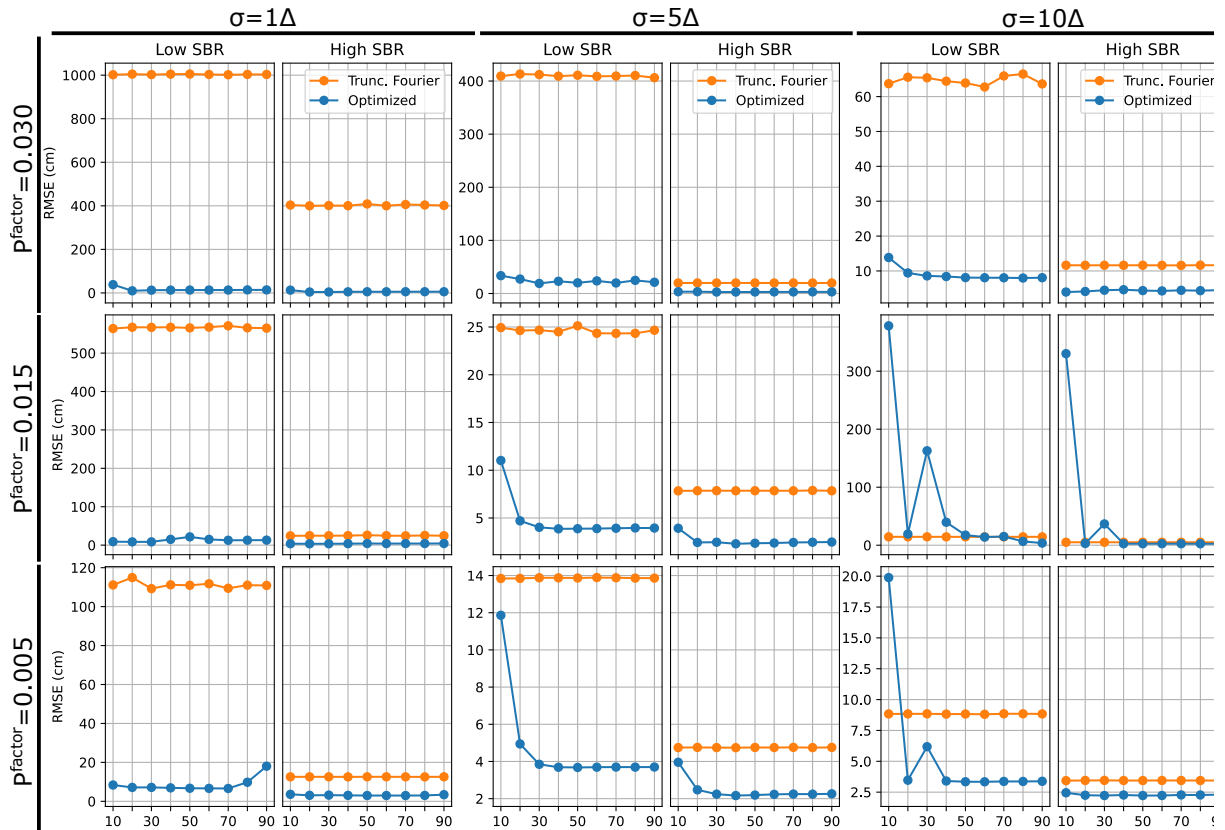


Figure 20: **Fourier analysis of peak power-limited coding functions:** This figure shows the RMSE as a function of the number of Fourier coefficients to reconstruct the coding matrix at peak powers of  $p^{\text{factor}} = [0.05, 0.015, 0.030]$  and bandwidths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$ . For each bandwidth and peak power, RMSE is reported under both high and low SBR conditions. Our optimized coding functions exhibit negligible performance degradation with fewer than 30 Fourier coefficients, regardless of SBR

## 6 Additional $K$ Coding Function Results

In this section, we evaluate the performance of the coding functions for  $K = 4$  and  $K > 8$ .

Supp. Fig. 21 shows RMSE (top row) and MAE (bottom row) from Monte Carlo simulations using band-limited, infinite peak power coding functions for  $K = [10, 12, 14]$ . Supp. Fig. 22 presents the same errors but for band-limited and peak power-limited coding functions when preserving pulse energy. Gray-based codes are limited to  $K = 10$  since the maximum number of encodings is  $K = \log_2(N)$ . We observe that the optimized coding functions show the largest performance gains with increasing  $K$  compared to compressive histograms. Truncated Fourier sees minimal gain in band-limited systems and no gain in peak power-limited systems due to increased pulse width. Furthermore, the higher RMSE values at low photon counts are mitigated by increasing  $K$ , as shown in the figure. Note that the illumination is held fixed for  $K = [10, 12, 14]$ , and only the coding matrix is optimized.

Supp. Fig. 23 shows depth estimation results for each of the compressive histograms, along with our optimized coding functions at  $K = 4$ . While our learned coding improves upon Truncated Fourier performance, Continuous Gray outperformed the other compressive codes in this setting. This contrasts with the  $K = 8$  case, where our optimized coding achieved the best performance. The strong performance of Gray codes at  $K = 4$  can be attributed to their efficiency at low dimensions. However, increasing to  $K = 8$  leads to significant performance improvements across all coding schemes, with only a minimal decrease in the compression rate.

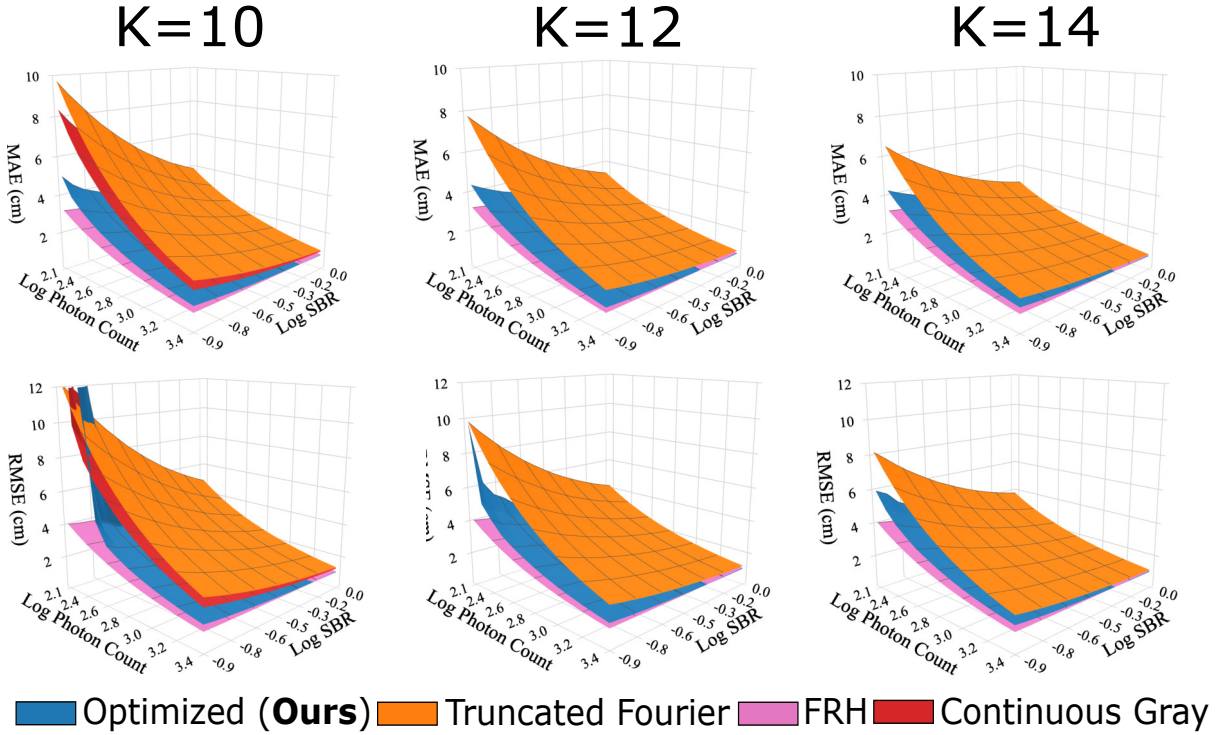


Figure 21: **Simulated depth map results for  $K > 8$  (band-limited coding functions).** This figure plots the MAE (top row) and RMSE (bottom row) results from Monte Carlo experiments using the compressive codes for  $K = [10, 12, 14]$  ( $K = 10$  for Gray codes), IRF width  $\sigma = 20\Delta$ , and infinite peak power. The optimized coding functions see the most performance improvement for increasing  $K$  compared to Truncated Fourier.

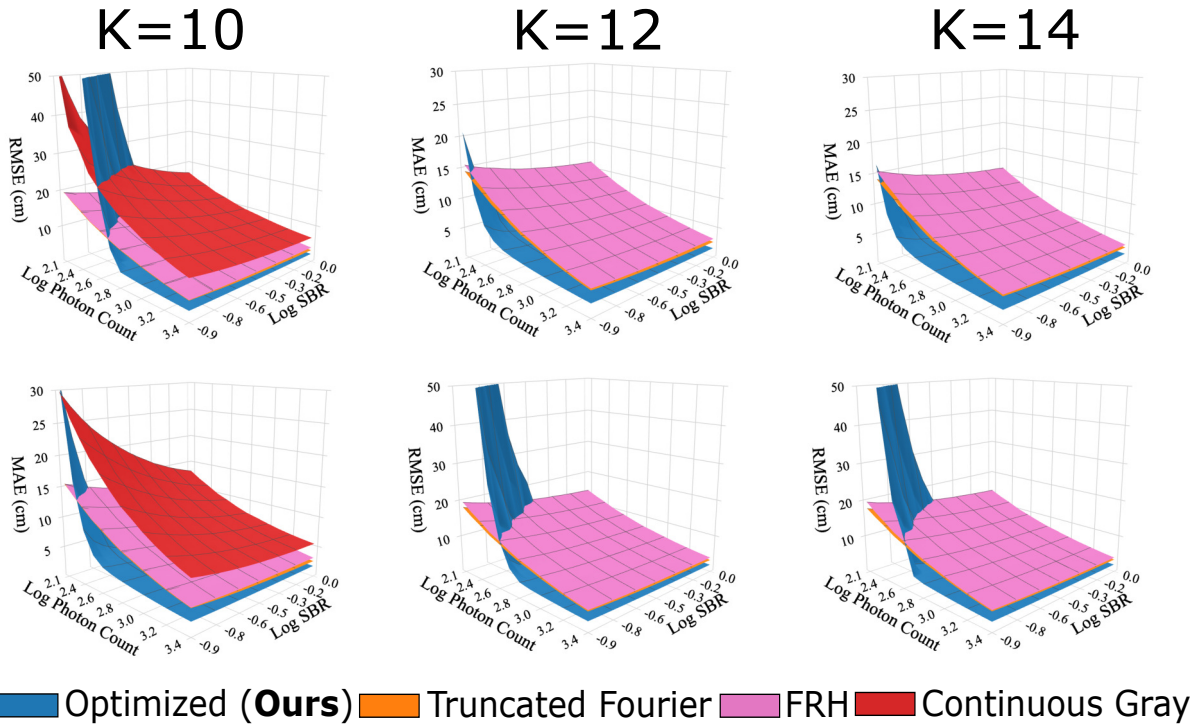


Figure 22: **Simulated depth map results for  $K > 8$  (moderate peak power-limited coding functions)**. This figure plots the MAE (top row) and RMSE (bottom row) results from Monte Carl experiments using the compressive codes for  $K = [10, 12, 14]$  ( $K = 10$  for Gray codes),  $\sigma = 10\Delta$  and peak factor  $p^{\text{factor}} = 0.015$ . Here, we give the error when preserving pulse energy for the pulse-based methods. The optimized coding functions see the most performance improvement for increasing  $K$ , whereas truncated Fourier coding sees no gain.




Figure 23: **Simulated depth map results for  $K = 4$ .** This figure plots a simulated depth map obtained from our optimized coding functions ( $K = 4$ ), Truncated Fourier ( $K = 4$ ), Gray-based codes ( $K = 4$ ), and full-resolution histograms ( $N = 1024$ ) with an IRF width  $\sigma = 20\Delta$  and infinite peak power. Top plots depth map with high SNR ( $SBR = 1.0$ ) and bottom plots depth map with low SNR ( $SBR = 0.1$ ). The RMSE and MAE are given for each respective coding. Black pixels indicate outliers.

## 7 Additional Implementation Details

This section outlines additional implementation of the constrained optimization described in the paper. To solve the optimization, we implemented a neural network trained via gradient descent. 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 follows:

$$\begin{aligned} \min_{f, D} \quad & l((t_{d_j})_{j=0}^{J-1}, (\hat{t}_{d_j})_{j=0}^{J-1}) \\ \text{subject to} \quad & (f \otimes h)_i = s_i \\ & f_i \leq \Phi^{\max}, i = 0, \dots, N-1 \end{aligned} \tag{4}$$

where  $l$  is the loss function between  $J$  true depths  $(t_{d_j})_{j=0}^{J-1}$  and predicted depths  $(\hat{t}_{d_j})_{j=0}^{J-1}$ .  $\Phi^{\max}$  is the peak photon count, expressed as  $\Phi^{\max} = p^{\text{factor}} * \Phi^{\text{sig}}$ , where  $\Phi^{\text{sig}} = \sum_i^{N-1} \Phi_i^{\text{sig}}$ . Bandwidth and IRF constraints are addressed in Line 2 of 4, whereas Line 3 contains the peak power limitations. To train, we create a dataset of labels that consists of ground truth depths,  $G \times G$  photon counts, and  $G \times G$  signal-to-background ratios. To ensure that we are covering the entire depth range, we create  $N \times R$  true depth labels for each photon count and SBR level, where  $R$  is the number of sweeps over the depth range. In the optimization,  $J$  is the total number of labels in the dataset, which equals  $J = G \times G \times N \times R$ .

Since the illumination is a learnable parameter, we simulate histogram formation during the forward pass. First, the illumination is filtered using a 1D convolutional layer to model the IRF and passed through a ReLU to ensure non-negativity. The filtered illumination is then scaled based on the desired photon count and SBR level. If peak power is limited, the illumination is clamped accordingly and re-filtered to suppress high-frequency artifacts introduced by clipping. Next, Gaussian noise with variance equal to the mean is added to the illumination, which approximates Poisson noise. We then compute the coded values from the simulated histogram using a 1D convolutional layer with kernel size 1 and shape  $K \times N$ , where the weights are the coding matrix  $D$ , which is equivalent to a fully-connected linear layer. We add TV regularization to the convolution weights to promote smooth functions. Finally, we estimate depths using the ZNCC reconstruction implemented with a soft arg max. In training, we set  $G = 10$ ,  $N \times R = 15,000$ . The total number of training labels for  $K = 8$  simulated optimized coding functions is  $J = 1,500,000$  and  $J = 3,000,000$  for the real-world coding functions.

### 7.1 Loss Functions

We tested three different loss functions for  $l$ . They are the L1 loss (mean depth error), RMSE loss, and Charbonnier loss [9]. The charbonnier loss is defined as:

$$l(t_d, \hat{t}_d) = \sqrt{||t_d - \hat{t}_d||^2 + \epsilon^2} \tag{5}$$

where  $\epsilon$  balances the tradeoff between the L1 loss and L2 loss. For testing we set  $\epsilon = 1e - 8$ . Supp. Fig. 25 shows the MAE for different loss functions used to train the optimized coding functions. Among the three, the main results in our paper are based on the L1 loss, which we found to perform best overall. Despite training with the L1 loss, our method still achieves lower RMSE compared to the compressive histogram baselines for photon counts above 150 (as shown in the main paper).

### Coding Scheme with IRF width $\sigma=30\Delta$

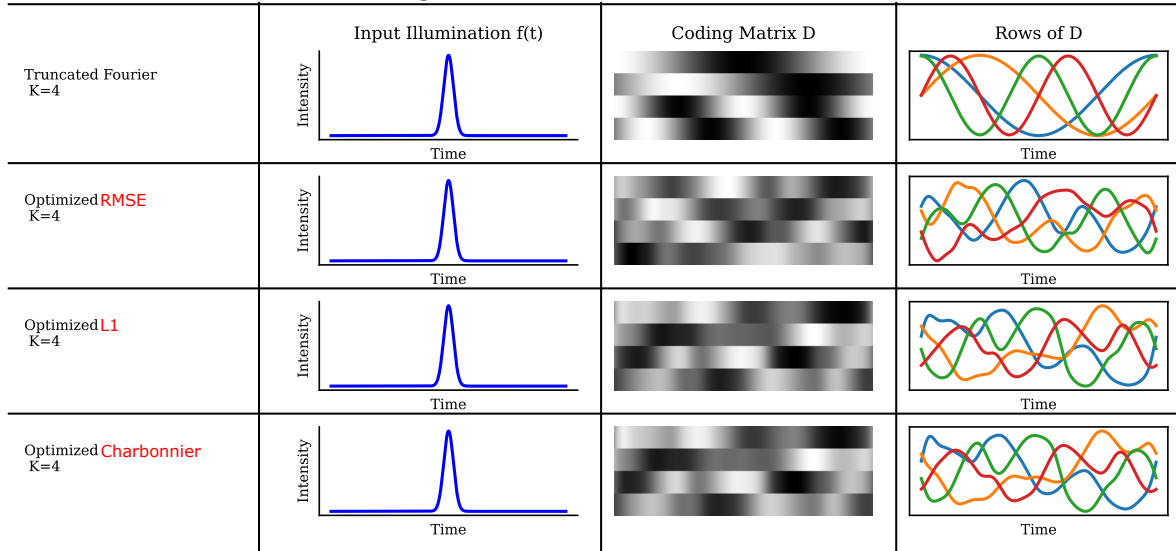


Figure 24: This figure plots the optimized coding functions for the RMSE loss, L1 Loss, and Charbonnier Loss.

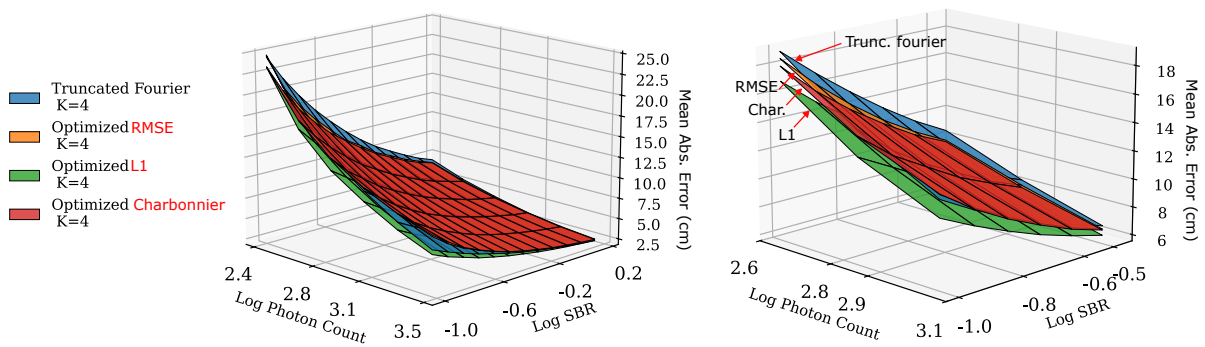


Figure 25: This figure plots the mean absolute error for the RMSE loss, L1 Loss, and Charbonnier Loss optimized coding functions. The L1 loss achieves the lowest error amount the loss functions.

## 8 Low Photon Count Regime

A limitation of our method is the optimized coding functions' suboptimal performance in the low photon count regime. This is because at low photon counts, the additive Gaussian noise approximation differs significantly from the Poisson noise model. Supp. Fig 26 plots the RMSE and MAE for different coding schemes at photon counts of  $\leq 100$ . Our optimized coding and Gray codes produce excessive outliers below 100 mean photons, while the FRH and Trunc. Fourier only produces a moderate number of outliers and maintains overall a lower depth error. Above 150 photons, our learned coding is able to achieve the closest accuracy to FRH regardless of the SBR.




Figure 26: **Optimized codes in photon-starved regime.** This figure shows the RMSE (bottom row) and MAE (top row) obtained from Monte Carlo simulations of our different coding schemes for photon count under 150, evaluated at IRF widths of  $\sigma = [10\Delta, 20\Delta, 30\Delta]$  from left to right.

## 9 Additional Results

In this section, we present additional results not shown in the main paper. We present both clipped pulse energy and constant pulse energy results and give them as surface plots.

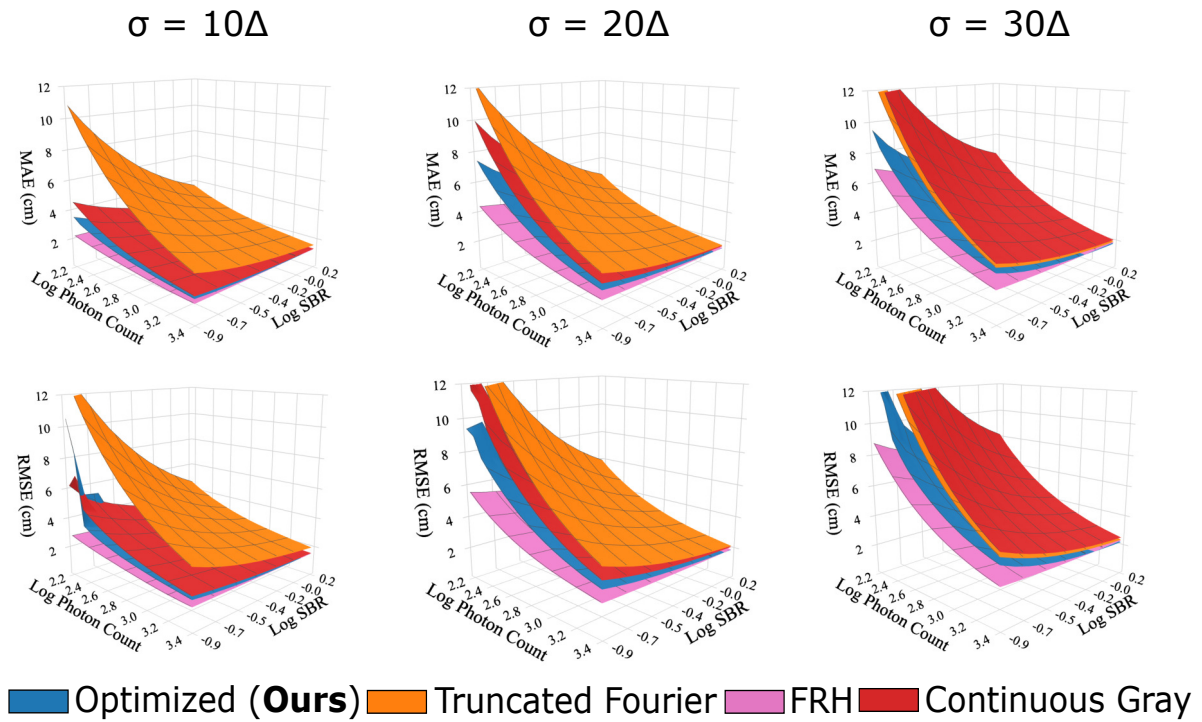


Figure 27: **Simulated results with finite bandwidth:** This figure shows the RMSE (bottom row) and MAE (top row) obtained from Monte Carlo simulations of our different coding schemes, evaluated at IRF widths of  $\sigma = [10\Delta, 20\Delta, 30\Delta]$  from left to right.

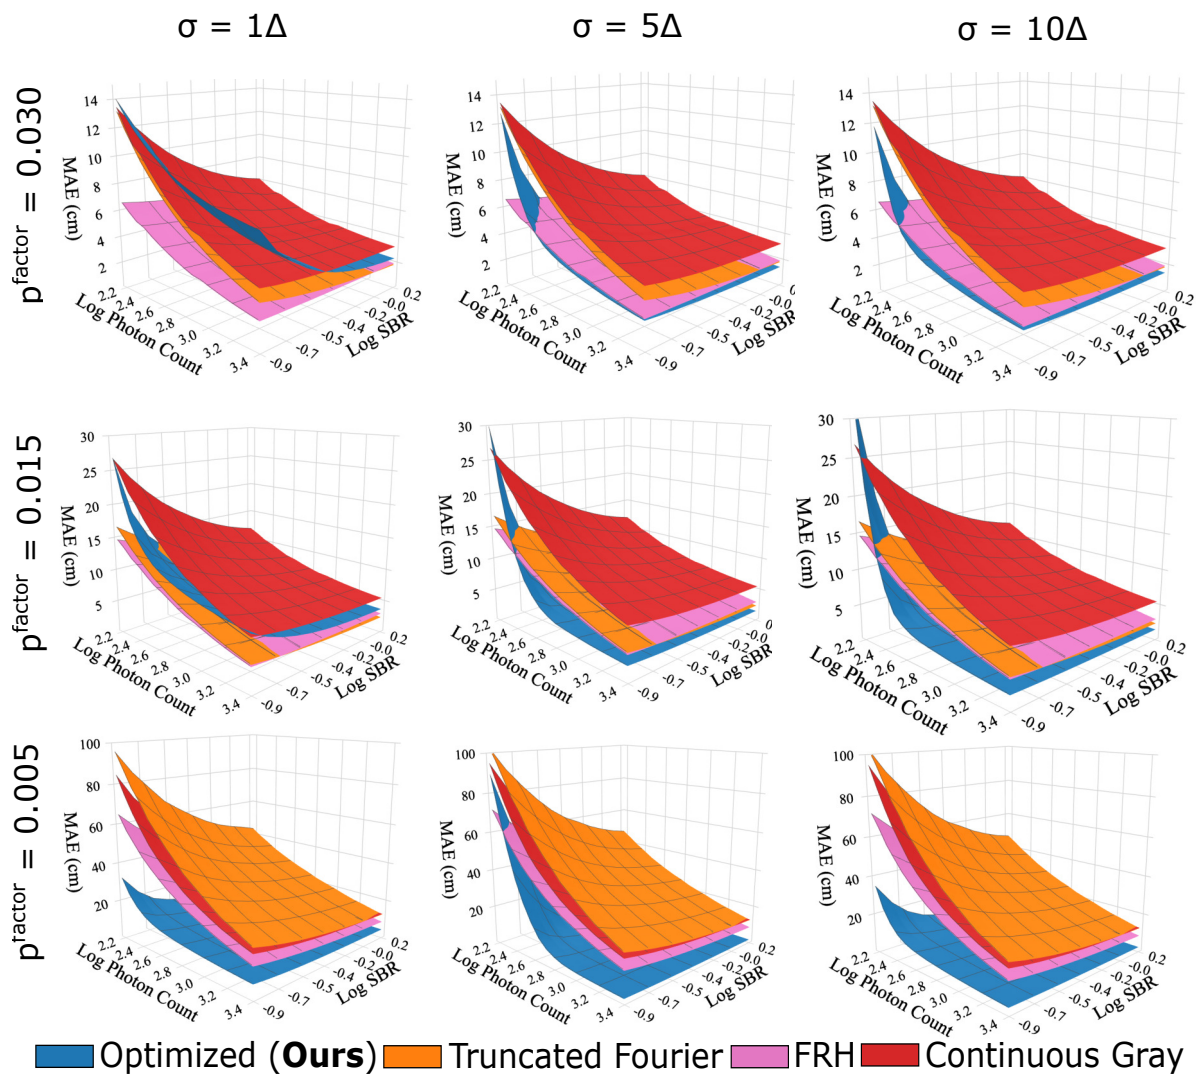


Figure 28: MAE results for peak power-limited coding functions while keeping pulse energy constant: This figure shows the mean absolute error (MAE) obtained from Monte Carlo simulations of our different coding schemes, evaluated at IRF widths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  from left to right and peak factors of  $p^{\text{factor}} = [0.005, 0.015, 0.030]$  from bottom to top when keeping pulse energy constant.

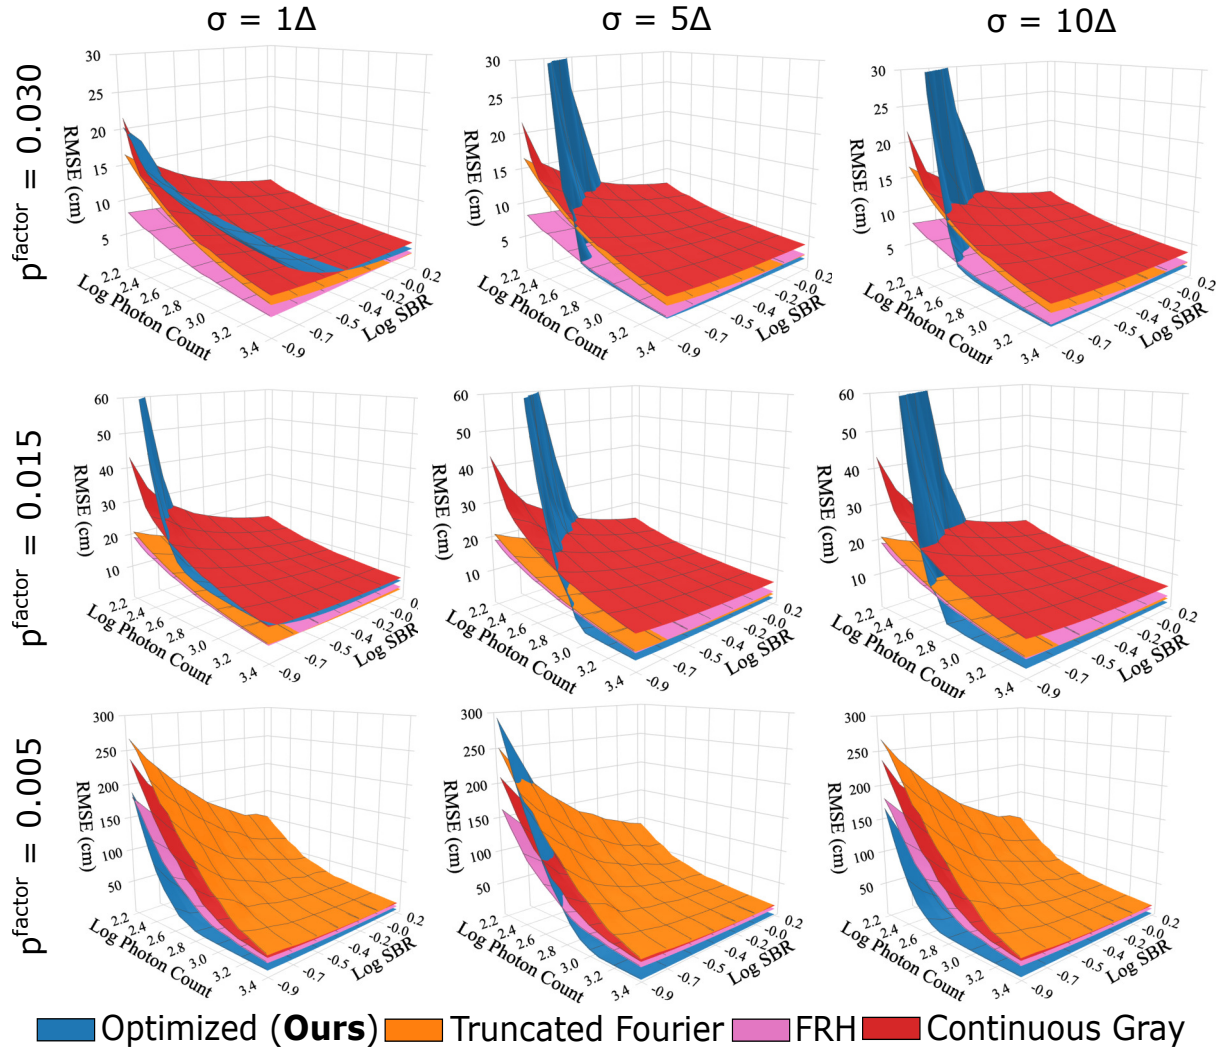


Figure 29: **RMSE results for peak power-limited coding functions while keeping pulse energy constant:** This figure shows the mean absolute error (MAE) obtained from Monte Carlo simulations of our different coding schemes, evaluated at IRF widths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  from left to right and peak factors of  $p^{\text{factor}} = [0.005, 0.015, 0.030]$  from bottom to top when keeping pulse energy constant.

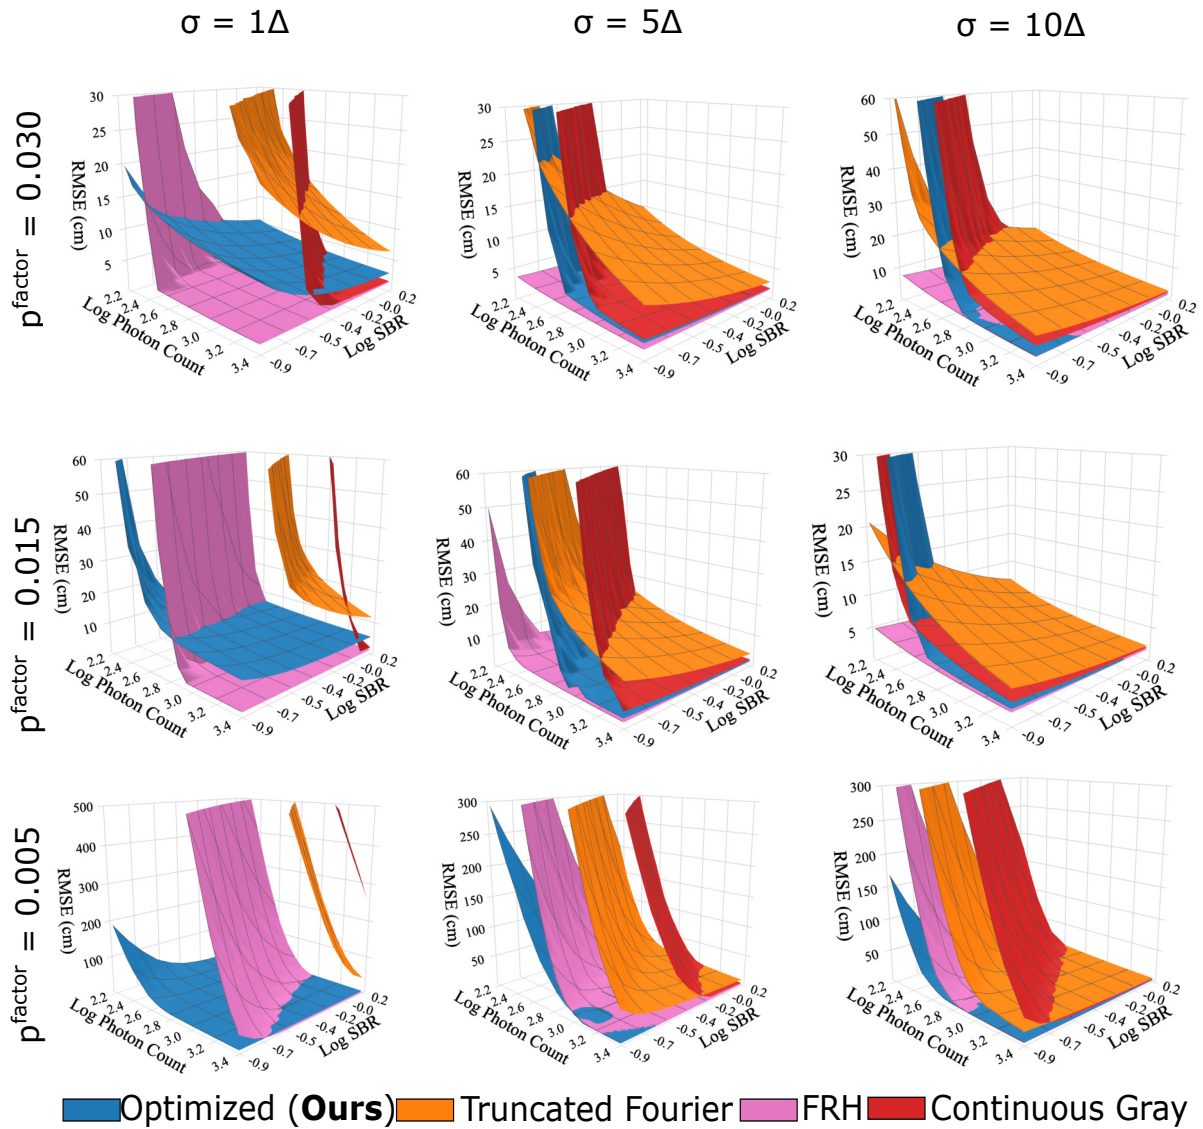


Figure 30: **RMSE results for peak power-limited coding functions when clipping pulse energy:** This figure shows the mean absolute error (MAE) obtained from Monte Carlo simulations of our different coding schemes, evaluated at IRF widths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  from left to right and peak factors of  $p^{\text{factor}} = [0.005, 0.015, 0.030]$  from bottom to top when clipping pulse energy.

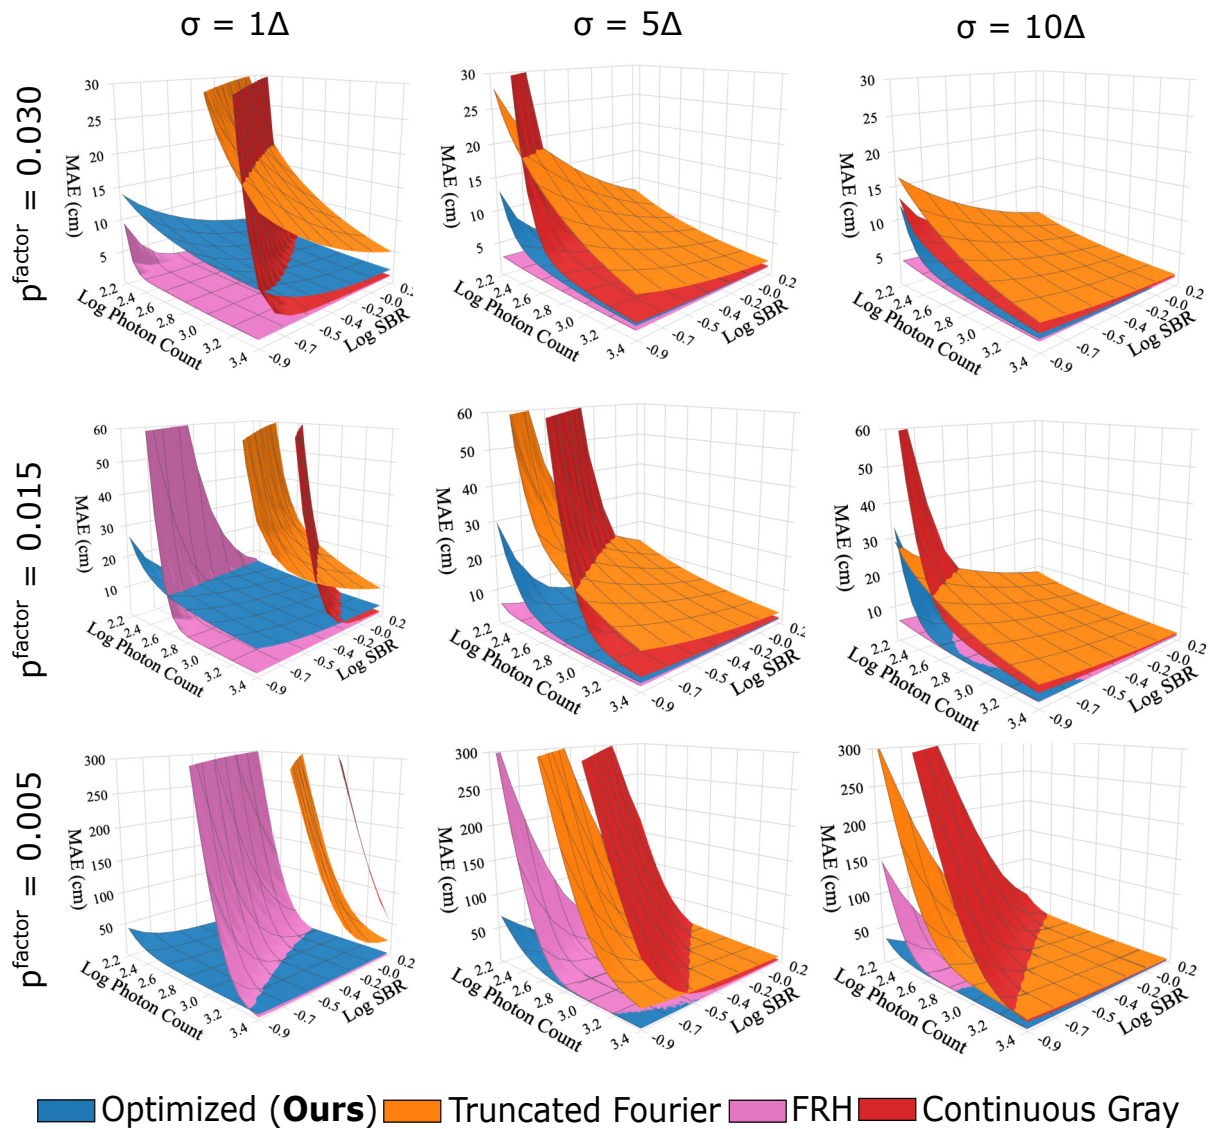


Figure 31: MAE results for peak power-limited coding functions when clipping pulse energy: This figure shows the mean absolute error (MAE) obtained from Monte Carlo simulations of our different coding schemes, evaluated at IRF widths of  $\sigma = [1\Delta, 5\Delta, 10\Delta]$  from left to right and peak factors of  $p^{\text{factor}} = [0.005, 0.015, 0.030]$  from bottom to top when clipping pulse energy.

## References

- [1] W. Rudin, *Real and Complex Analysis*. McGraw-Hill Book Company, 1987.
- [2] 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.
- [3] 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.
- [4] 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. 17 854–17 864.
- [5] 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.
- [6] 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.
- [7] A. Ardelean, “Computational imaging spad cameras,” PhD thesis, École Polytechnique Fédérale de Lausanne (EPFL), 2023.
- [8] 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. 10 756–10 766.
- [9] J. T. Barron, “A general and adaptive robust loss function,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4331–4339.