

Dual-branch Graph Feature Learning for NLOS Imaging

Xiongfei Su^{1,2,3*}, Tianyi Zhu^{2*}, Lina Liu², Zheng Chen⁴, Yulun Zhang⁴,
Siyuan Li^{1,3}, Juntian Ye⁵, Feihu Xu⁵, Xin Yuan^{3†}

¹Zhejiang University, Hangzhou, China, e-mail: xsuac@zju.edu.cn

²China Mobile Research Institute, Beijing, China

³Westlake University, Hangzhou, China

⁴Shanghai Jiao Tong University, Shanghai, China

⁵University of Science and Technology of China, Anhui, China

Abstract

The domain of non-line-of-sight (NLOS) imaging is advancing rapidly, offering the capability to reveal occluded scenes that are not directly visible. However, contemporary NLOS systems face several significant challenges: (1) The computational and storage requirements are profound due to the inherent three-dimensional grid data structure, which restricts practical application. (2) The simultaneous reconstruction of albedo and depth information requires a delicate balance using hyperparameters in the loss function, rendering the concurrent reconstruction of texture and depth information difficult. This paper introduces the innovative methodology, DG-NLOS, which integrates an albedo-focused reconstruction branch dedicated to albedo information recovery and a depth-focused reconstruction branch that extracts geometrical structure, to overcome these obstacles. The dual-branch framework segregates content delivery to the respective reconstructions, thereby enhancing the quality of the retrieved data. To our knowledge, we are the first to employ the GNN as a fundamental component to transform dense NLOS grid data into sparse structural features for efficient reconstruction. Comprehensive experiments demonstrate that our method attains the highest level of performance among existing methods across synthetic and real data.

Introduction

Non-line-of-sight (NLOS) imaging transcends the limitations of rectilinear propagation by examining the diffuse reflections from a relay surface to image occluded objects (Faccio, Velten, and Wetzstein 2020; Maeda et al. 2019). This technique has extensive applications across various domains, including medical imaging, autonomous driving, and robotic vision. The swift progression in photon-sensitive sensors and sophisticated imaging algorithms has engendered solutions that span a comprehensive range of ray propagation paradigms, such as occluder-based recovery (Xu et al. 2018; Saunders, Murray-Bruce, and Goyal 2019), Time-of-Flight (ToF) methodologies (Velten et al. 2012; Heide et al. 2014), wavefront shaping (Cao et al.

*These authors contributed equally.

†Corresponding author.

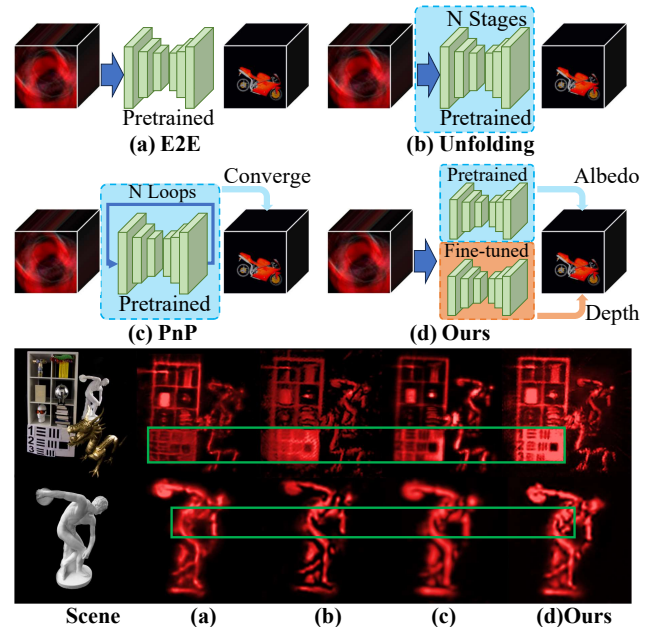


Figure 1: Various types of Deep NLOS reconstruction. Blue arrows denote the transform with physical models. The two real scene reconstructions are compared, where green rectangles highlight the superior performance of DG-NLOS.

2022), and alternative non-optical techniques (Lindell, Wetzstein, and Koltun 2019; Scheiner et al. 2020). Several methods exist for reconstructing hidden scenes, including Fermat flow (Xin et al. 2019), temporal focusing (Pediredla, Dave, and Veeraraghavan 2019), optimization-based reconstruction (Heide et al. 2019; Tsai, Sankaranarayanan, and Gkioulekas 2019; Ye et al. 2021), deconvolution (O’Toole et al. 2018; Young et al. 2020), wave-based techniques (Liu et al. 2019; Lindell, Wetzstein, and O’Toole 2019), and deep learning (Grau Chopite et al. 2020; Nam et al. 2021; Pei et al. 2021; Liu et al. 2022; Ye et al. 2024).

NLOS reconstructions commonly face the following challenges: (i) **Low signal-to-noise ratio (SNR)**: as only a small fraction of scattered photons are captured (Zhu et al. 2023; Li et al. 2024). (ii) **Ill-posed solution**: Signals travel through

multiple paths to reach the receiver, causing time and phase differences between the paths (Wt, Zhangt, and Mou 2021), resulting in multipath interference. **(iii) Complex signal detection and estimation:** NLOS environments require complex signal processing algorithms (Yu et al. 2023; Liu et al. 2023), channel estimation methods, and modulation/demodulation techniques to handle the multipath effects, angle spread, and time-varying characteristics.

Three deep learning-based frameworks are proposed to solve these problems, shown in Fig. 1: **(a) End-to-end (E2E):** (Chen et al. 2020) introduces a feature embedding method for NLOS tasks, including reconstruction, imaging, classification, and object detection. (Grau Chopite et al. 2020) generates training images using noise and rendering models for transient NLOS imaging with Time-of-Flight (ToF) technology, modifying U-Net with a 3D tensor to convert transient measurements into depth maps. (Metzler et al. 2020) use the spectral estimation theory for NLOS correlography, employing a deep neural network to handle noisy phase retrieval and detect hidden objects through indirect reflections. (Li et al. 2023) introduces Transformer to extract global and local relationship. **(b) Deep Unfolding** and **(c) Plug-and-Play (PnP)** explore hybrid approaches combining the benefits of both optimization and deep learning paradigms. Deep Unfolding (Su et al. 2023) trains a sequence of interconnected small sub-networks, emulating the iterative process of traditional optimization. This method “unfolds” optimization iterations with end-to-end training and offers interpretability by mapping stages to iterations. PnP algorithms (Ye et al. 2024) use pre-trained deep denoising networks as priors in iterative algorithms, requiring no additional training for new applications. They are adaptable to various systems and ensure algorithm stability under certain loss functions and denoiser conditions.

The **motivations** of this work are: (1) Previous methods rely on regular grids or sequences, whereas NLOS requires flexible data representation due to its dependence on *geometrical structure*. Regular grids often lead to redundancy and overhead. (2) Objects can be seen as *sparse* parts; for instance, a marble statue’s head, upper body, and limbs form a graph structure connected by joints. Graph neural network (GNN) methods can address this sparse structure. (3) The *coupled reconstruction* of albedo and depth information requires balanced optimization via hyper-parameters in a mixed loss function. Specific contributions are as follows:

- Firstly, we introduce a dual-branch graph learning framework DG-NLOS for NLOS reconstruction, incorporating a two-stage training mechanism, shown in Fig. 1(d), which decouples albedo and depth reconstruction to achieve the best results respectively.
- Secondly, to effectively extract geometrical features, we develop a graph block and a channel fusion block specifically tailored for NLOS feature, where dense grid data is converted into a sparse graph structure of the objects.
- Lastly, extensive experiments validate DG-NLOS’s robustness across various scenarios, achieving state-of-the-art performance including mainstream real-world data, with less GPU memory.




Figure 2: A schematic diagram of the C-NLOS system.

Related Work

Deep NLOS Reconstruction: Compared to traditional algorithms (Nam et al. 2021; Pei et al. 2021), deep learning algorithms can learn scene priors, extract features, and reconstruct hidden objects. Hybrid methods (Meng, Jalali, and Yuan 2020; Mou, Wang, and Zhang 2022; Wt, Zhangt, and Mou 2021; Zhang and Ghanem 2018) introduce several multi-scale iterative model-guided unfolding networks for Confocal NLOS (C-NLOS) reconstruction. Recent work (Su et al. 2023; Ye et al. 2024) transforms traditional optimization into iterative learning but focuses solely on reconstruction performance without testing on corrupted cases. Inspired by but different from previous researches (Shen et al. 2021; Liu et al. 2021; Wt, Zhangt, and Mou 2021; Mou, Wang, and Zhang 2022; Fujimura et al. 2023; Zhu et al. 2023), our proposed DG-NLOS develops a multi-scale graph-based network for C-NLOS reconstruction.

Graph Learning: GNNs are designed to process graph data by establishing long-range correlations in non-Euclidean space. Micheli (Micheli 2009) introduced the spatial graph convolutional network with nonrecursive layers. Instead of directly aggregating features from neighboring nodes, EdgeConv (Wang and Solomon 2021) obtains local neighborhood information by subtracting the central vertex’s feature from that of neighboring vertices. (Yang et al. 2020) introduced a highway GNN for user geo-location in social media graphs, using ‘highway’ gates to enhance gradient flow. Their research indicated a decline in performance beyond 6 layers. (Xu et al. 2017) proposed a *Jump Knowledge Network* to determine graph neighbors for each node based on the graph’s structure. GNN applications in computer vision (Landrieu and Simonovsky 2018) include point cloud classification, scene graph generation, and action recognition. Point clouds are 3D points typically collected by LiDAR scans. GCNs have been used for classifying and segmenting point clouds (Landrieu and Simonovsky 2018; Wang et al. 2019).

Physical Forward Model

The C-NLOS imaging system typically consists of a scanning pulsed laser and a single photon time-resolved detector,

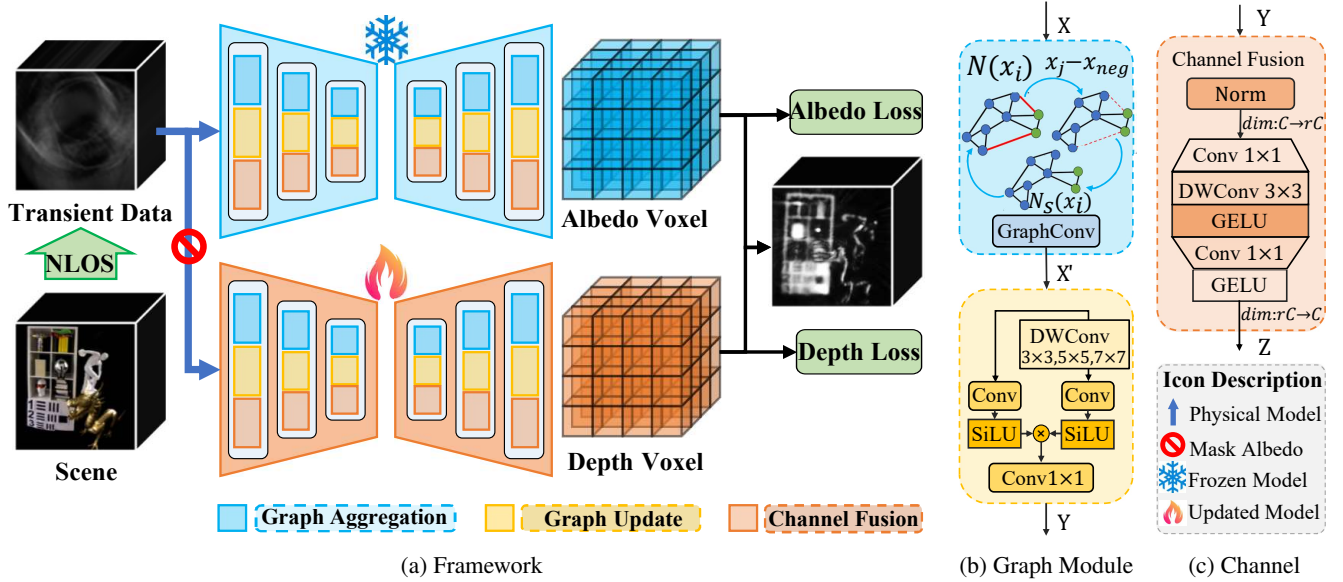


Figure 3: (a) **Structure of the two-stage learning pipeline:** The loss functions including Albedo and Depth are calculated in triple scales, respectively. Finally, the output voxels of two branches are combined for the optimized reconstruction. (b) **Structure of graph module.** (c) **Structure of channel fusion.** Each module adopts the Resnet skip connection mechanism.

which focuses on the same points on a diffuse reflective wall. As shown in Fig. 2, the directly illuminated points $(\xi, \kappa, 0)$ on the visible wall are considered as the sampling points, and the first diffuse reflection ray propagates to the points $(x, y, z) \in \Omega$ on the hidden object. The second reflected wave from the object, $\phi(x, y, z)$, and the third reflections, which have the same reflective position $(\xi, \kappa, 0)$ as the first one, result in a time-resolved diffusive intensity $d(\xi, \kappa, t)$ received by the detector, where t represents the time of photon flight between the first and third reflections. Finally, a three-dimensional (3D) light transient $d(\xi, \kappa, t)$ is measured by an $m \times m$ array sampling. As derived in (O’Toole et al. 2018), the 3D continuous signal is formulated as

$$d(\xi, \kappa, t) = \iiint_{\Omega} \frac{1}{r^4(x-\xi, y-\kappa, z)} \phi(x, y, z) \delta(2r - ct) dx dy dz, \quad (1)$$

$$r(x - \xi, y - \kappa, z) = \sqrt{(x - \xi)^2 + (y - \kappa)^2 + z^2},$$

where the Dirac delta function δ models the light propagation, c is the speed of light. Note that r is the distance between the sampling points and the corresponding points on the surface of the hidden object. Combining all the detected photon arrival events into a single histogram results in a discrete inhomogeneous variable as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{b}, \quad (2)$$

where $\mathbf{y} \in \mathbb{R}^{n_x n_y n_t}$ represents the discretized measurement by scanning point (n_x, n_y) with respect to the discretized time bin n_t . $\mathbf{A} \in \mathbb{R}^{n_x n_y n_t \times n_x n_y n_z}$ is the discretized version of the volumetric Albedo model in (1). $\mathbf{x} \in \mathbb{R}^{n_x n_y n_z}$ represents the discretized Albedo of the hidden object. $\mathbf{b} \in \mathbb{R}^{n_x n_y n_t}$ denotes the dark count of the detector and background noise (Bronzi et al. 2016).

Proposed DG-NLOS

Graphing is an effective (Jia et al. 2022) technique for preserving structural features and propagation knowledge during occluded object reconstruction. A key component of our DG-NLOS is the graph blocks, which generate learnable parameters and guide the model in the NLOS reconstruction process. Here, we outline the DG-NLOS pipeline:

Network Framework

Following learning-based methods (Chen et al. 2020; Mu et al. 2022; Li et al. 2023), we transform spatial-temporal transient data to the 3D spatial domain grid feature \mathbf{F} using a physics-based prior in the feature transform layer following (Chen et al. 2020). The proposed DG-NLOS features a dual-branch (Zhou et al. 2023) symmetric architecture for efficient hierarchical representation learning, as shown in Figure 3a. Both the encoder and decoder networks have three scales, with the decoder progressively reconstructing hidden objects from the input grid features \mathbf{F} .

The upper albedo branch extracts albedo representation by disentangling features into texture and style components. We mask the albedo of grid feature \mathbf{F} by retaining the maximum value voxel along the z axis to generate depth-focused grid feature \mathbf{F}' . The lower depth branch is trained to suppress depth-irrelevant structures with the upper branch parameters frozen. Finally, the depth-focused voxel provides depth information and albedo-focused voxel provides albedo information to render the depth-focused voxel with the corresponding albedo value, followed by (Ye et al. 2024) Algorithm 2. Next, we describe graph block and core modules.




Figure 4: Graph construction and selection. Cyan edges connect neighbor vertices, while red edges denote negative connections. From (a) to (b), original neighbor vertices 7 and 5 are removed due to closer distance with negative vertices.

Graph Block

The overall process of graph construction is defined as:

$$\mathcal{G} = G(\mathbf{F}) \quad (3)$$

Then, we explore the graph feature representations of NLOS. The graph aggregation block consists of

$$\begin{aligned} \mathcal{G}' &= F(\mathcal{G}, \mathcal{W}) \\ &= \text{Update}(\text{Aggregate}(\text{Select}(\mathcal{G}), W_{\text{agg}}), W_{\text{update}}), \end{aligned} \quad (4)$$

where W_{agg} and W_{update} are the learnable weights of the aggregation and update operations, respectively. Concretely, the aggregation operation computes the representation of a vertex by aggregating features of neighbor vertices, and the update operation further merges the aggregated features.

(1) Graph Construction For a grid feature voxel $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$, we divided it into N grids. By transforming each grid into a feature vector $\mathbf{x}_i \in \mathbb{R}^D$, we have $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$ where D is the feature dimension and $i = 1, 2, \dots, N$. These features can be viewed as a set of unordered vertices which are denoted as $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$. For each vertex v_i , we find its K nearest neighbors $\mathcal{N}(v_i)$ and add an edge e_{ji} directed from v_j to v_i for all $v_j \in \mathcal{N}(v_i)$. In this problem, we use \mathbf{x} instead of v for convenience, and $\mathcal{N}(\mathbf{x}_i)$ is the set of NLOS voxel neighbor vertices of \mathbf{x}_i , shown in Fig. 4(a). Here we adopt dynamic k -NN, (Li et al. 2019) to construct the dynamic edges between points at every layer in the feature space. Then we obtain a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{E} denotes all the edges.

(2) Graph Selection To extract more effective information between vertices by aggregating features from its neighbor vertices, we define negative vertices \mathbf{x}_{neg} from the constructed graph and apply them to yield selected graph, shown in Fig. 4(b):

$$\mathcal{N}_s(\mathbf{x}_i) = \max_k(\{\mathbf{x}_j - \mathbf{x}_{\text{neg}} | \mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)\}), \quad (5)$$

where \mathbf{x}_{neg} are the average pooling value of 8 black vertices of the NLOS voxel, being the vertices with the least effective information. In our implementation, we compute a pairwise distance matrix in feature space and then take the k vertices of maximum distance from \mathbf{x}_{neg} as the selected graph.

(3) Graph Aggregation. Given the graph feature $\mathbf{X}' \in \mathbb{R}^{D \times N}$ in Fig. 3b, which is resize to $\mathbf{X}' \in \mathbb{R}^{D \times \sqrt{N} \times \sqrt{N}}$ for the aggregation operation:

$$\mathbf{X}'' = \text{GraphConv}([\mathbf{X}, \mathbf{X}']W_{\text{in}})W_{\text{out}} + \mathbf{X}, \quad (6)$$

where $\mathbf{X}' \in \mathbb{R}^{kD \times N}$ is calculated by Eq. (7), W_{in} and W_{out} are the weights of fully-connected layers. After these layers, we resize the graph feature for the next operation.

$$\mathbf{x}'_i = [\mathbf{x}_i, \{\mathbf{x}_j - \mathbf{x}_i | \mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)\}], \quad (7)$$

and we test four types of GCN (Li et al. 2019) *ResEdgeConv*, *GraphSAGE*, *GIN* and *Max-Relative* respectively.

(4) Graph Update We Further employ three different update operation with depth-wise convolutions (DWConv) layers in dilation ratios $d \in \{1, 2, 3\}$ in parallel to capture different weights respectively: Given the aggregated feature $\mathbf{X}'' \in \mathbb{R}^{D \times N}$, $\text{DW}_{5 \times 5, d=1}$ is first applied for low-order features; then, the output is factorized into $\mathbf{X}_l \in \mathbb{R}^{D_l \times N}$, $\mathbf{X}_m \in \mathbb{R}^{D_m \times N}$, and $\mathbf{X}_h \in \mathbb{R}^{D_h \times N}$ along the channel dimension, where $D_l + D_m + D_h = D$; afterward, \mathbf{X}_m and \mathbf{X}_h are assigned to $\text{DW}_{5 \times 5, d=2}$ and $\text{DW}_{7 \times 7, d=3}$, respectively, while \mathbf{X}_l serves as identical mapping; finally, all the branches can be updated in parallel and concatenated as the final value: $\mathbf{X}_D = \text{Concat}(\mathbf{X}_{l, 1:D_l}, \mathbf{X}_m, \mathbf{X}_h)$. Multi-order update operation allows the model to update information in multiple representation subspaces for feature diversity.

Then, to alleviate over-smoothing phenomenon in deep GCNs (Li, Han, and Wu 2018; Oono and Suzuki 2020), we introduce more feature transformations and nonlinear activations in our block. We utilize the gating operation to adaptively fuse priors and structure features:

$$\mathbf{Y} = \text{SiLU}(\text{Conv}_{1 \times 1}(\mathbf{X}')) \otimes \text{SiLU}(\text{Conv}_{1 \times 1}(\text{DW}(\mathbf{X}_D))).$$

The nonlinear activation function SiLU is inserted after graph aggregation to avoid layer collapse.

Channel Fusion Module

Prevalent architectures perform channel fusion mainly by two linear projections, e.g., 2-layer channel-wise MLP (Dosovitskiy et al. 2021; Tolstikhin et al. 2021) with a expand ratio r or the MLP with a 3×3 DWConv between (Pan et al. 2022; Pan, Cai, and Zhuang 2022). Due to information redundancy across channels (Cao et al. 2019; Tan and Le 2019; Wang et al. 2020), vanilla MLP requires a number of parameters (r default to 4 or 8) to achieve expected performance. To address this issue, most current methods directly insert a channel enhancement module, e.g., SE module (Hu, Shen, and Sun 2018), into MLP. Unlike these designs that require an additional MLP bottleneck, we introduce a lightweight channel fusion module $\text{CF}(\cdot)$ to adaptively reallocate channel-wise features in high-dimensional hidden spaces. As shown in Fig. 3c:

$$\begin{aligned} \mathbf{Z}' &= \text{GELU}(\text{DW}_{3 \times 3}(\text{Conv}_{1 \times 1}(\text{Norm}(\mathbf{Y})))) , \\ \mathbf{Z} &= \text{GELU}(\text{Conv}_{1 \times 1}(\mathbf{Z}')) + \mathbf{Y}. \end{aligned} \quad (8)$$

Concretely, GELU is used to gather and reallocate channel-wise information with the complementary interactions.

Decoupled Optimization Strategy

It has been demonstrated in (Lu et al. 2020; Chen et al. 2021) that structural information is closely associated with depth, whereas elements like style and lighting act as interference




Figure 7: Qualitative visualization of two color samples in $256 \times 256 \times 512$ size reconstructed by the proposed DG-NLOS and previous strong baseline methods. “Meas”, and “GT” stands for measurement and ground truth respectively.

Methods	Venue	Type	PSNR(dB) \uparrow	SSIM \uparrow	RMSE \downarrow	TIME(S) \downarrow	Mem(G) \downarrow
<i>SP</i>	Nat. Com.'2021		18.34	0.76	0.65	-	-
<i>FBP</i>	Siggraph'2017	Non-	17.63	0.62	0.73	0.65	15.6
<i>LCT</i>	Nature'2018	Learning-	19.02	0.85	0.59	0.89	17.7
<i>FK</i>	TOG'2019	Based	23.51	0.87	0.54	1.63	21.0
<i>Phasor</i>	Nature'2019		24.53	0.90	0.47	1.92	10.8
<i>LFE</i>	TOG'2020		27.54	0.89	0.08	0.96	5.0
<i>MIMU</i>	PG'2023	Learning-	28.77	0.92	0.06	1.12	5.8
<i>NLOST</i>	CVPR'2023	Based	28.17	0.90	0.07	0.24	20.3
<i>PnP</i>	TOG'2024		25.37	0.83	0.22	3.70	13.7
DG-NLOS	Ours		29.93	0.92	0.04	0.18	9.6

Table 1: Quantitative results on the motorbikes dataset.

Gray Scale. To compare the C-NLOS reconstruction results fairly in the same size with existing methods, we evaluate the performance in quality and quantity on the same test data. We quantitatively evaluate the results among these baseline approaches in Tab. 1, where we can see that the accuracy of the proposed DG-NLOS exceeds existing methods.

RMSE Evaluation Compared with the ground truth, the proposed method generates the least distance error and sharpest edges. The RMSE in Tab. 1 shows that it especially outperforms existing methods thanks to the decoupled albedo-depth training strategy.




Figure 8: Depth estimation on real dataset. Brighter regions are closer to the observer than darker regions.

Modules	PSNR	SSIM	Params. (M)	FLOPs (G)
Baseline	28.37	0.86	15.46	146.06
+Select(\cdot)	28.68	0.89	15.46	146.92
+Update(\cdot)	29.41	0.90	15.84	147.16
+Fusion(\cdot)	29.93	0.92	16.04	148.02

Table 2: Ablation of modules: Baseline uses *ResEdgeConv*, MLP as Update(\cdot) and Fusion(\cdot) following (Li et al. 2019).

Results on Real Data

Albedo Reconstruction. The results in Fig. 9 using data from (O’Toole et al. 2018) show that DG-NLOS restores hidden details, high-frequency textures, and weak illumination, achieving superior quality. DG-NLOS reconstructs object surfaces with sharper, cleaner details, including the bike front wheel, dragon tail, and bookshelf, surpassing previous methods. In the bottom row of Fig. 9, the statue is hard to reconstruct with existing methods due to its distance from foreground objects. Our graph structure dynamically adjusts geometric weights, enhancing object perception.

Depth Estimation. Most deep learning methods can achieve good results on simulated data sets, shown in Tab. 1, but perform poorly on the real data sets, shown in Fig. 8. We show the depth estimation from Phasor for comparison. LFE only estimates outlines of some objects; MIMU makes an error in estimating the front and rear distance of an object because the bookshelf should be between the closest dragon and furthest statue; NLOST estimates depth in details, but fails to eliminate the influence of texture with number on the board. Our DG-NLOS can revise it with fewer artifacts due to the decoupled strategy analyzed before.

Ablation Study

Effects of Graph and Channel Block. We first add the graph selection module and the graph update module to the baseline framework in Tab. 2, which leads a 1.04dB increase in performance. We found that all proposed modules yield improvements with favorable costs.

Comparison with alternative Graph Aggregation. We further demonstrate the superiority of our Graph Aggregation by replacing it with three popular graph aggregation methods. As represented in Tab. 3, the *GIN* and *GraphSAGE* lead to performance degradation. Considering the balance of




Figure 9: Reconstructions from real measurements at 32ps time resolution. Zoom in for more details.

Method	PSNR	SSIM	Params. (M)	FLOPs (G)
(a) <i>GIN</i>	27.37	0.82	15.35	140.58
(b) <i>GraphSAGE</i>	27.68	0.87	16.38	151.30
(c) <i>Max-Relative</i>	28.42	0.91	16.00	142.62
(d) <i>ResEdgeConv</i>	29.93	0.92	16.04	148.02

Table 3: Comparison with various Graph Aggregation.

Strategy	Albedo		Depth	
	PSNR	SSIM	RMSE	MAD
Single Branch	28.28	0.88	0.07	0.04
Depth First	29.65	0.90	0.03	0.02
Albedo First	29.93	0.92	0.04	0.02

Table 4: Various training strategy. “Single Branch” adopts albedo-depth mixed loss function in (Chen et al. 2020).

performance and overload, *ResEdgeConv* (Li et al. 2019) is used in our final version.

Effects of dual-branch framework. We test the single branch in the same architecture and the training sequence in Tab. 4 and found that training albedo branch first and then training the depth branch can achieve the best performance.

Multi-scale Loss Function. Tab. 5 left part shows how the multi-scale loss function influences the performance. We can find that performance increases with the larger number of

Scale	PSNR	SSIM	Type	PSNR	SSIM
1	28.41	0.89	L1	29.93	0.92
2	28.85	0.90	MSE	26.90	0.86
3	29.93	0.92	L1+MSE	28.63	0.91

Table 5: Various loss scale (left) and types (right).

loss scales, demonstrating the effectiveness of the proposed multi-scale loss framework, which takes advantage of the global content awareness and local content interaction in different scales to enhance the perception of the weak signal. Furthermore, we investigate the effectiveness of the “L1” by replacing it with “MSE” and “L1+MSE” in Tab. 5 right part, showing the effect of our designs.

Conclusion

To our knowledge, we are the first to introduce a graph-based architecture, DG-NLOS, designed to transform previous 3D grid features to flexible graph features, ensuring the reconstruction process is sparser than previous work. This method not only efficiently extracts geometric information, but also greatly reduces overload. The visual results and quantitative evaluations demonstrate that our proposed DG-NLOS consistently surpasses existing learning-based and nonlearning-based methods. In the future, we aim to broaden the model’s applicability to encompass a wider range of corruptions, generating universal and robust models.

Acknowledgements

This work was supported by The National Key R&D Program of China (grant number 2024YFF0505603, 2024YFF0505600), the National Natural Science Foundation of China (grant number 62271414), Zhejiang Outstanding Youth Science Foundation (grant number LR23F010001), Zhejiang “Pioneer” and “Leading Goose” R&D Program (grant number 2024SDXHDX0006, 2024C03182), the Key Project of Westlake Institute for Optoelectronics (grant number 2023GD007), the 2023 International Sci-tech Cooperation Projects of the “Innovation Yongjiang 2035” R&D Program (grant number 2024Z126).

References

- Arellano, V.; Gutierrez, D.; and Jarabo, A. 2017. Fast back-projection for non-line-of-sight reconstruction. In *ACM SIG-GRAPH 2017 Posters*, 1–2. ACM.
- Bronzi, D.; Villa, F.; Tisa, S.; Tosi, A.; and Zappa, F. 2016. SPAD Figures of Merit for Photon-Counting, Photon-Timing, and Imaging Applications: A Review. *IEEE Sensors Journal*, 16(1): 3–12.
- Cao, R.; de Goumoens, F.; Blochet, B.; Xu, J.; and Yang, C. 2022. High-resolution non-line-of-sight imaging employing active focusing. *Nature Photonics*, 1–7.
- Cao, Y.; Xu, J.; Lin, S.; Wei, F.; and Hu, H. 2019. GC-Net: Non-Local Networks Meet Squeeze-Excitation Networks and Beyond. In *ICCVW*, 1971–1980.
- Chen, W.; Wei, F.; Kutulakos, K. N.; Rusinkiewicz, S.; and Heide, F. 2020. Learned feature embeddings for non-line-of-sight imaging and recognition. *ToG*, 39(6): 1–18.
- Chen, X.; and He, K. 2020. Exploring Simple Siamese Representation Learning. ArXiv:2011.10566 [cs].
- Chen, X.; Wang, Y.; Chen, X.; and Zeng, W. 2021. S2R-DepthNet: Learning a Generalizable Depth-specific Structural Representation. ArXiv:2104.00877 [cs].
- Cui, Y.; Tao, Y.; Bing, Z.; Ren, W.; Gao, X.; Cao, X.; Huang, K.; and Knoll, A. 2023. Selective frequency network for image restoration. In *ICLR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Faccio, D.; Velten, A.; and Wetzstein, G. 2020. Non-line-of-sight imaging. *Nature Reviews Physics*, 2(6): 318–327.
- Fujimura, Y.; Kushida, T.; Funatomi, T.; and Mukaigawa, Y. 2023. NLOS-NeuS: Non-line-of-sight Neural Implicit Surface. In *2023 ICCV*, 10498–10507.
- Grau Chopite, J.; Hullin, M. B.; Wand, M.; and Iseringhausen, J. 2020. Deep non-line-of-sight reconstruction. In *CVPR*, 960–969.
- Heide, F.; O’Toole, M.; Zang, K.; Lindell, D. B.; Diamond, S.; and Wetzstein, G. 2019. Non-line-of-sight imaging with partial occluders and surface normals. *ToG*, 38(3): 1–10.
- Heide, F.; Xiao, L.; Heidrich, W.; and Hullin, M. B. 2014. Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors. In *CVPR*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *CVPR*, 7132–7141.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *ECCV 2022*, 709–727. Springer.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv:1412.6980*.
- Landrieu, L.; and Simonovsky, M. 2018. Large-scale point cloud semantic segmentation with superpoint graphs. In *CVPR*, 4558–4567.
- Li, G.; Muller, M.; Thabet, A.; and Ghanem, B. 2019. Deep-gcns: Can gcns go as deep as cnns? In *ICCV*, 9267–9276.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*, 3538–3545.
- Li, Y.; Peng, J.; Ye, J.; Zhang, Y.; Xu, F.; and Xiong, Z. 2023. NLOST: Non-Line-of-Sight Imaging With Transformer. *CVPR*.
- Li, Y.; Zhang, Y.; Ye, J.; Xu, F.; and Xiong, Z. 2024. Deep Non-line-of-sight Imaging from Under-scanning Measurements. *NIPS*, 36.
- Lindell, D. B.; Wetzstein, G.; and Koltun, V. 2019. Acoustic non-line-of-sight imaging. In *CVPR*, 6780–6789.
- Lindell, D. B.; Wetzstein, G.; and O’Toole, M. 2019. Wave-based non-line-of-sight imaging using fast f-k migration. *ToG*, 38(4): 116:1–116:13.
- Liu, P.; Yu, Y.; Pan, Z.; Peng, X.; Li, R.; Wang, Y.; Yu, J.; and Li, S. 2022. HiddenPose: Non-Line-of-Sight 3D Human Pose Estimation. In *2022 IEEE ICCP*, 1–12. IEEE.
- Liu, X.; Guillén, I.; La Manna, M.; Nam, J. H.; Reza, S. A.; Huu Le, T.; Jarabo, A.; Gutierrez, D.; and Velten, A. 2019. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature*, 572(7771): 620–623.
- Liu, X.; Wang, J.; Li, Z.; Shi, Z.; Fu, X.; and Qiu, L. 2021. Non-line-of-sight reconstruction with signal-object collaborative regularization. *Light: Science & Applications*, 10(1).
- Liu, X.; Wang, J.; Xiao, L.; Shi, Z.; Fu, X.; and Qiu, L. 2023. Non-line-of-sight imaging with arbitrary illumination and detection pattern. *Nature Communications*, 14(1): 3230.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Lu, K.; Barnes, N.; Anwar, S.; and Zheng, L. 2020. From Depth What Can You See? Depth Completion via Auxiliary Image Reconstruction. In *CVPR*, 11303–11312. IEEE.
- Maeda, T.; Satat, G.; Swedish, T.; Sinha, L.; and Raskar, R. 2019. Recent advances in imaging around corners. *arXiv preprint arXiv:1910.05613*.
- Meng, Z.; Jalali, S.; and Yuan, X. 2020. GAP-net for Snapshot Compressive Imaging. ArXiv:2012.08364 [eess].

- Metzler, C. A.; Heide, F.; Rangarajan, P.; Balaji, M. M.; Viswanath, A.; Veeraraghavan, A.; and Baraniuk, R. G. 2020. Deep-inverse correlography: towards real-time high-resolution non-line-of-sight imaging. *Optica*, 7(1): 63–71.
- Micheli, A. 2009. Neural network for graphs: A contextual constructive approach. *TNN*, 20(3): 498–511.
- Mou, C.; Wang, Q.; and Zhang, J. 2022. Deep Generalized Unfolding Networks for Image Restoration. In *CVPR*. IEEE.
- Mu, F.; Mo, S.; Peng, J.; Liu, X.; Nam, J. H.; Raghavan, S.; Velten, A.; and Li, Y. 2022. Physics to the Rescue: Deep Non-line-of-sight Reconstruction for High-speed Imaging. ArXiv:2205.01679 [cs, eess].
- Nam, J. H.; Brandt, E.; Bauer, S.; Liu, X.; Renna, M.; Tosi, A.; Sifakis, E.; and Velten, A. 2021. Low-latency time-of-flight non-line-of-sight imaging at 5 frames per second. *Nature communications*, 12(1): 1–10.
- Oono, K.; and Suzuki, T. 2020. Graph Neural Networks Exponentially Lose Expressive Power for Node Classification. In *ICLR*.
- O’Toole, M.; Lindell, D. B.; a; and Wetzstein, G. 2018. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555(7696): 338–341.
- Pan, Z.; Cai, J.; and Zhuang, B. 2022. Fast Vision Transformers with HiLo Attention. In *NIPS (NeurIPS)*.
- Pan, Z.; Zhuang, B.; He, H.; Liu, J.; and Cai, J. 2022. Less is More: Pay Less Attention in Vision Transformers. In *AAAI*.
- Pediredla, A.; Dave, A.; and Veeraraghavan, A. 2019. Snlos: Non-line-of-sight scanning through temporal focusing. In *2019 IEEE ICCP*, 1–13. IEEE.
- Pei, C.; Zhang, A.; Deng, Y.; Xu, F.; Wu, J.; David, U.; Li, L.; Qiao, H.; Fang, L.; and Dai, Q. 2021. Dynamic non-line-of-sight imaging system based on the optimization of point spread functions. *Optics Express*, 29(20): 32349–32364.
- Saunders, C.; Murray-Bruce, J.; and Goyal, V. K. 2019. Computational periscopy with an ordinary digital camera. *Nature*, 565(7740): 472–475.
- Scheiner, N.; Kraus, F.; Wei, F.; Phan, B.; Mannan, F.; Appenrodt, N.; Ritter, W.; Dickmann, J.; Dietmayer, K.; Sick, B.; and Heide, F. 2020. Seeing around street corners: Non-line-of-sight detection and tracking in-the-wild using doppler radar. In *CVPR*, 2068–2077.
- Shen, S.; Wang, Z.; Liu, P.; Pan, Z.; Li, R.; Gao, T.; Li, S.; and Yu, J. 2021. Non-line-of-Sight Imaging via Neural Transient Fields. *TPAMI*, 43(7): 2257–2268.
- Su, X.; Yu, H.; Juntian, Y.; Feihu, X.; and Xin, Y. 2023. Multi-scale Iterative Model-guided Unfolding Network for NLOS Reconstruction. In *PG*. Wiley Online Library.
- Tan, M.; and Le, Q. 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *ICML*.
- Tolstikhin, I. O.; Houlsby, N.; Kolesnikov, A.; Beyer, L.; Zhai, X.; Unterthiner, T.; Yung, J.; Keysers, D.; Uszkoreit, J.; Lucic, M.; and Dosovitskiy, A. 2021. MLP-Mixer: An all-MLP Architecture for Vision. *NIPS*, 34: 12116–12128.
- Tsai, C.-Y.; Sankaranarayanan, A. C.; and Gkioulekas, I. 2019. Beyond Volumetric Albedo—A Surface Optimization Framework for Non-Line-Of-Sight Imaging. In *CVPR*.
- Tu, Z.; Talebi, H.; Zhang, H.; Yang, F.; Milanfar, P.; Bovik, A.; and Li, Y. 2022. MAXIM: Multi-Axis MLP for Image Processing. In *CVPR*, 5769–5780.
- Velten, A.; Willwacher, T.; Gupta, O.; Veeraraghavan, A.; Bawendi, M. G.; and Raskar, R. 2012. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications*, 3(1): 1–8.
- Wang, J.; Chen, Y.; Chakraborty, R.; and Yu, S. X. 2020. Orthogonal Convolutional Neural Networks. In *CVPR*.
- Wang, Y.; and Solomon, J. M. 2021. Object dgcnn: 3d object detection using dynamic graphs. *NIPS*, 34: 20745–20758.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic graph cnn for learning on point clouds. *ToG*, 38(5): 1–12.
- Wt, Z.; Zhangt, J.; and Mou, C. 2021. Dense Deep Unfolding Network with 3D-CNN Prior for Snapshot Compressive Imaging.
- Wu, C.; Liu, J.; Huang, X.; Li, Z.-P.; Yu, C.; Ye, J.-T.; Zhang, J.; Zhang, Q.; Dou, X.; Goyal, V. K.; Xu, F.; and Pan, J.-W. 2021. Non-line-of-sight imaging over 1.43 km. *PNAS*.
- Xin, S.; Nousias, S.; Kutulakos, K. N.; Sankaranarayanan, A. C.; Narasimhan, S. G.; and Gkioulekas, I. 2019. A theory of Fermat paths for non-line-of-sight shape reconstruction. In *CVPR*, 6800–6809.
- Xu, D.; Zhu, Y.; Choy, C. B.; and Fei-Fei, L. 2017. Scene graph generation by iterative message passing. In *CVPR*.
- Xu, F.; Shulkind, G.; Thrampoulidis, C.; Shapiro, J. H.; Torralba, A.; Wong, F. N.; and Wornell, G. W. 2018. Revealing hidden scenes by photon-efficient occlusion-based opportunistic active imaging. *Optics express*, 26(8): 9945–9962.
- Yang, Y.; Qiu, J.; Song, M.; Tao, D.; and Wang, X. 2020. Distilling knowledge from graph convolutional networks. In *CVPR*, 7074–7083.
- Ye, J.; Hong, Y.; Su, X.; Yuan, X.; and Xu, F. 2024. Plug-and-Play Algorithms for Dynamic Non-line-of-sight Imaging. *ToG*.
- Ye, J.-T.; Huang, X.; Li, Z.-P.; and Xu, F. 2021. Compressed sensing for active non-line-of-sight imaging. *Optics Express*, 29(2): 1749–1763.
- Young, S. I.; Lindell, D. B.; Girod, B.; Taubman, D.; and Wetzstein, G. 2020. Non-line-of-sight surface reconstruction using the directional light-cone transform. In *CVPR*.
- Yu, Y.; Shen, S.; Wang, Z.; Huang, B.; Wang, Y.; Peng, X.; Xia, S.; Liu, P.; Li, R.; and Li, S. 2023. Enhancing Non-line-of-sight Imaging via Learnable Inverse Kernel and Attention Mechanisms. In *ICCV 2023*, 10563–10573.
- Zhang, J.; and Ghanem, B. 2018. ISTA-Net: Interpretable Optimization-Inspired Deep Network for Image Compressive Sensing. In *2018 CVPR*, 1828–1837. Salt Lake City, UT: IEEE. ISBN 978-1-5386-6420-9.
- Zhou, Y.; Lu, R.; Xue, F.; and Gao, Y. 2023. Occlusion Relationship Reasoning with a Feature Separation and Interaction Network. In *Visual Intelligence 1, Article no. 23*.
- Zhu, S.; Sua, Y. M.; Bu, T.; and Huang, Y.-P. 2023. Compressive non-line-of-sight imaging with deep learning. *Physical Review Applied*, 19(3): 034090.