photonics

MDPI

Article

Non-Line-of-Sight Imaging via Sparse Bayesian Learning Deconvolution

Yuyuan Tian [1,2], Weihao Xu [1], Dingjie Wang [1], Ning Zhang [3], Songmao Chen [1,4], Peng Gao [5], Xiuqin Su [1,4] and Wei Hao [1,4,*]

[1] Key Laboratory of Space Precision Measurement Technology, Xi'an Institute of Optics and Precision Mechanics, Chinese Academy of Sciences, Xi'an 710119, China; tianyuyuan18@mails.ucas.ac.cn (Y.T.); xuweihao@opt.cn (W.X.); wangdingjie17@mails.ucas.ac.cn (D.W.); chensongmao@opt.ac.cn (S.C.); suxiuqin@opt.ac.cn (X.S.)
[2] University of Chinese Academy of Sciences, Beijing 100049, China
[3] Key Laboratory of Spectral Imaging Technology, Xi'an Institute of Optics and Precision Mechanics, Chinese Academy of Sciences, Xi'an 710119, China; zhangning@opt.ac.cn
[4] Laoshan Laboratory for Marine Science and Technology, Qingdao 266237, China
[5] Xi'an Institute of Optics and Precision Mechanics, Chinese Academy of Sciences, Xi'an 710119, China; gaopeng1@opt.ac.cn
* Correspondence: hwei@opt.ac.cn

Abstract

By enhancing transient fidelity before geometric inversion, this work revisits the classical LCT-based non line-of-sight (NLOS)imaging paradigm and establishes a unified Bayesian sparse-enhancement framework for reconstructing hidden objects under photon-starved and hardware-limited conditions. We introduce sparse Bayesian learning (SBL) as a dedicated front-end transient restoration module, leveraging adaptive sparsity modeling to suppress background fluctuations while preserving physically consistent multipath returns. This lightweight and geometry-agnostic design enables seamless integration into existing LCT processing pipelines, granting the framework strong compatibility with diverse acquisition configurations. Comprehensive simulations and experiments on complex reflective targets demonstrate significant improvements in spatial resolution, boundary sharpness, and robustness to IRF-induced temporal blurring compared with traditional LCT and f-k migration methods. The results validate that transient quality remains a critical bottleneck in practical NLOS deployment, and addressing it via probabilistic sparsity inference offers a scalable and computationally affordable pathway toward stable, high-fidelity NLOS reconstruction. This study provides an effective signal-domain enhancement solution that strengthens the practicality of NLOS imaging in real-world environments, paving the way for future extensions toward dynamic scenes, multi-view fusion, and high-throughput computational sensing.

Keywords: non line-of-sight (NLOS) imaging; LiDAR; computational imaging; SPADs; sparse Bayesian learning

Check for updates

Received: 25 November 2025
Revised: 30 December 2025
Accepted: 30 December 2025
Published: 6 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Non-line-of-sight (NLOS) imaging aims to reconstruct targets that are not directly visible to the sensor. Such scenarios occur when objects are hidden behind corners or located inside enclosed environments. In these cases, light travels to the hidden scene via multiple diffuse reflections on an intermediate surface and returns to the detector only after significant attenuation. The recorded transient signals are usually weak due to photon

losses over multipath propagation and are also affected by noise and instrument response effects. Furthermore, the temporal profiles are broadened by convolution with the system impulse response function (IRF), which reduces temporal resolution and mixes information in the histogram domain [1,2]. As a result, computational methods are required to extract useful information from the degraded observations and recover the concealed structure.

Active illumination solutions based on time-correlated single-photon counting (TCSPC) are widely employed for their single-photon sensitivity and picosecond-level temporal resolution [1]. Various time-resolved detectors have been implemented in practice, including streak cameras [1,3], intensified CCDs [4,5], single-photon avalanche diode (SPAD) sensors [6–8], and superconducting nanowire devices [9,10]. By exploiting the time-of-flight (ToF) characteristics of triple-bounce photon paths, these systems enable the recovery of hidden 3D structures even under low-light or nighttime conditions. Recent studies have also explored system-level modifications of existing detection architectures to improve acquisition efficiency, such as spatial multiplexing detection based on a passive single-pixel camera [11]. At the detection level, time-sequential first-photon acquisition has also been explored to improve photon-event efficiency in photon-starved NLOS imaging [12]. By exploiting the time-of-flight (ToF) characteristics of triple-bounce photon paths, these systems enable the recovery of hidden 3D structures even under low-light or nighttime conditions. To improve reconstruction quality in challenging photon-starved regimes, various computational methods have been proposed. Classical inversion techniques such as filtered back-projection [1], light cone transform (LCT) [2], phasor field propagation [13], and f-k migration [14] rely on geometric constraints of light transport but typically retain blur and noise in the reconstructed transient domain. Beyond these classical formulations, recent efforts have extended the physical forward model and inversion assumptions. For example, Liu et al. formulated NLOS reconstruction under arbitrary illumination and detection patterns as a generalized linear inverse problem, relaxing the requirement of regular scanning geometries [15], and Young et al. proposed the directional light-cone transform (DLCT) to incorporate directional light transport cues for enhanced surface reconstruction within the LCT framework [16]. Signal restoration methods attempt to compensate for IRF influences and better recover fine structural information. For example, Richardson–Lucy deconvolution [17] improves temporal sharpness under Poisson noise but may become unstable under low signal-to-noise ratio conditions. Regularization-based approaches embed prior knowledge such as sparsity, smoothness, or curvature to enhance robustness [18–20]; however, the performance can be sensitive to parameter settings and initial estimation accuracy. Operator-level refinements including product convolution and low-rank acceleration [21] improve computational efficiency while potentially compromising weak signal components. Learning-based reconstruction strategies [22,23] have shown encouraging results in speed and image quality, yet their reliability may be reduced when encountering measurement conditions not well represented in the training dataset, and the inference process may require considerable computational resources.

Despite these advances, reliably recovering fine transient structures under photon-starved conditions and IRF-induced temporal blurring remains a fundamental challenge in practical NLOS systems. In particular, degradation introduced at the transient level often propagates through geometric inversion, leading to spatial blur, depth bias, and background artifacts in the final reconstruction.

To address this issue, we introduce sparse Bayesian learning (SBL) as a plug-and-play transient restoration module preceding LCT reconstruction. The key contributions of this work are summarized as follows:

- Bayesian sparsity-driven transient enhancement: SBL selectively suppresses noise while preserving physically consistent multipath photon returns, effectively restoring informative temporal structures;
- Hardware-free compatibility with existing LCT pipelines: The proposed module is lightweight and geometry-agnostic, allowing direct integration without acquisition or system modifications;
- Improved reconstruction fidelity under photon-starved conditions: Enhanced transients provide more reliable input for light-cone propagation, yielding clearer boundaries, finer spatial details, and stronger robustness to IRF distortions;
- Validated effectiveness on both simulated and real measurements: Results confirm that transient quality is a critical bottleneck in practical NLOS imaging, and addressing it promotes scalable and reliable real-world deployment.

In summary, this work offers a practical and computationally efficient pathway to improving the robustness and fidelity of NLOS reconstruction in challenging environments.

2. Light-Cone Transform Algorithm

2.1. Non-Line-of-Sight Imaging Problem

In non-line-of-sight (NLOS) imaging, photons emitted from a pulsed laser undergo three successive diffuse reflections before reaching the detector. The contribution of a hidden scene point $\mathbf{x}$ to the received transient signal is formulated as:

$$\tau(\mathbf{x}', \mathbf{x}'', t) = \iiint_{\Omega} \rho(\mathbf{x}) \cdot \frac{\delta(\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{x} - \mathbf{x}''\| - tc)}{\|\mathbf{x} - \mathbf{x}'\|^2 \cdot \|\mathbf{x} - \mathbf{x}''\|^2} \, d\mathbf{x}, \tag{1}$$

where $\rho(\mathbf{x})$ is the reflectance at voxel $\mathbf{x} \in \Omega$. Due to the cubic propagation loss, the signal arriving at the detector is attenuated by 10^6–10^9 [2], leading to an extremely sparse photon distribution along the temporal axis.

As shown in Figure 1, this work considers a confocal non-line-of-sight (NLOS) imaging scenario, in which the illumination and detection points coincide. Accordingly, in a confocal configuration ($\mathbf{x}' = \mathbf{x}''$), Equation (2) is simplified as

$$\tau(\mathbf{x}', t) = \iiint_{\Omega} \rho(\mathbf{x}) \cdot \frac{\delta(2\|\mathbf{x}' - \mathbf{x}\| - tc)}{\|\mathbf{x}' - \mathbf{x}\|^4} \, d\mathbf{x}. \tag{2}$$

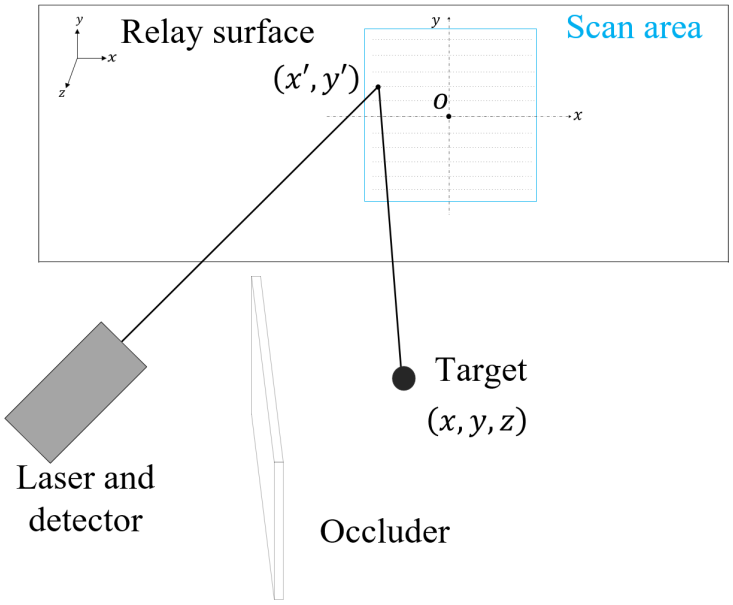

Figure 1. Confocal NLOS imaging geometry.

Despite this inherent sparsity, the temporal reflectance function is significantly blurred and overlapped due to convolution with the instrument response function (IRF), obscuring the exact photon arrival times and reducing the robustness of inverse reconstruction.

2.2. Light Cone Transform (LCT) Algorithm

Following O'Toole et al. [2], mappings $z = \sqrt{u}$ and $v = (ct/2)^2$ reformulate Equation (2) into a shift-invariant 3D convolution:

$$R_t\{\tau\}(\mathbf{x}', v) = h * R_z\{\rho\}, \tag{3}$$

where $R_t\{\tau\}$ and $R_z\{\rho\}$ represent nonuniform resampling and geometric correction along the temporal and depth axes, respectively.

The discrete form is expressed as

$$\tilde{\boldsymbol{\tau}} = \mathbf{H}\,\mathbf{R}_z\,\boldsymbol{\rho}, \tag{4}$$

and the closed-form Wiener inversion is

$$\boldsymbol{\rho}_* = \mathbf{R}_z^{-1}\mathcal{F}^{-1}\left(\frac{\hat{\mathbf{H}}^*}{|\hat{\mathbf{H}}|^2 + 1/\alpha}\,\mathcal{F}(\tilde{\boldsymbol{\tau}})\right). \tag{5}$$

However, backprojection over the light-cone surface accumulates measurement noise and IRF-induced artifacts. Therefore, temporal sparsity and localization of the input transient signals are critical for preserving geometric consistency in the volumetric reconstruction.

3. Sparse Bayesian Learning in NLOS

3.1. Sparse Bayesian Learning

In NLOS imaging, the measured transient waveform is broadened by the system impulse response function (IRF) and contaminated by stochastic photon-counting noise. The deconvolution problem can be modeled as

$$\boldsymbol{\tau} = \mathbf{K}\mathbf{y} + \boldsymbol{n}, \tag{6}$$

where $\mathbf{K}$ is a convolution operator derived from the IRF, $\mathbf{y}$ denotes the latent transient signal, and $\boldsymbol{n}$ represents measurement noise dominated by photon-counting fluctuations.

The sparsity considered in this work does not correspond to strict ℓ_0 sparsity assumption at the single-photon event level. In a typical confocal NLOS system, the measured transient at each scanning position is a time histogram accumulated over many laser pulses. Only a small subset of temporal bins corresponds to physically valid multi-bounce photon returns, while the majority of bins are dominated by background photons, dark counts, or stochastic fluctuations. Even after convolution with the IRF, these informative components remain temporally localized, resulting in an effectively sparse structure along the time axis. From a statistical perspective, although individual photon arrivals follow Poisson statistics, the accumulated transient histograms under moderate photon counts and temporal accumulation can be reasonably approximated by a Gaussian distribution. Accordingly, a Gaussian likelihood is adopted for stable and tractable inference, while the hierarchical prior in SBL serves as an adaptive mechanism to distinguish physically consistent temporal components from noise-dominated or uncertain fluctuations.

To account for the physical sparsity of temporally localized photon returns, Sparse Bayesian Learning (SBL) introduces a hierarchical Gaussian prior

$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}), \qquad \boldsymbol{\Gamma} = \mathrm{diag}(\gamma_1, \ldots, \gamma_T), \tag{7}$$

where γ_i denotes the adaptive variance associated with the i-th temporal component.

Posterior inference yields the Gaussian posterior distribution

$$\boldsymbol{\Sigma} = \left(\boldsymbol{\Gamma}^{-1} + \tfrac{1}{\sigma^2}\mathbf{K}^\top\mathbf{K}\right)^{-1}, \tag{8}$$

$$\boldsymbol{\mu} = \tfrac{1}{\sigma^2}\boldsymbol{\Sigma}\mathbf{K}^\top\boldsymbol{\tau}, \tag{9}$$

where σ^2 approximates the measurement noise variance.

The hyperparameters are iteratively updated via type-II maximum likelihood as

$$\gamma_i^{\text{new}} = \mu_i^2 + \Sigma_{ii}. \tag{10}$$

Components associated with low posterior energy, characterized by small values of $\mu_i^2 + \Sigma_{ii}$, are progressively suppressed through the adaptive variance update, whereas physically meaningful signal structures with stable temporal support are retained through iterative evidence maximization.

Therefore, the SBL-estimated transient signal is obtained as

$$\hat{\mathbf{y}} = \boldsymbol{\mu}. \tag{11}$$

Through uncertainty-aware sparsity modeling and adaptive variance regularization, the proposed SBL approach provides a physically consistent statistical enhancement of temporally localized photon returns, thereby serving as a well-conditioned input for subsequent geometric inversion.

It is worth noting that the proposed SBL formulation does not aim to explicitly invert or remove the instrument response function (IRF). Instead, the IRF is incorporated as a fixed forward convolution operator within the inference model. Under this formulation, moderate IRF modeling inaccuracies, such as mild temporal broadening or imperfect tail characterization, do not necessarily introduce additional instability into the restored transients. Their impact is primarily reflected as increased temporal uncertainty, while dominant multipath components that remain consistent with the forward model are preserved.

From an implementation perspective, compared with classical deconvolution methods such as Richardson–Lucy or Tikhonov filtering, SBL introduces additional computational cost due to adaptive variance estimation. In this work, SBL is applied independently to each one-dimensional transient waveform and executed only once as a preprocessing step prior to LCT reconstruction. As a result, the additional overhead does not alter the computational complexity of the subsequent geometric inversion. In typical NLOS imaging scenarios where reconstruction is performed offline, this overhead remains computationally affordable and does not affect the overall reconstruction pipeline.

3.2. Enhanced LCT Reconstruction via SBL Preprocessing

To clarify the origin of reconstruction blur, it is important to distinguish temporal-domain degradation from spatial-domain artifacts in confocal NLOS imaging. In practice, the dominant degradation is introduced in the temporal domain rather than by an explicit spatial convolution. The finite width of the system impulse response function (IRF), together with the nonzero laser pulse duration and detector timing jitter, causes ideal photon returns to be broadened along the time axis.

This temporal broadening can be accurately modeled as a per-pixel one-dimensional convolution,

$$y_p(t) = h(t) * x_p(t) + n_p(t), \tag{12}$$

where $h(t)$ denotes the measured IRF, $x_p(t)$ represents the latent transient signal at the p-th scanning position, and $n_p(t)$ denotes noise.

This per-pixel formulation is consistent with the vectorized model in Section 3.1. In particular, $y_p(t)$ corresponds to the latent transient waveform at the p-th scanning position, and stacking all $y_p(t)$ across spatial positions yields the vector $\boldsymbol{y}$ used in the SBL formulation.

Accordingly, the convolution operator $\mathbf{K}$ in Section 3.1 can be interpreted as a block-diagonal operator induced by the IRF $h(t)$, where each block performs a one-dimensional temporal convolution at an individual scanning position.

When temporally broadened transients are directly used for geometric reconstruction, the temporal uncertainty propagates into the spatial domain through the light-cone integration process. In the Light Cone Transform (LCT), photon energy is accumulated along ellipsoidal iso-time surfaces; temporal broadening effectively increases the thickness of these surfaces, leading to spatial diffusion, blurred object boundaries, and background artifacts. Therefore, the observed spatial blur is not caused by a spatial convolution itself, but rather arises from the propagation of temporal convolution effects through geometric mapping.

This analysis indicates that improving transient fidelity prior to geometric inversion is essential for mitigating reconstruction blur in LCT-based NLOS imaging.

Light Cone Transform (LCT) reconstructs hidden geometry by integrating photon arrival times along ellipsoidal iso-time surfaces, making its performance highly dependent on the sparsity and temporal accuracy of the input transient signals. If the measurement is substantially broadened by the IRF or corrupted by background fluctuations, light-cone accumulation becomes biased, resulting in blurry object contours, depth misregistration, and cluttered background artifacts. The proposed SBL-based preprocessing introduces adaptive variance modulation within a Bayesian framework, where sparsity priors promote compact and physically meaningful photon returns while suppressing noise-driven fluctuations. As a result, weak multi-bounce reflections, which are often overwhelmed in classical deconvolution, are better preserved, and the recovered peaks show more accurate temporal localization. This yields a more reliable temporal-to-spatial mapping for LCT inversion, reducing spurious backprojection in empty voxels while enhancing path consistency across different depths. Consequently, the SBL–LCT reconstruction pipeline achieves cleaner background, sharper boundary transitions, and improved geometric fidelity even under low-photon or high-noise conditions, providing favorable conditions for high-quality and robust 3D NLOS reconstruction in practical scenarios.

Furthermore, the proposed SBL–LCT pipeline is designed in a modular form. SBL acts as an independent transient enhancement front-end without altering the core geometric constraints of the LCT inversion. This makes the proposed method seamlessly compatible with existing NLOS reconstruction frameworks and hardware configurations. The modular design reduces deployment complexity and provides a flexible interface for further extensions, such as dynamic scene recovery, multi-view fusion, or integration with neural rendering backends.

In light of this modular design, we further summarize how the proposed SBL-based transient enhancement is incorporated into the overall NLOS reconstruction pipeline in the following subsection.

3.3. Overall Processing Workflow

For clarity of practical implementation, the complete processing pipeline of the proposed method is summarized as follows and illustrated in Figure 2. The pipeline consists of four main stages: transient data acquisition, temporal degradation modeling, SBL-based transient restoration, and geometric reconstruction.

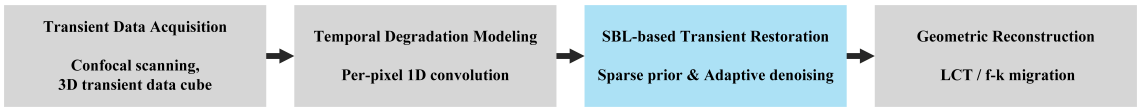

Figure 2. The overall processing workflow of the proposed SBL-based NLOS reconstruction pipeline. The pipeline consists of four stages: transient data acquisition, temporal degradation modeling, SBL-based transient restoration, and geometric reconstruction (e.g., LCT or f-k migration).

First, time-resolved transient signals are acquired using a confocal scanning configuration, where each scanning position on the relay wall corresponds to a one-dimensional transient waveform. All measurements are organized into a three-dimensional transient data cube. Meanwhile, the system impulse response function (IRF) is experimentally calibrated and used as the convolution kernel for temporal degradation modeling.

Second, temporal degradation at each scanning position is modeled as a one-dimensional convolution between the latent transient signal and the measured IRF, forming a physically interpretable forward model for signal restoration.

Third, Sparse Bayesian Learning (SBL) is applied independently to each transient waveform to perform deconvolution and noise suppression. By introducing adaptive sparsity priors, SBL selectively suppresses background fluctuations while preserving physically meaningful multipath photon returns. The restored transients are then reassembled into an enhanced transient data cube.

Finally, the SBL-enhanced transient data cube is fed into standard geometric reconstruction algorithms, such as LCT or f-k migration, without modifying the underlying geometric model. This modular design enables seamless integration with existing NLOS reconstruction pipelines.

4. Results and Discussion

4.1. Simulation Results

In order to verify the effectiveness of the SBL algorithm for signal preprocessing in the NLOS problem, both simulation and experimental data are used to compare reconstruction results.

4.1.1. Signal-Domain Comparison of SBL, RL and Tikhonov

To assess the transient enhancement capability prior to volumetric NLOS reconstruction, we take experimental data of the hidden letter **E** as a representative example, and compare SBL with two classical deconvolution approaches: Richardson–Lucy (RL) and Tikhonov inverse filtering (TK).

Figure 3 visualizes the processed transient signals from all 33×33 scanning positions, rearranged into a time–pixel matrix for global inspection (Figure 3a). Figure 3b further shows the same data cropped to the ROI of 1050–1450 bins, matching the subsequent waveform-level analysis. For practical reference, Table 1 summarizes the per-pixel processing time of the three transient preprocessing methods applied to the same $33 \times 33 \times 4096$ measurement, where the SBL-based enhancement is applied to a temporally gated window centered around the dominant echo energy. The timings were recorded on a system with an Intel Core i7-8750H CPU and 16 GB RAM, using parallel processing with 4 workers.

As observed, the three deconvolution approaches lead to clearly different transient behaviors. SBL produces a visibly more compact and cleaner response with strong suppression of background fluctuations, which makes the dominant return components more separable in time. RL sharpens the peak structure but introduces oscillatory ringing patterns near strong transitions, consistent with iterative amplification of high-frequency components. Tikhonov inverse filtering can also sharpen the main peak, yet the recovered traces exhibit increased low-amplitude oscillations, indicating sensitivity to noise and model mismatch.

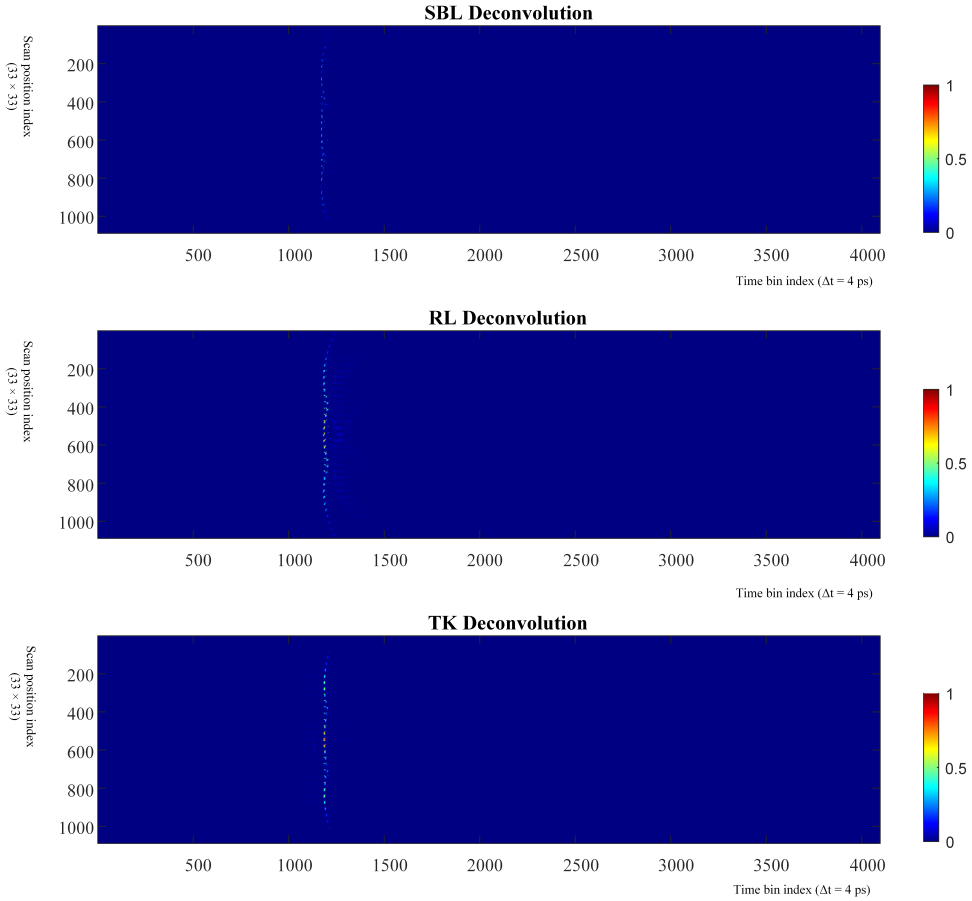

(**a**) Global maps (SBL/RL/TK)

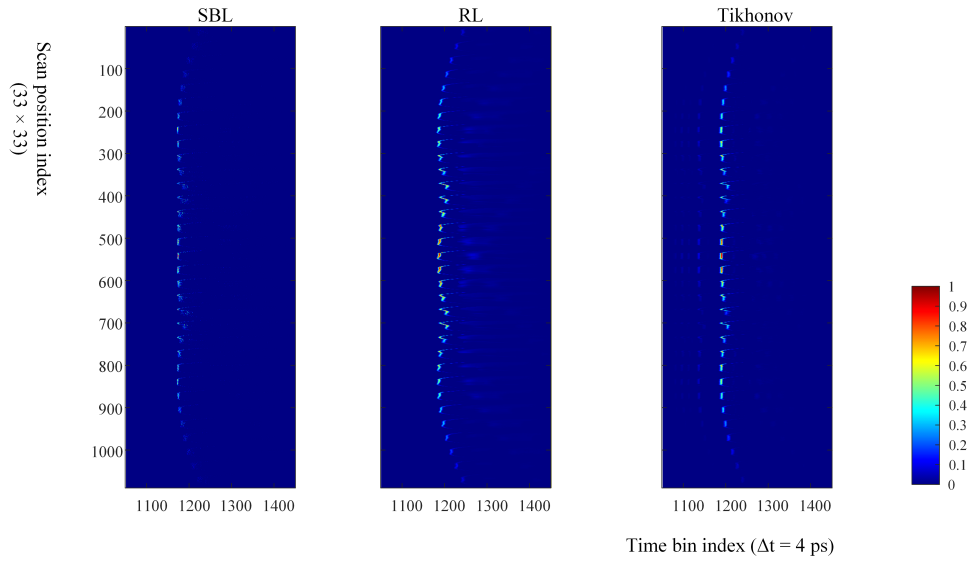

(**b**) Region of Interest (1050–1450 bin)

Figure 3. Comparison of transient signals across all 33 × 33 scanning positions. (**a**) Global time–pixel maps produced by SBL, RL, and Tikhonov deconvolution (as arranged in the figure). (**b**) Temporal Region of Interest (ROI) cropped to bins 1050–1450, consistent with the subsequent signal-domain evaluation. For visualization clarity, each transient map is normalized by its own maximum amplitude. All panels are displayed using an identical color scale ranging from 0 to 1, enabling a consistent interpretation of within-method contrast and background suppression (note that absolute energy across methods is not compared after per-map normalization).

Table 1. A runtime comparison of per-pixel transient preprocessing on the $33 \times 33 \times 4096$ experimental transient cube (letter E).

Method	Time/Pixel (s)
SBL (windowed, $D \in \mathbb{R}^{256\times256}$)	0.016231
RL (Richardson–Lucy)	0.009528
TK (Tikhonov inverse filtering)	0.000398

Implementation details: All methods process identical per-pixel 1D transients with the same IRF calibration. RL and TK perform IRF peak alignment internally, while SBL constructs the corresponding convolution operator once and then performs iterative evidence maximization per pixel.

As shown in Figure 4, the raw measurement exhibits an evidently broadened and delayed peak due to IRF convolution, resulting in poor temporal localization. SBL restores a markedly sharper and more concentrated response, while simultaneously suppressing the low-level background fluctuations, leading to a waveform that is closer to the ideal temporally compact return. RL yields a relatively high waveform similarity in terms of global morphology, but it also introduces a noticeable temporal shift and ringing artifacts around the main transition. Tikhonov inverse filtering sharpens the peak as well; however, the recovered background becomes more oscillatory, suggesting reduced stability in low-return regions.

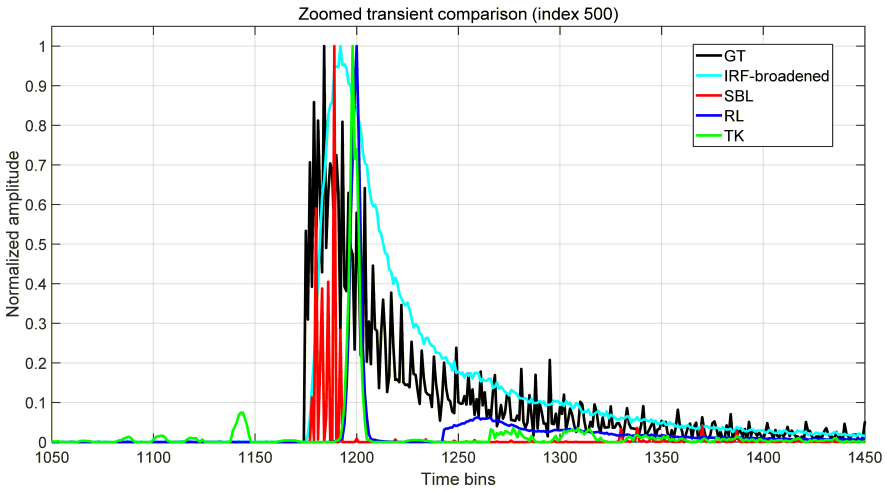

Figure 4. Zoomed-in view around the main peak (1050–1450 bins).

Quantitative results in Table 2 further corroborate these observations. The peak location and temporal shift evaluate the accuracy of dominant-return localization. FWHM (Full Width at Half Maximum) reflects temporal resolution, where smaller values indicate a more concentrated peak. BG_STD (Background Standard Deviation) characterizes robustness against background fluctuations, whereas SNR (Signal-to-Noise Ratio) quantifies the separability between the dominant return and the background. SSIM (Structural Similarity Index Measure) measures the overall waveform similarity with respect to the GT (Ground Truth) trace.

Among the compared methods, SBL achieves the smallest temporal shift (5 bins) and the narrowest peak width (FWHM = 1.27), together with a very low background fluctuation (BG_STD = 0.0001) and the highest SNR (105.40 dB). RL attains a higher SSIM, indicating better preservation of global waveform morphology, but it suffers from a much larger temporal shift (16 bins) and visible ringing artifacts. Tikhonov inverse filtering provides moderate peak narrowing, yet its background fluctuation is noticeably higher (BG_STD = 0.0009), implying higher sensitivity to noise. Notably, the raw signal can present a relatively high SSIM due to retaining the broadened waveform shape, but this comes at

the cost of severely degraded temporal resolution (FWHM = 31.22). These results suggest that the sparsity-guided Bayesian strategy offers a balanced trade-off among resolution, robustness, and structural fidelity.

These signal-domain results indicate that using SBL as a transient enhancement front-end can provide temporally more compact and higher-SNR measurements for downstream light-cone transform (LCT) inversion. Since LCT integrates photon arrivals along iso-time light-cone manifolds, temporal broadening and background clutter can be mapped into spatial blurring and artifacts. By reducing IRF-induced spreading and suppressing background fluctuations, the SBL-processed transients help improve temporal localization and mitigate clutter propagation in the subsequent reconstruction, especially under low-photon experimental conditions.

Table 2. Single-pixel transient metrics and SSIM at index 500 (vs. GT).

Method	Peak Location	Shift	FWHM	BG_STD	SNR (dB)	SSIM
GT	1184	0	1.86	0.0000	313.07	1.0000
Raw	1192	8	31.22	0.0004	67.34	0.9564
SBL	1189	**5**	**1.27**	**0.0001**	**105.40**	0.9323
RL	1200	16	5.25	**0.0001**	78.99	**0.9537**
TK	1198	14	4.75	0.0009	61.20	0.9363

Note: Bold values indicate the best quantitative performance among the compared methods.

4.1.2. Non-Line-of-Sight Reconstruction Results

Based on the SBL-processed transient signals, non-line-of-sight 3D reconstruction was performed on the simulated data. Figure 5 presents the reconstruction results of letters E and F. For each target, the first column shows the ground truth (GT), followed by the results of SBL-LCT, LCT, and f-k migration. The comparison clearly demonstrates the superior noise suppression and structural fidelity achieved by SBL-assisted signal processing.

To quantify the reconstruction fidelity, we report structural similarity index measure (SSIM), peak signal-to-noise ratio (PSNR), and standard deviation (STD) as evaluation metrics, where larger SSIM and PSNR values generally indicate higher similarity to the reference and lower distortion in the reconstructed images. The quantitative results are summarized in Table 3. As shown, the proposed SBL-LCT method achieves the highest SSIM and PSNR values for both targets, alongside the lowest standard deviation (STD), indicating superior consistency. Together, the quantitative and visual comparisons demonstrate that SBL preprocessing leads to sharper strokes, richer details and lower background noise than LCT alone and f-k migration, confirming the effectiveness of SBL-driven transient enhancement in simulated NLOS imaging tasks.

Table 3. Reconstruction quality metrics for simulated letters E and F.

Target	Algorithm	SSIM	PSNR (dB)	STD
E	**SBL-LCT**	**0.2005**	**11.2058**	**0.2146**
	LCT	0.1472	10.6974	0.2300
	f-k	0.1750	10.7851	0.2388
F	**SBL-LCT**	**0.1471**	**10.6429**	**0.1771**
	LCT	0.1040	10.0420	0.2384
	f-k	0.1334	10.3214	0.2620

Note: The bold formatting is used to highlight the proposed method.

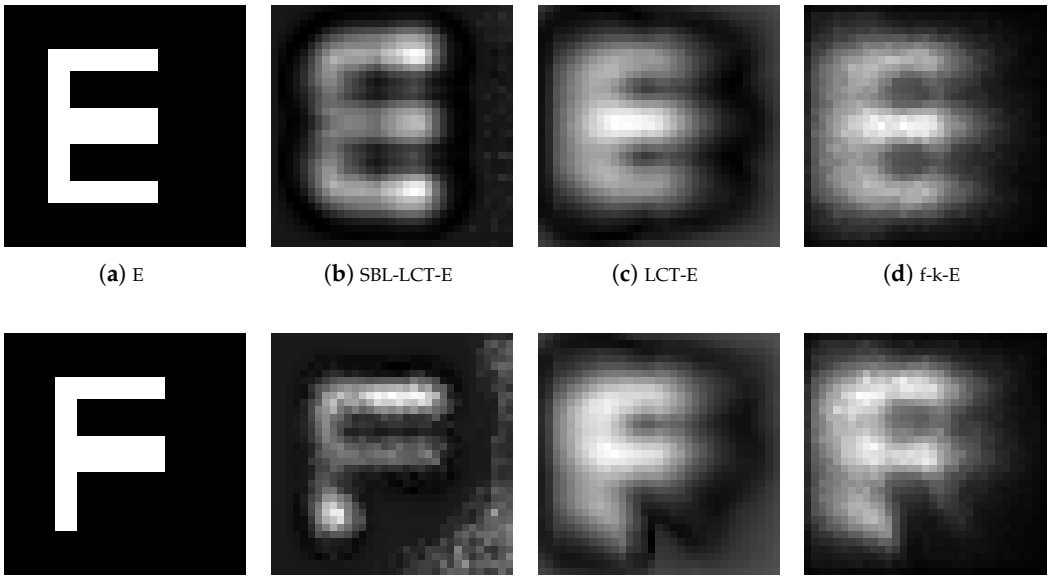

(**a**) E (**b**) SBL-LCT-E (**c**) LCT-E (**d**) f-k-E

(**e**) F (**f**) SBL-LCT-F (**g**) LCT-F (**h**) f-k-F

Figure 5. Reconstruction results of simulated letters **E** (top) and **F** (bottom).

4.2. Experiment Results

The experimental system adopts a coaxial optical path. All experiments were conducted in a controlled dark-room environment to suppress ambient light interference. The illumination source was an active pulsed laser scanning system rather than ambient or passive light. A PicoQuant LDH-D-C-850 pulsed laser was employed as the active illumination source, and its internal synchronization signal was used as the temporal reference for TCSPC measurements.

The system configuration used in the experiments is summarized in Table 4. A curtain was used to separate the hidden target region from the NLOS imaging system, while additional light shields and baffles were used in the detection path to reduce the influence of stray light.

Table 4. Experimental System Configuration.

Hardware	Setup
Laser source	PicoQuant LDH-D-C-850 (PicoQuant GmbH, Berlin, Germany) Pulse width: 500 ps (FWHM) Repetition rate: 2.5 MHz (experimental setting) Average power: 0.1 mW@1 MHz
SPAD	EXCELITAS SPCM-AQRH-16-FC (Excelitas Technologies Corp., Vaudreuil-Dorion, QC, Canada) Photon detection efficiency: 45%@850 nm Dark count rate: <25 cps Dead time: 22 ns
TCSPC module	PicoQuant PicoHarp 300 (PicoQuant GmbH, Berlin, Germany) Selected temporal resolution: 16 ps

The LDH-D-C-850 pulsed laser operates at a center wavelength of approximately 850 nm and emits optical pulses with a temporal width of about 500 ps (FWHM). In our experiments, the laser repetition rate was set to 2.5 MHz. The average output power is specified as 0.1 mW at a repetition rate of 1 MHz according to the manufacturer's datasheet, which corresponds to the nominal single-pulse energy level of the laser source.

The system impulse response function (IRF) used in the experiments is obtained through dedicated calibration on the same confocal NLOS setup. The IRF is measured with an extended integration time to improve the signal-to-noise ratio, and subsequently normalized and truncated to remove invalid leading bins before being applied as the deconvolution kernel in the SBL preprocessing stage. The calibrated IRF is also adopted in the simulation section, ensuring consistency between the experimental and simulated processing pipelines. Figure 6 shows the normalized IRF measured from our system and used throughout this work.

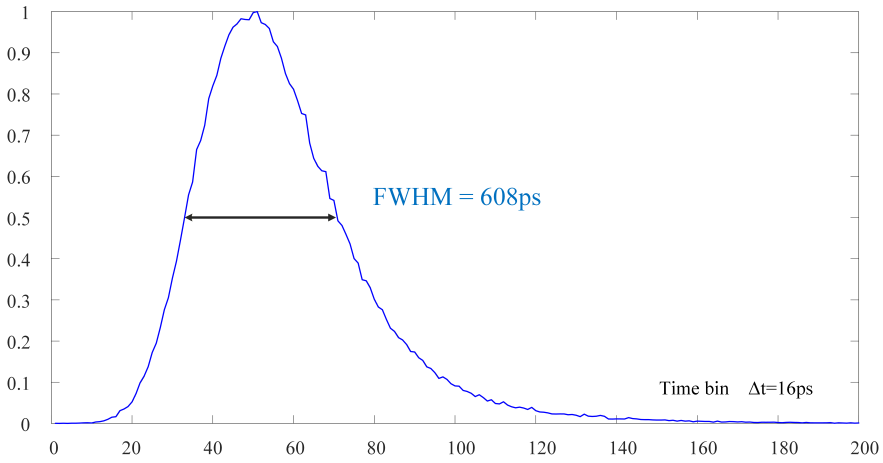

Figure 6. Normalized IRF. The experimentally calibrated system impulse response function (IRF) used in both simulations and experiments. The IRF is normalized and truncated by removing invalid leading bins before being used as the deconvolution kernel in the SBL preprocessing stage.

As shown in Figure 7a, the experimental setup consists of a scanning system (laser source, optical components, scanning mirror, and detector) illuminating a 1.12 m × 1.12 m area on the relay wall, with the hidden target placed behind an occluding curtain. The confocal detection scheme in Figure 7b shows that both the emission (red) and reception (green) paths pass through a perforated mirror and share the same scanning mirror.

Specifically, the laser beam emitted from the source is first guided by reflective mirrors for beam alignment, then passes through the central aperture of a perforated mirror (5 mm diameter), and is steered by a scanning mirror (GVS-012, Thorlabs Ltd., Ely, UK). By electronically controlling the scanning mirror, the laser beam performs point-by-point raster scanning over the relay wall. The same scanning mirror is also used for signal collection. Backscattered photons are redirected toward the perforated mirror, coupled into a collection lens system, and guided into an optical fiber connected to the SPAD detector. A narrowband optical filter (Semrock FF01-850-10-25, Rochester, NY, USA) centered at 850 nm is inserted in the detection path to suppress out-of-band background photons. Photon arrival times are recorded by the TCSPC module, forming a time-resolved transient histogram at each scanning position.

Foam boards were precisely cut into alphabetic shapes and assembled to form a series of planar experimental targets, including the letter combinations **CH**, **HS**, **C** with a tilted **H**, and **UCAS**. All targets were fabricated from identical material and mounted at fixed positions approximately 0.8 m in front of the diffuse relay wall. The illuminated and scanned area on the relay wall was defined as 1.12 m × 1.12 m, ensuring full coverage of the backscattered light corresponding to the hidden targets.

During data acquisition, raster scanning was performed over this region with spatial resolutions configured as follows: 33 × 33 sampling points for **CH** and **HS**, 65 × 65 for **C** with a tilted **H**, and 64 × 64 for **UCAS**. These sampling grids represent the discrete

observation points on the relay wall, balancing measurement time and spatial resolution for each target. Representative reconstruction results of these experimental scenes are presented in Figure 8.

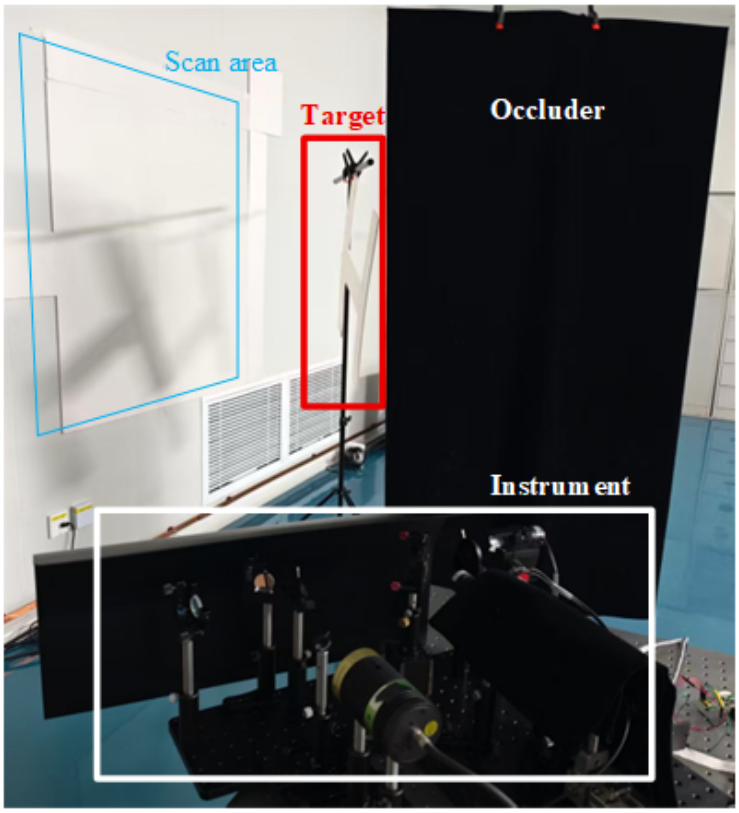

(**a**) Experimental scene setup

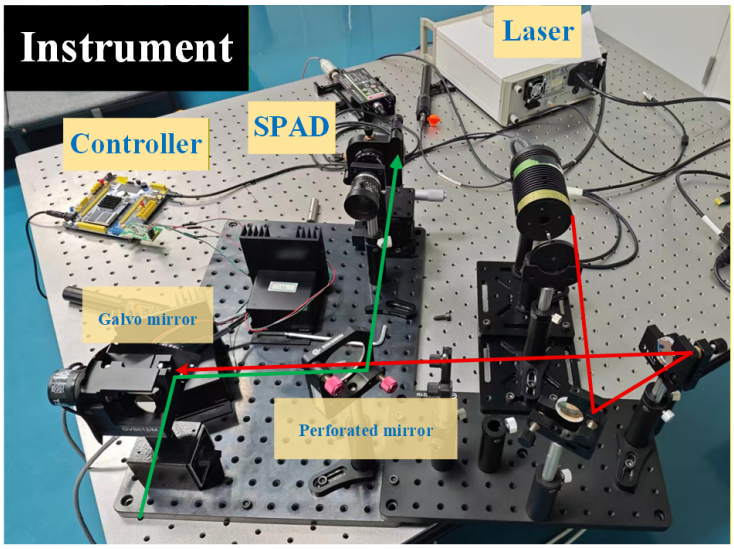

(**b**) Confocal light path and instrument

Figure 7. An illustration of the experimental system, where (**a**) shows the scene setup and (**b**) presents the confocal light path with the instrument, the green and red arrows denote the illumination and detection optical paths, respectively.

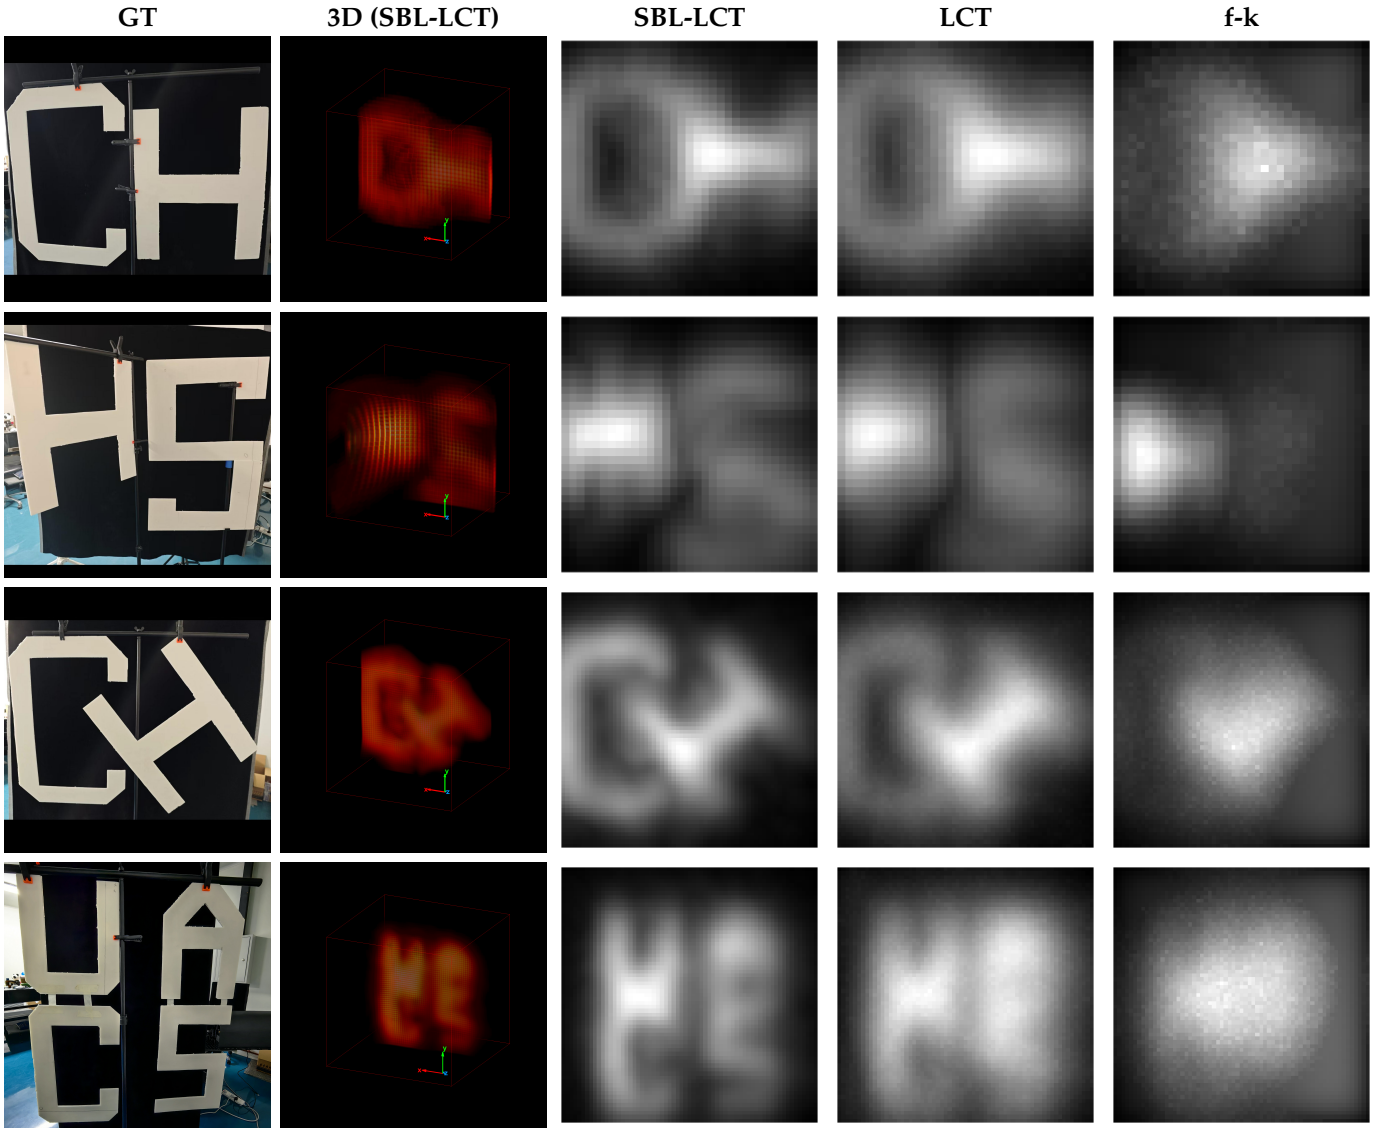

Figure 8. Reconstruction results on four target scenes: CH, SH, C with tilted H, and UCAS. Columns correspond to ground truth (GT), 3D visualization of SBL-LCT, SBL-LCT, LCT, and f-k.

We employ two no-reference evaluation metrics—the standard deviation (STD) and information entropy (Entropy)—to compare the reconstruction results of different algorithms. STD reflects pixel-intensity fluctuations, where a lower value generally indicates reduced background noise and smoother image regions. Entropy, on the other hand, measures the information richness and structural diversity of the reconstructed image, with higher values typically representing clearer details and stronger feature preservation. Since STD emphasizes noise suppression while entropy emphasizes detail retention, a desirable reconstruction must achieve a proper balance between the two. As summarized in Table 5, SBL maintains lower background variations while sustaining informative structural complexity, indicating that noise is effectively suppressed without over-smoothing target details. This balanced behavior leads to clearer object boundaries and more faithful spatial structures compared with LCT and f-k migration, demonstrating that SBL provides a more reliable compromise between background cleanliness and detail preservation for high-quality NLOS reconstruction.

Based on the quantitative results in Table 5, the LCT algorithm generally produces higher entropy values, indicating stronger retention of structural variations, yet its high STD scores reveal noticeable background fluctuations in the reconstructed images. The f-k method yields the lowest STD values and thus stronger noise suppression, but this comes

at the cost of reduced entropy and oversmoothed scene details. In comparison, SBL-LCT achieves a desirable compromise between the two metrics: background variations are reduced more effectively than in LCT, while informative structures are better preserved than in f-k reconstruction. Especially for the HS and UCAS scenes, SBL-LCT maintains detail-rich representations while noticeably reducing noise, leading to clearer boundaries and more faithful visual perception of the concealed shapes. It is worth noting that, in the NLOS context, a slightly lower entropy does not necessarily imply a loss of fine structures, as noise-induced high-frequency intensity fluctuations can also increase entropy values. In the UCAS scene, the entropy reduction observed for SBL-LCT primarily reflects effective suppression of background noise rather than the removal of meaningful structural details. These observations confirm that SBL-LCT achieves a favorable balance between background cleanliness and structural fidelity, enabling consistently improved reconstruction quality across different datasets.

Table 5. No-reference image quality assessment metrics (STD and Entropy).

Target	Algorithm	STD	Entropy
CH	SBL-LCT	0.215	7.43
	LCT	0.221	**7.54**
	f-k	**0.188**	7.18
HS	SBL-LCT	0.205	**7.31**
	LCT	0.196	7.29
	f-k	**0.188**	6.84
C + tilted H	SBL-LCT	0.230	7.45
	LCT	0.231	**7.51**
	f-k	**0.226**	7.35
UCAS	SBL-LCT	**0.234**	7.46
	LCT	0.263	**7.75**
	f-k	0.254	7.69

Note: Bold values indicate the best quantitative performance among the compared methods.

To further assess robustness under practical NLOS conditions, we evaluate the three reconstruction strategies on diverse hidden scenes, including both synthetic letter targets and real-object measurements involving complex surfaces and irregular geometry. Under challenging scenarios with weak photon returns, strong multipath interference, or tilted object configuration, SBL-LCT demonstrates improved adaptability by preserving target geometry and suppressing clutter-induced artifacts. The results show that SBL-based transient enhancement provides more physically reliable and informative input for geometric inversion, supporting stable and high-quality LCT reconstructions under a wide range of real-world constraints.

4.3. Discussion

We first discuss the role of IRF calibration in the proposed framework. In this work, the instrument response function (IRF) is obtained through a one-time experimental calibration using the same confocal NLOS setup and is consistently applied across both simulation and experimental evaluations. The proposed SBL-based transient enhancement operates under a given forward convolution model and does not attempt to explicitly invert or estimate the IRF.

While accurate IRF calibration can improve temporal localization and depth accuracy, the primary role of the SBL module is to suppress background fluctuations and enhance the relative temporal concentration of dominant photon returns. Accordingly, when IRF modeling errors do not introduce significant systematic timing shifts, the transient enhancement process

typically does not amplify instability arising from model mismatch. Substantial IRF inaccuracies, such as pronounced temporal shifts or severe shape mismatch, may still lead to biased peak localization and geometric errors, which represent a common limitation of convolution-based NLOS observation models. Incorporating IRF uncertainty into a joint or hierarchical Bayesian estimation framework therefore constitutes a promising direction for future work.

From the perspective of subsequent geometric inversion, the light-cone transform (LCT) primarily relies on accurate temporal localization of the main echo energy and on the separability of multipath components, rather than on strict recovery of an ideal impulse response. As a result, improvements achieved at the signal level mainly enhance reconstruction stability under the assumed geometric model, while violations of the geometric assumptions (e.g., strongly non-planar relay surfaces) are not compensated by the transient enhancement stage. This observation highlights the role of SBL as a signal-domain preprocessing module that complements, but does not replace, geometric modeling in NLOS reconstruction.

5. Conclusions

This work advances practical LCT-based NLOS imaging by addressing transient degradation at the signal front end. Through Bayesian sparsity-driven transient enhancement, the proposed SBL-LCT framework effectively restores temporally compact and geometrically consistent multipath returns, leading to improved spatial resolution, boundary sharpness, and robustness against IRF-induced blurring and photon-starved degradation. The ability to integrate SBL as an independent and lightweight preprocessing module ensures strong compatibility with existing acquisition hardware and reconstruction pipelines, enabling low-cost deployment and incremental upgrades for real-world systems.

Extensive experiments on diverse scenes, including complex reflective objects and tilted geometries, validate that transient quality remains a fundamental bottleneck in practical NLOS deployment and that enhancing it through probabilistic inference significantly strengthens reconstruction reliability under challenging conditions. Looking ahead, accelerating SBL inference via parallelization and efficient hardware implementations, as well as coupling with physics-informed learning models, may further extend the framework toward dynamic scenarios, multi-view fusion, and high-throughput computational sensing. Overall, this study establishes transient front-end enhancement as a promising and scalable direction for enabling stable, flexible, and high-fidelity NLOS visual sensing in realistic environments.

Author Contributions: Conceptualization, Y.T. and W.X.; methodology, Y.T.; software, Y.T.; validation, Y.T., W.X. and D.W.; formal analysis, N.Z.; investigation, S.C.; resources, S.C.; data curation, Y.T. and D.W.; writing—original draft preparation, Y.T.; writing—review and editing, Y.T.; visualization, P.G.; supervision, X.S.; project administration, S.C. and W.H.; funding acquisition, S.C., N.Z. and P.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China (grant numbers 62201568 and 62305375), the China Postdoctoral Science Foundation (grant number 2020M683600), the Strategic High Technology Innovation Project of the Chinese Academy of Sciences (grant number GQRC-19-19), the Key Program of the Chinese Academy of Sciences (grant number KGFZD-145-25-15), and the Natural Science Basic Research Program of Shaanxi Province (grant number 2025JC-YBQN-118). The APC was funded by the National Natural Science Foundation of China (grant numbers 62201568 and 62305375).

Data Availability Statement: The data presented in this study are not publicly available due to privacy, and access can be requested from [tianyuyuan18@mails.ucas.ac.cn] upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Velten, A.; Willwacher, T.; Gupta, O.; Bawendi, M.; Freeman, W.T.; Raskar, R. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nat. Commun.* **2012**, *3*, 745. [CrossRef]
2. O'Toole, M.; Lindell, D.B.; Wetzstein, G. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature* **2018**, *555*, 338–341. [CrossRef] [PubMed]
3. Velten, A.; Wu, D.; Jarabo, A.; et al. Femto-photography: Capturing and visualizing the propagation of light. *ACM Trans. Graph.* **2013**, *32*, 44. [CrossRef]
4. Heide, F.; Xiao, L.; Heidrich, W.; Hullin, M.B. Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In Proceedings of the CVPR, Columbus, OH, USA, 23–28 June 2014; pp. 3222–3229.
5. Laurenzis, M.; Velten, A. Nonline-of-sight laser gated viewing of scattered photons. *Opt. Eng.* **2014**, *53*, 023102. [CrossRef]
6. Buttafava, M.; Zeman, J.; Tosi, A.; Eliceiri, K.; Velten, A. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Opt. Express* **2015**, *23*, 20933–20946. [CrossRef] [PubMed]
7. Gariepy, G.; Krstajic, N.; Henderson, R.; Li, C.; Thomson, R.R.; Buller, G.S.; Heshmat, B.; Raskar, R.; Leach, J.; Faccio, D. Single-photon sensitive light-in-flight imaging. *Nat. Commun.* **2015**, *6*, 6021. [CrossRef] [PubMed]
8. O'Toole, M.; Heide, F.; Lindell, D.B.; Zang, K.; Diamond, S.; Wetzstein, G. Reconstructing transient images from single-photon sensors. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 1539–1547.
9. Li, T.; Luo, Y.; Zhao, K. Weak non-line-of-sight target echoes extraction without accumulation. *Opt. Express* **2023**, *31*, 36209–36218. [CrossRef] [PubMed]
10. Feng, Y.; Cui, X.; Meng, Y. Non-line-of-sight imaging at infrared wavelengths using a superconducting nanowire single-photon detector. *Opt. Express* **2023**, *31*, 42240–42254. [CrossRef] [PubMed]
11. Yang, W.; Zhang, C.; Jiang, W.; Zhang, Z.; Sun, B. None-line-of-sight imaging enhanced with spatial multiplexing. *Opt. Express* **2022**, *30*, 5855–5867. [CrossRef] [PubMed]
12. Li, Z.; Liu, X.; Wang, J.; Shi, Z.; Qiu, L.; Fu, X. Fast non-line-of-sight imaging based on first photon event stamping. *Opt. Lett.* **2022**, *47*, 1928–1931. [CrossRef] [PubMed]
13. Lindell, D.B.; Wetzstein, G.; O'Toole, M. Wave-based non-line-of-sight imaging using fast f-k migration. *ACM Trans. Graph.* **2019**, *38*, 1–13. [CrossRef]
14. Liu, X.; Guillén, I.; La Manna, M. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature* **2019**, *572*, 620–623. [CrossRef] [PubMed]
15. Liu, X.; Wang, J.; Xiao, L.; Shi, Z.; Fu, X.; Qiu, L. Non-line-of-sight imaging with arbitrary illumination and detection pattern. *Nat. Commun.* **2023**, *14*, 3230. [CrossRef] [PubMed]
16. Young, S.I.; Lindell, D.B.; Girod, B.; Taubman, D.; Wetzstein, G. Non-line-of-sight surface reconstruction using the directional light-cone transform. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Virtual, 14-19 June 2020; pp. 1407–1416.
17. Jin, C.; Hu, X.; Cao, X.; Zhang, Z.J.; Zhao, Y. Richardson–Lucy deconvolution of time histograms for high-resolution non-line-of-sight imaging based on a back-projection method. *Opt. Lett.* **2018**, *43*, 5885–5888. [CrossRef] [PubMed]
18. Liu, X.; Wang, J.; Li, Z.; Shi, Z.; Fu, X.; Qiu, L. Non-line-of-sight reconstruction with signal–object collaborative regularization. *Light. Sci. Appl.* **2021**, *10*, 198. [CrossRef] [PubMed]
19. Ding, R.; Ye, J.; Gao, Q.; Xu, F.; Duan, Y. Curvature regularization for non-line-of-sight imaging from under-sampled data. *IEEE Trans. Pattern Anal. Mach. Intell.* **2024**, *46*, 8474–8485. [CrossRef] [PubMed]
20. Xu, W.; Chen, S.; Wang, D.; Tian, Y.; Zhang, N.; Hao, W.; Su, X. Feature enhanced non-line-of-sight imaging using graph model in latent space. *Opt. Laser Technol.* **2025**, *181*, 111538. [CrossRef]
21. Xu, W.; Chen, S.; Tian, Y.; Wang, D.; Su, X. Fast non-line-of-sight imaging based on product-convolution expansions. *Opt. Lett.* **2022**, *47*, 4680–4683. [CrossRef] [PubMed]
22. Li, Y.; Zhang, Y.; Ye, J.; Xu, F.; Xiong, Z. Deep non-line-of-sight imaging from under-scanning measurements. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 59095–59106.
23. Ye, J.; Sun, Y.; Li, W.; Zeng, J.W.; Hong, Y.; Li, Z.P.; Huang, X.; Xue, X.; Yuan, X.; Xu, F.; et al. Real-time non-line-of-sight computational imaging using spectrum filtering and motion compensation. *Nat. Comput. Sci.* **2024**, *4*, 920–927. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.