

Robust Single-shot Structured Light 3D Imaging via Neural Feature Decoding

JIAHENG LI*, Wangxuan Institute of Computer Technology, Peking University, China

QIYU DAI*, School of Intelligence Science and Technology, Peking University, China

LIHAN LI, Yuanpei College, Peking University, China

PRANEETH CHAKRAVARTHULA, University of North Carolina at Chapel Hill, The United States

HE SUN, College of Future Technology, Peking University, China

BAOQUAN CHEN[†], School of Intelligence Science and Technology, Peking University, China and State Key Laboratory of General Artificial Intelligence, Peking University, China

WENZHENG CHEN[†], Wangxuan Institute of Computer Technology, Peking University, China and State Key Laboratory of General Artificial Intelligence, Peking University, China

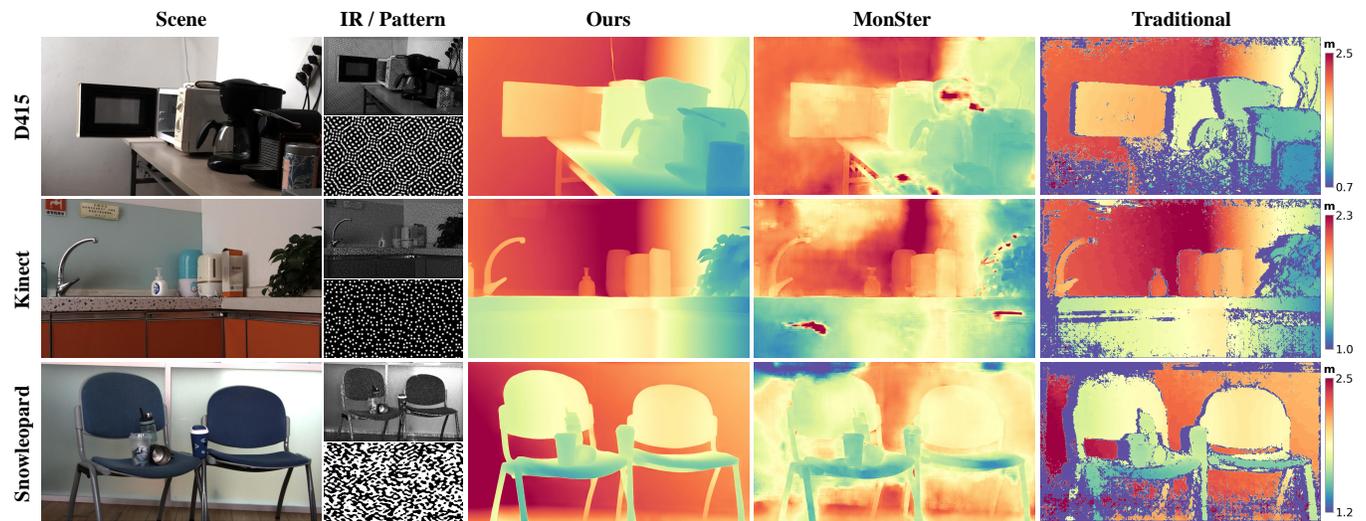


Fig. 1. Trained entirely on synthetic data, our structured light 3D imaging method generalizes well to real-world indoor scenes. It supports various pattern types (rows) and effectively handles challenging cases, including low-texture regions, fine structural details, and reflective or transparent surfaces. Our method significantly outperforms both RGB-based stereo methods (e.g., MonSter) and traditional pixel-matching-based structured light decoding approaches.

*Equal contribution

[†]Co-corresponding authors

Authors' Contact Information: Jiaheng Li, jiahengli@buaa.edu.cn, Wangxuan Institute of Computer Technology, Peking University, Beijing, China; Qiyu Dai, qiuyudai@pku.edu.cn, School of Intelligence Science and Technology, Peking University, Beijing, China; Lihan Li, lh_li@stu.pku.edu.cn, Yuanpei College, Peking University, Beijing, China; Praneeth Chakravarthula, cpk@cs.unc.edu, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, The United States; He Sun, hesun@pku.edu.cn, College of Future Technology, Peking University, Beijing, China; Baoquan Chen, baoquan@pku.edu.cn, School of Intelligence Science and Technology, Peking University, Beijing, China and State Key Laboratory of General Artificial Intelligence, Peking University, Beijing, China; Wenzheng Chen, wenzhengchen@pku.edu.cn, Wangxuan Institute of Computer Technology, Peking University, Beijing, China and State Key Laboratory of General Artificial Intelligence, Peking University, Beijing, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

We consider the problem of active 3D imaging using single-shot structured light systems, which are widely employed in commercial 3D sensing devices such as Apple Face ID and Intel RealSense. Traditional structured light methods typically decode depth correspondences through pixel-domain matching algorithms, resulting in limited robustness under challenging scenarios like occlusions, fine-structured details, and non-Lambertian surfaces. Inspired by recent advances in neural feature matching, we propose a learning-based structured light decoding framework that performs robust correspondence matching within feature space rather than the fragile pixel domain. Our method extracts neural features from the projected patterns and captured infrared (IR) images, explicitly incorporating their geometric priors by building cost volumes in feature space, achieving substantial performance improvements over pixel-domain decoding approaches. To further enhance depth quality, we introduce a depth refinement module that leverages strong

SA Conference Papers '25, Hong Kong, Hong Kong

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-2137-3/2025/12

<https://doi.org/10.1145/3757377.3763924>

priors from large-scale monocular depth estimation models, improving fine detail recovery and global structural coherence. To facilitate effective learning, we develop a physically-based structured light rendering pipeline, generating nearly one million synthetic pattern-image pairs with diverse objects and materials for indoor settings. Experiments demonstrate that our method, trained exclusively on synthetic data with multiple structured light patterns, generalizes well to real-world indoor environments, effectively processes various pattern types without retraining, and consistently outperforms both commercial structured light systems and passive stereo RGB-based depth estimation methods. Code and data are available at <https://github.com/Namintismpot/NSL>

CCS Concepts: • **Computing methodologies** → **3D imaging**.

Additional Key Words and Phrases: Structured Light, Depth Estimation, Stereo Vision

ACM Reference Format:

Jiaheng Li, Qiyu Dai, Lihan Li, Praneeth Chakravarthula, He Sun, Baoquan Chen, and Wenzheng Chen. 2025. Robust Single-shot Structured Light 3D Imaging via Neural Feature Decoding. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3757377.3763924>

1 Introduction

Active structured light (SL) 3D imaging has made significant progress over the past two decades, enabling accurate and efficient 3D reconstruction for numerous applications, including augmented reality, robotics, and industrial automation. Single-shot structured light systems, which are widely adopted in commercial devices such as Apple Face ID [Apple 2025], Microsoft Kinect [Microsoft 2025], and Intel RealSense [Corporation 2025], operate by projecting spatially-coded infrared patterns onto the scene and decoding depth information by finding correspondence between the projected pattern and the captured image. These systems have attracted growing attention due to their simplicity, high efficiency, and suitability for real-time 3D sensing.

Traditional single-shot structured light decoding methods predominantly rely on pixel-domain correspondence matching, i.e., decoding depth from local patch intensity cues. Commercial systems such as Intel RealSense [Corporation 2025] and Kinect V1 [Tölgessy et al. 2021] project dot patterns and use patch-based template matching. However, these approaches are fundamentally limited: image patches contain only low-level, local intensity information, which is often insufficient for robust correspondence, especially under challenging conditions such as occlusions, low-texture surfaces, or non-Lambertian reflectance. As a result, pixel-level decoding becomes highly unstable, compromising 3D acquisition quality in complex scenes.

On the other side, although deep learning has significantly advanced passive RGB-based stereo tasks, its application to structured light remains limited and ineffective, even though the core of both tasks is stereo matching. This is mainly due to constraints in both data and method design. On the data side, SL lacks large-scale datasets comparable to those for RGB-based stereo settings. Existing approaches often train on small synthetic datasets or use semi-supervised learning on limited real SL data without ground-truth depth, resulting in poor generalization [Baek and Heide 2021; Riegler et al. 2019;

Xu et al. 2022; Zhang et al. 2018]. On the methodological side, most prior works ignore the known projected patterns, treating SL as merely RGB stereo with added texture, rather than leveraging the spatial priors encoded in the projections.

In this work, we present **NSL**: a novel **Neural Structured Light** decoding framework that simultaneously addresses above challenges. At the core of NSL is to shift stereo matching from the fragile pixel domain to a robust neural feature space. Inspired by advances in learning-based passive stereo [Chang and Chen 2018; Chen et al. 2024; Guo et al. 2019; Jing et al. 2023; Li et al. 2024; Lipson et al. 2021; Xu and Zhang 2020; Zhang et al. 2019], we argue that structured light can similarly benefit from learned features. Our key idea is to match projected patterns and IR images via deep features rather than raw intensities, which moves beyond traditional structured light decoding methods that operate purely in the pixel intensity domain, and substantially improves robustness and accuracy. In contrast to passive stereo methods, which rely solely on images of natural scenes without projected patterns, NSL extracts matching features mainly from the actively projected, spatially encoded patterns.

Following setups in commercial SL devices, our method supports monocular (pattern + single IR) and binocular (pattern + stereo IR) inputs. It consists of two neural stages: a learned feature matching module and a monocular depth refinement module. The first stage, inspired by RAFT [Lipson et al. 2021], uses a stereo network to extract features from both pattern and IR inputs, build hierarchical cost volumes, and iteratively estimate an initial dense depth map. Importantly, leveraging the pattern features allows the network to exploit spatial priors encoded in the projection in an end-to-end manner, outperforming conventional stereo in textureless regions even using only one camera and one projector (monocular structured light).

To enhance depth quality, we introduce a monocular refinement module as a second stage, aimed at recovering structural detail and correcting mismatches from triangulation. We adopt a fine-tuned monocular depth estimation (MDE) backbone [Lin et al. 2024], using the initial depth as a prompt. By combining visual context from the IR image with geometric cues from the coarse depth, the module outputs sharper, more coherent depth maps, particularly around boundaries and fine structures.

To support the training of NSL, we develop a physically-based structured light simulation platform in Blender and generate a large-scale, high-fidelity synthetic indoor dataset containing 953K pattern-image pairs. The dataset covers diverse indoor layouts, materials, textures, lighting conditions, projected patterns and hardware settings, and includes synchronized RGB, IR, and pattern with ground-truth depth labels. In training, we use a mixture of projected patterns, enabling NSL to be trained once and then process multiple patterns at inference without retraining. Importantly, although our method is trained entirely on synthetic data, it generalizes well to real-world scenes without requiring any fine-tuning. This strong generalization capability comes from the fact that the network learns correspondence matching from the projected pattern cues, which exhibit minimal domain gaps between simulation and reality.

Extensive experiments show that our method consistently outperforms traditional structured light systems, existing neural structured light approaches, and passive stereo-based learning methods, particularly in challenging regions such as occlusions, fine structural details, reflective surfaces, and low-texture areas.

2 Related Work

In this section, we review three related areas. First, we discuss the evolution of depth imaging techniques, highlighting the contrast between passive and active methods. Next, we introduce traditional structured light systems, focusing on their core decoding principles. Lastly, we examine recent deep learning-based approaches for single-shot structured light decoding.

2.1 Depth Imaging

Depth imaging methods can be broadly categorized into *passive* and *active* approaches, depending on the use of external illumination. Passive methods infer depth from naturally available image cues such as shading [Tao et al. 2015], parallax [Godard et al. 2017], defocus [Hazirbas et al. 2019], polarization [Kadambi et al. 2015], and silhouettes [Laurentini 1994]. These cues are typically exploited in multiview stereo or structure-from-motion pipelines. Although advances in dense camera arrays [Levoy and Hanrahan 2023], neural scene representations [Mildenhall et al. 2021] and deep feature-based matching strategies [Zhan et al. 2018] have improved reconstruction fidelity, passive approaches still struggle with low-texture surfaces and lighting variations due to their reliance on natural image features and dense viewpoint coverage.

To overcome these challenges, active systems augment the scene with controlled illumination to enhance correspondence. Time-of-flight sensors [Sun et al. 2023], LiDAR [Huang et al. 2023], and structured light (SL) [Geng 2011] are commonly used. SL systems, in particular, offer high-resolution depth sensing at low cost, ideal for near-field real-time applications. Nevertheless, conventional SL techniques often depend on pixel-wise intensity matching, which is sensitive to ambient illumination, specular reflections, and textureless regions.

2.2 Traditional Active Structured Light Depth Imaging

Structured Light (SL) has been a cornerstone of active depth sensing for several decades [Posdamer and Altschuler 1982; Will and Pennington 1971], with methods broadly classified into temporal and spatial encoding [Salvi et al. 2010]. Temporal SL projects time-varying patterns (e.g., binary [Scharstein and Szeliski 2003], gray [Aliaga and Xu 2008; Posdamer and Altschuler 1982], or fringe patterns [Kawasaki et al. 2008; Koninckx and Van Gool 2006; Sagawa et al. 2011; Taguchi et al. 2012]) and decodes per-pixel signals for fast depth recovery, but suffers from motion artifacts due to multi-frame requirements.

Spatial SL, on the other hand, uses a single static 1D or 2D pattern [Le Moigne and Waxman 1988; Maruyama and Abe 1993; Vuylsteke and Oosterlinck 1990; Zhang et al. 2002], enabling one-shot capture and better motion robustness. This shifts the challenge to correspondence estimation via image matching. For example, Kinect

V1 [Martinez and Stiefelhagen 2013] employs local block matching against a reference, but such correlation-based approaches can struggle in non-ideal regions due to assumptions like local photo-consistency.

2.3 Deep Learning for Single-Shot Structured Light Decoding

Despite the success of deep learning in passive stereo, its application to single-shot structured light (SL) decoding remains under-explored. Early methods lacked end-to-end learning frameworks and achieved limited performance [Fanello et al. 2016, 2017]. ActiveStereoNet [Zhang et al. 2018] was among the first to apply CNN-based training for SL stereo matching, but it excluded the projected pattern from inference and relied on semi-supervised learning due to limited data that lacks ground-truth depth, resulting in unstable training and poor generalization. Connecting [Riegler et al. 2019], though categorized under neural SL, is not a decoding method. It formulates the problem as monocular depth estimation without using the projected pattern during inference, leading to poor cross-device generalization.

Polka [Baek and Heide 2021] introduced a parametric diffractive optical element (DOE) model for jointly optimizing pattern design and depth decoding, while MonoStereoFusion [Xu et al. 2022] proposed a two-stage pipeline where the coarse depth is generated through traditional decoding between the projected pattern and the left IR image, and subsequently used to guide stereo matching between binocular IR images. Both methods highlight the benefit of incorporating pattern information, but Polka is restricted to DOE-generated dot patterns, and MonoStereoFusion excludes the pattern from end-to-end learning, underutilizing SL priors. Furthermore, both approaches synthesize training data by overlaying patterns on passive stereo datasets—an approach that lacks physical realism and fails to capture challenges such as reflectance and ambient light, thus limiting generalization.

In contrast, our method extracts neural features from both the projected pattern and IR image in an end-to-end framework, using a physically based simulator to generate large-scale synthetic data with diverse materials and pattern types. This results in improved accuracy, robustness, and sim-to-real generalization across occlusions, reflective surfaces, and low-texture regions.

3 Methods

We now describe the pipeline of NSL, a single-shot structured light decoding framework that replaces traditional pixel-domain template matching with robust neural feature decoding. Given either a single or stereo pair of IR images captured by a structured light system and the corresponding projected pattern, NSL produces a dense, high-quality depth map of the scene. As illustrated in Figure 2, NSL consists of two main components: (1) a Neural Feature Matching module that extracts deep features from the IR image(s) and pattern to estimate an initial depth map, and (2) a Monocular Depth Refinement module that leverages priors from a monocular depth estimation model to improve depth quality. Below, we describe each module in detail.

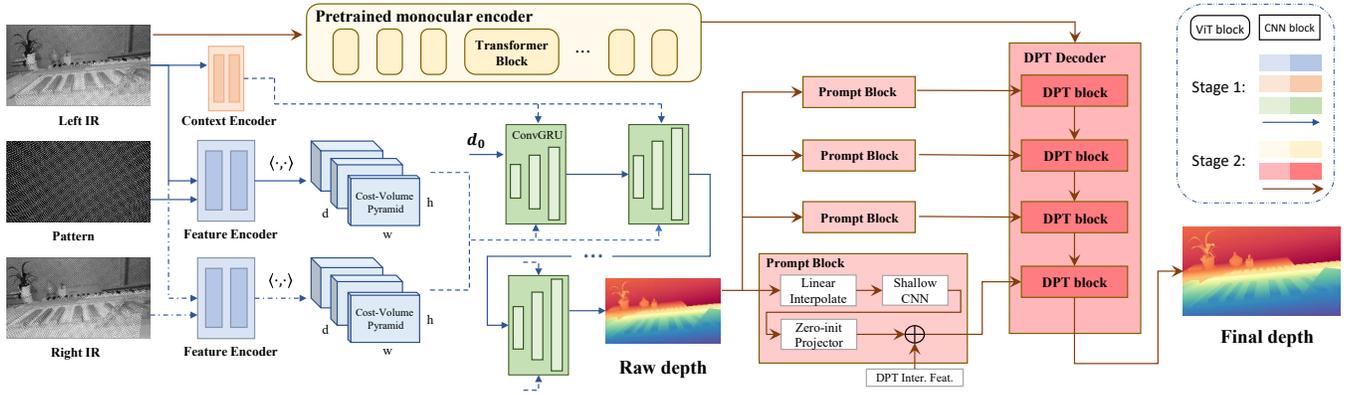


Fig. 2. **The pipeline of NSL.** Given a single or stereo pair of IR images and a projected pattern, NSL first estimates an initial raw depth map via the Neural Feature Matching module, which extracts deep features from both IR and pattern inputs, followed by cost volume construction and GRU-based iterative refinement (left path). Next, the Monocular Depth Refinement module incorporates priors from a monocular depth estimation model, using the initial depth as a prompt (right path), and generates a final depth map with enhanced structural detail.

3.1 Neural Feature Matching

To replace fragile pixel-domain decoding in structured light (SL), we propose a neural feature matching module that estimates dense correspondences in learned feature space. Inspired by RAFT-Stereo [Lipson et al. 2021], our model extracts deep features from the projected pattern and IR image(s), builds multi-level cost volumes, and iteratively refines depth predictions using a GRU-based updater.

Feature Extraction. Let $I_l \in \mathbb{R}^{H \times W}$ denote the left IR image (used in both monocular and binocular SL modes), $I_r \in \mathbb{R}^{H \times W}$ denote the right IR image (used only in binocular SL mode), and $I_p \in \mathbb{R}^{H \times W}$ denote the projected pattern. A CNN-based *Feature Encoder* Enc^{lp} extracts 1/4-resolution feature maps from the Left IR image and the projected pattern. For stereo IR input, another encoder Enc^{lr} with the same architecture is also used to process the Left and Right IR images.

$$\begin{aligned} F_l^{lp} &= \text{Enc}^{lp}(I_l), & F_p &= \text{Enc}^{lp}(I_p), \\ F_l^{lr} &= \text{Enc}^{lr}(I_l), & F_r &= \text{Enc}^{lr}(I_r). \end{aligned} \quad (1)$$

To utilize global image priors in depth reasoning, we further use a separate *Context Encoder* to extract multi-scale context features from I_l :

$$F_l^c = \text{ContextEnc}(I_l), \quad (2)$$

where F_l^c provides multi-resolution cues used in the GRU-based refinement stage.

Cost Volume Construction. In traditional RGB-based stereo matching, a 3D cost volume $C \in \mathbb{R}^{H \times W \times W}$ is constructed by computing the dot product between feature maps F_L and F_R :

$$C(i, j, k) = \sum_h F_{L_{i,j,h}} \cdot F_{R_{i,k,h}}, \quad (3)$$

where (i, j) denotes the pixel location in the left image and k enumerates correspondence candidates in the right image. This gives a matching score between pixel (i, j) in the left image and pixel (i, k) in the right image.

In our NSL framework, F_L and F_R correspond to the left IR image feature F_l^{lp} and the pattern feature F_p . This formulation enables pattern-to-image matching, allowing the network to leverage the spatial information encoded in the pattern. For stereo IR input, a second cost volume is additionally constructed from F_l^{lr} and F_r .

To capture multi-scale cues, we build a 4-level cost volume pyramid by applying average pooling along the last dimension of C :

$$C^{(l)} = \text{AvgPool}^{(l)}(C), \quad \text{for } l = 0, 1, 2, 3. \quad (4)$$

In the binocular SL setup, features from the two cost volume pyramids, one for left-pattern matching and one for stereo ir matching, are concatenated together and used jointly in the refinement stage.

GRU-based Depth Prediction. Following RAFT-Stereo, we apply a GRU-based iterative update to predict disparity. Starting from an initial disparity estimate d_0 , the module generates a sequence of disparity maps $\{d_1, \dots, d_N\}$ over N iterations. The final predicted disparity map d_N is then converted into the initial raw depth map D_{init} .

At each iteration t , we sample correlation features from the multi-scale cost volume at the current disparity d_t , denoted as $\phi(C, d_t)$. These features, along with the context features F_l^c , are fed into a convolutional GRU:

$$h_{t+1} = \text{GRU}(h_t, \phi(C, d_t), F_l^c), \quad (5)$$

$$\Delta d_t = \text{Regress}(h_{t+1}), \quad d_{t+1} = d_t + \Delta d_t, \quad (6)$$

where h is the hidden feature in each iteration. After N iterations, the final disparity prediction d_N is converted to an initial depth D_{init} , which is used as the initial depth estimate for the refinement module.

3.2 Monocular Depth Refinement

While neural feature matching significantly improves structured light decoding, the resulting depth can remain suboptimal, particularly in challenging regions such as thin structures, occlusions, or

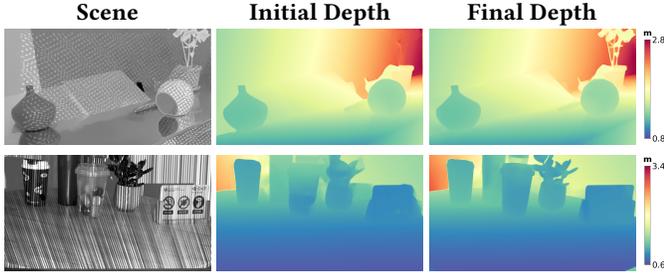


Fig. 3. **Effectiveness of Monocular Depth Refinement.** The first row shows results on synthetic data, and the second row shows results on a real-captured scene. The refinement module uses monocular depth priors to recover challenging regions, such as thin structures and sharp edges.

indirect illumination. To address this, we introduce a Monocular Depth Refinement module that leverages global image priors from monocular depth estimation (MDE) foundation models [Lin et al. 2024]. These models, trained on millions of images, can infer plausible depth from a single RGB input, but typically suffer from scale ambiguity. We resolve this issue by using the structured light initial depth D_{init} as a geometric prior to guide the model and produce metrically accurate depth maps. Fig. 3 presents a visual comparison with and without this module.

Depth Foundation Model. We adopt Depth Anything v2 [Yang et al. 2024] as our base model. It first extracts multi-scale image features using a Vision Transformer (ViT) [Oquab et al. 2024], then aggregates them using a DPT-style decoder [Ranftl et al. 2021]. The final depth is predicted via a convolutional head.

Previous works [Lin et al. 2024] demonstrate that foundation models like Depth Anything v2 supports the injection of auxiliary geometric signals, such as LiDAR, by incorporating them into the DPT decoder. We then modify its input and prompt settings to accommodate to our structured light setting.

Structured Light Prompt Injection. In our case, we replace the RGB image I with the IR input I_l . Additionally, we use the initial depth D_{init} as a depth prompt. The prompt network adopts the architecture proposed in PromptDA (see Fig. 2). For more efficient training and inference, our approach employs a backbone model at the ViT-Base scale.

3.3 Structured Light Simulation and Data Generation

A key bottleneck preventing the application of neural methods to SL decoding is the lack of large-scale labeled data. Robust neural decoders typically require training on millions of diverse image pairs to generalize well across real-world variations. However, collecting such real-world SL data at scale is time-consuming and costly.

To address this, we train our model entirely on synthetic data. We develop a physically-based SL simulation platform using Blender’s Cycles ray-tracing engine for rendering a large amount of synthetic SL data for training.

Dataset Composition. We construct a large-scale indoor dataset tailored for structured light decoding. The dataset includes:

- **Scenes:** 2860 indoor scenes generated from Proctor[Deitke et al. 2022], illuminated with realistic HDR environment maps (100 maps).
- **Objects:** 5000+ 3D objects from sources like ShapeNet or Objaverse, covering a wide range of categories with diffuse, specular, and transparent materials.
- **Patterns:** eight structured light patterns (e.g., random binary textures, dot patterns), of which six are used for training and two for testing.

During rendering, we randomize object placement, materials, scene layout, camera poses, projected patterns, camera parameters, baseline distance, and lighting conditions. This diversity helps simulate the wide range of challenging scenarios encountered in real-world structured light capture, such as occlusions, low-texture surfaces, non-Lambertian materials, and complex illumination. We generate a total of 953K samples. Each sample includes binocular RGB images, binocular IR images, a pattern image, and a ground-truth depth map. We show the dataset in the Supp.

Sim2Real Generalization. Although our model is trained entirely on synthetic data, it generalizes effectively to real-world structured light systems. We attribute this to the physically accurate simulation process and the fact that our decoder learns correspondences directly from pattern-guided feature matching, which minimizes the domain gap between synthetic and real data. And the monocular depth refinement module also maintains this generalization, leveraging its robust foundation model backbone and focusing on localized corrections guided by the initial depth, as shown in Fig. 1, 3 and 8.

3.4 Network Training

We train the Neural Feature Matching and Monocular Depth Refinement modules in two stages: 200k iterations for the former, followed by 100k iterations fine-tuning the latter using D_{init} from stage one.

Neural Feature Matching. The feature matching module predicts a sequence of disparity maps $\{d_1, \dots, d_N\}$ through iterative GRU refinement. We supervise each stage with a weighted L1 loss against the ground truth disparity d_{gt} :

$$\mathcal{L}_{\text{stage 1}} = \sum_{t=1}^N \lambda_t \cdot |d_t - d_{gt}|_1, \quad (7)$$

where λ_t denotes the weight for iteration t , with larger weights assigned to later stages for better supervision on refined outputs.

Monocular Depth Refinement. The refinement module is initialized from the pre-trained Depth Anything v2 [Yang et al. 2024] and fine-tuned on our dataset. It takes the IR image I_l and the initial depth D_{init} as input and is supervised by two terms:

$$\mathcal{L}_{\text{stage 2}} = |D - D_{gt}|_1 + \alpha \cdot |\nabla D - \nabla D_{gt}|_1, \quad (8)$$

where the first term ensures metric accuracy, and the second enforces edge consistency. α is a fixed weight (set to 0.5), and ∇ denotes the spatial gradient.

4 Experimental Results

We conduct comprehensive experiments on both synthetic and real-world datasets to evaluate the effectiveness and generalizability of the proposed NSL framework. Section 4.1 details the evaluation metrics and baseline methods. Section 4.2 presents results on the synthetic dataset, followed by ablation studies in Section 4.3, which analyze the impact of various design choices. Finally, Section 4.4 demonstrates the framework’s practical applicability using data captured from both a custom structured light setup and the commercially available RealSense D435 RGB-D camera.

4.1 Comparison Details

Configuration. As our method is very flexible, it supports various input configuration including monocular (single IR + Pattern) and binocular (dual IR + Pattern) SL settings. We train two models under both settings, noted as **NSL-Mono (monocular, single IR + Pattern)** and **NSL-Bino (binocular, dual IR + Pattern)**.

Baselines. We compare against both **intensity-based SL** and more recent **learning-based SL**. For the intensity-based baseline, we adopt a **traditional template-matching SL method (TM)** [Fixstars 2025], which is a pixel intensity based approach widely used in commercial products. For the learning-based baseline, we compare with **ActiveStereoNet** [Zhang et al. 2018], a CNN-based SL approach. We note that other learning-based SL works, MonoStereoFusion [Xu et al. 2022] (using a non-standard RGB+IR+pattern setup) and Polka [Baek and Heide 2021] (requiring pattern-decoder co-optimization), are not included since they provide no public-available code and are incompatible with our evaluation settings..

ActiveStereoNet is retrained and tested on our dataset using only image pairs rendered with binary patterns—where pixel values are strictly 0 or 1—since it crashes when pairs with non-binary patterns are used. Even under this restricted setup, training remains highly unstable, and the results reported are the best among multiple runs.

Additionally, to contextualize NSL in the broader depth estimation landscape and provide broader comparison, we include a state-of-the-art passive RGB stereo network MonSter [Cheng et al. 2025]. We fine-tune it using the passive RGB pairs in our dataset.

Evaluation Metrics. We use standard depth estimation metrics: RMSE(Root Mean Squared Error), MAE(Mean Absolute Error), and REL(Mean Relative Error) to measure overall, average, and relative errors. We also report δ -accuracy (thresholds $\delta \in \{1.05, 1.10, 1.25\}$), which measures the percentage of pixels where $\max(\frac{D_i}{D_i^*}, \frac{D_i^*}{D_i}) < \delta$. Though disparity based metrics like EPE (Endpoint error) are also commonly used in stereo systems, we choose depth based metrics because depth is the ultimate target and is insensitive to irrelevant conditions like image resolution.

4.2 Synthetic data evaluation

Experimental Setup. We validate our method on the proposed synthetic dataset. We use a resolution of 1280×720 for evaluation. During training, we only use a mixture of 6 patterns, while we evaluate the performance under all 8 patterns (including 2 unseen patterns). See Supp. for full details. Below, we compare baselines under both monocular setting and binocular setting.

Table 1. Quantitative Comparison of Our Method with Traditional Template Matching Methods (TM) and RAFT.

Methods	MAE(m) ↓	RMSE ↓	REL ↓	$\delta_{1.25}$ ↑	$\delta_{1.10}$ ↑	$\delta_{1.05}$ ↑
Monocular Setting						
TM(IR+Pattern) [†]	0.5250	1.1984	0.2886	0.7467	0.6955	0.5716
Ours(NSL-Mono)	0.0408	0.0907	0.0357	0.9704	0.9437	0.9142
Binocular Setting						
TM(Dual IR) [†]	0.3568	1.9008	0.2988	0.8001	0.7790	0.7532
ActiveStereoNet [‡]	0.5693	75.5485	0.4984	0.8455	0.7912	0.7287
MonSter(passive)	0.0862	0.1457	0.1839	0.9059	0.8611	0.8087
Ours(NSL-Bino)	0.0237	0.0634	0.0210	0.9817	0.9666	0.9502

[†] For the pixel-matching based traditional method (TM) [Fixstars 2025], large depth outliers are excluded before metric calculation; other methods do not undergo outlier removal.

[‡] ActiveStereoNet is evaluated only on the subset rendered with binary patterns. Others are evaluated on all patterns.

Performance Discussion. Quantitative results are summarized in Table 1. Our method significantly outperforms the classical single-shot pixel-based matching decoding approach, the representative SL neural decoding method ActiveStereoNet and the state-of-the-art passive RGB-based stereo method, across all metrics.

Specifically, even monocular SL configuration (NSL-Mono, IR + pattern) already surpasses MonSter and ActiveStereoNet, which takes dual RGB or dual IR as input. Extending to the binocular SL setting (dual IR + pattern) further enhances our method’s performance, demonstrating the advantage of utilizing SL pattern priors when combined with feature-space correspondence learning.

Qualitative results are shown in Figure 6. Compared to the baselines, our method produces sharper object boundaries, finer structural details, smoother surfaces, and more accurate recovery in textureless and specular regions. Traditional decoding methods rely on low-level pixel matching, which is easily disrupted and unstable under weak signal conditions. ActiveStereoNet fails to resolve fine structures, often producing oversmoothed and indistinct results due to ineffective feature matching. MonSter, lacking pattern-specific local cues, exhibits noticeable depth shifts and missing geometry. In contrast, NSL maintains strong structural consistency and captures fine geometric details, demonstrating the effectiveness of pattern-aware stereo decoding.

4.3 Ablation Studies

We conduct a series of ablation experiments to analyze the impact of different input configurations, the effectiveness of the monocular depth refinement module, and the generalization capability of our neural decoder across various structured light patterns.

Impact of Input Configurations. Our method is very flexible and can work under various settings. To evaluate the importance of different input settings and the role of the structured light pattern in decoding, besides NSL-Mono (monocular, single IR + pattern) and NSL-Bino (binocular, dual IR + pattern), we further train a new model that only takes dual IR as input (only with captured SL images, without SL pattern), noted as **NSL-Stereo (stereo, dual IR only)**. As shown in Table 2, regardless of whether the monocular

Table 2. Comparison of Different Input Configurations on Our Dataset.

Method	MAE(m) ↓	RMSE ↓	REL ↓	$\delta_{1.25}$ ↑	$\delta_{1.10}$ ↑	$\delta_{1.05}$ ↑
Stage 1 without depth refinement module						
NSL-Mono (IR+Pat.)	0.0599	0.3340	0.0637	0.9612	0.9387	0.9159
NSL-Stereo (Dual IR)	0.0347	0.1005	0.0433	0.9718	0.9539	0.9351
NSL-Bino (Dual IR+Pat.)	0.0253	0.1073	0.0299	0.9788	0.9675	0.9560
Stage 2 with depth refinement module						
NSL-Mono (IR+Pat.)	0.0408	0.0907	0.0357	0.9704	0.9437	0.9142
NSL-Stereo (Dual IR)	0.0303	0.0751	0.0251	0.9764	0.9526	0.9258
NSL-Bino (Dual IR+Pat.)	0.0237	0.0634	0.0210	0.9817	0.9666	0.9502

depth refinement module is included, the trend consistently holds: IR + Pattern < Dual IR < Dual IR + Pattern.

Notably, even in the stereo setup where complete binocular cues are available, incorporating the projected pattern still brings a clear performance gain. This demonstrates that patterns are not redundant—even with dual-view geometry—and should not be ignored.

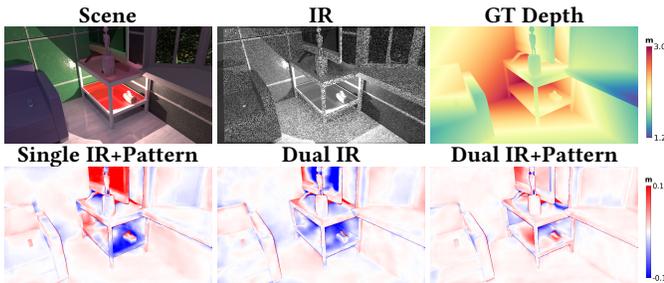


Fig. 4. **Comparison of Different Input of NSL.** The second row shows error maps. The depth error with pattern input (dual IR+Pattern) is lower than without (dual IR), validating the effectiveness of incorporating pattern features for structured light decoding.

Effectiveness of Monocular Depth Refinement. We compare results with and without the monocular refinement module across all three input settings (Rows 1–3 vs. 4–6 in Table 2; see also Figure 3). The refinement consistently improves performance, especially in the monocular IR + pattern setting, where global priors help compensate for missing stereo cues. A slight drop in $\delta_{1.1}$ and $\delta_{1.05}$ is observed, likely due to bias in the initial estimate D_{init} . While the module reduces MAE and RMSE by refining geometry and suppressing outliers (indicated by decreased MAE and RMSE), it may also propagate initial bias more broadly—highlighting a direction for future improvement.

Generalization Across Patterns. During training, we do not distinguish between structured light patterns and mix all 6 training patterns to improve robustness. To evaluate generalization, we test the trained model separately on each pattern in the test set, including two novel patterns (Kinect and RandomSquare) that were not seen during training. Results are summarized in Table 3. Full table can be found in the appendix. We observe only mild performance variation across patterns, indicating that our model has generalization capability to unseen patterns.

Table 3. NSL-Mono’s Performance Comparison of Different Patterns.

Method	MAE(m) ↓	RMSE ↓	REL ↓	$\delta_{1.25}$ ↑	$\delta_{1.10}$ ↑	$\delta_{1.05}$ ↑
Training Patterns						
D415	0.0388	0.0877	0.0354	0.9716	0.9458	0.9180
D435	0.0389	0.0878	0.0361	0.9721	0.9474	0.9195
Alacarte	0.0415	0.0913	0.0356	0.9661	0.9370	0.9068
Testing Patterns						
Randsquare	0.0395	0.0883	0.0338	0.9709	0.9450	0.9157

4.4 Real-World Evaluation

Hardware Setup. To validate the effectiveness of our method in real-world settings, we build a custom structured light capture system. The setup includes a Lenovo T8s projector and HikVision CS050-10UC left and right cameras. All components are calibrated using standard multi-view calibration techniques [Contributors 2025] to estimate intrinsic and extrinsic parameters.

In generating synthetic dual-view structured light data, we assume ideal optical axis alignment between the projector and both cameras, enabling perfect epipolar rectification for both dual IR and pattern inputs. However, such precise alignment is difficult to achieve with off-the-shelf projector-camera hardware. As a result, we evaluate only the Single IR + pattern and Dual IR settings on real-world captures.

In all real-data experiments, our method uses the model trained solely on the proposed synthetic dataset, without any additional training or fine-tuning. The evaluation resolution is 2040×1050 .

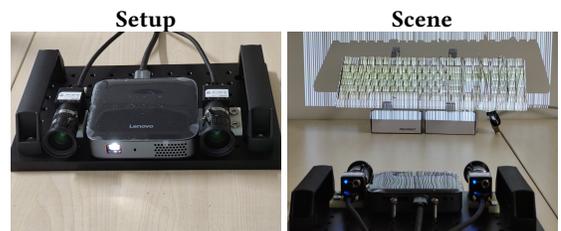


Fig. 5. **Hardware.** Our system consists of a Lenovo T8s projector and HikVision CS050-10UC left and right cameras.

Quantitative Real Results. We place the hardware on the desk and capture 11 real-world indoor scenes. We first apply traditional multi-shot structured light methods to estimate the GT depth. Pseudo ground truth depth maps were obtained by projecting 32 alacarte patterns and decoding using a ZNCC-based method [Mirdehghan et al. 2018]. We further create a GT mask to get rid of the occluded regions and outliers.

For evaluation, we captured all 8 structured light patterns from our dataset for each scene, resulting in a total of 88 structured light image pairs ($11 * 8 = 88$). All pattern types were included in the evaluation. Additionally, we captured passive stereo RGB pairs without pattern projection for each scene to enable comparison with RGB-based stereo methods. To ensure a fair comparison, we report both

Table 4. Quantitative Results on Real Data.

Methods	MAE(m)	$\delta_{1.05}$	EPE	Methods	MAE(m)	$\delta_{1.05}$	EPE
NSL-Mono	0.014	0.964	1.80	TM	0.111	0.901	12.55
NSL-Stereo	0.010	0.987	1.74	ASN	0.382	0.588	60.71
MonSter(z.s.)	0.849	0.008	228.8	MonSter(f.t.)	0.379	0.092	58.45

the zero-shot (z.s.) and fine-tuned (f.t.) results of MonSter. Fig. 7 and Tbl. 4 present the results.

Several interesting points can be observed. First, traditional decoding performs relatively well here compared to synthetic settings, potentially because pseudo ground truth is only available in unoccluded regions with non-reflective, non-transparent surfaces, and we evaluate only on these regions. Both ActiveStereoNet and MonSter fail under real-world conditions—the latter in particular suffers from severe scale inconsistencies (notably low $\delta_{1.05}$), which cannot be attributed to fine-tuning on synthetic data. In contrast, our method, even trained on synthetic data, demonstrates accurate, robust, and generalizable performance.

Qualitative Real Results. We further evaluate all methods under more challenging and diverse real-world scenes using hand-held capture conditions, and present qualitative results by visualizing depth maps across a broader range of scenarios. Representative examples are shown in Figure 1 and Figure 8. We also show the point clouds comparison in Figure 9. Since the traditional decoder’s point clouds are too noisy to be directly visualized, we apply the same point cloud denoising by removing points with anomalous depth gradient to all point clouds. Compared to the traditional decoding approach and state-of-the-art passive RGB stereo methods, our approach produces cleaner and more complete depth maps with fewer artifacts and sharper boundaries, particularly in difficult regions such as textureless surfaces and reflective or occluded areas.

Additionally, we evaluate our method on video sequences randomly recorded using a RealSense D435 depth camera, which demonstrates our method significantly outperforms commercial structured light cameras (Figure 10)

Limitation. Though our method generalizes well to high-quality inputs, its performance may degrade on lower-quality real data affected by factors such as motion-blur or hardware imperfections (e.g. lens blur). These are not considered during data synthesizing and result in a remaining sim2real gap.

Monocular depth refiner on low-quality images. While it performs well on real data from our high-end device, its effectiveness declines under low-quality imaging conditions like RealSense. Specifically, due to the absence of motion blur, noise, and defocus in the synthetic training data, the refiner tends to oversmooth depth predictions on RealSense D435 inputs. In contrast, the first-stage neural feature matching remains robust even under such degradations. As shown in Figure 11, applying deblurring and denoising as pre-processing may help mitigate this issue.

Object boundaries. Some artifacts can be observed on boundaries at depth discontinuities. Apart from the outliers caused by reprojection from IR view to RGB view in Realsense results, there may be oversmooth boundaries, especially under handheld condition

where motion-blur is severe. We will deal with non-ideal capture conditions in the future works.

4.5 Runtime Analysis

Table 5 reports the inference time on a 640×480 image pair with a single RTX4090 GPU. Though our model is not yet real-time (e.g., 30fps) due to its complexity and lack of speed optimization, it could be accelerated by standard techniques (e.g., quantization and distillation). ActiveStereoNet is faster but produces inaccurate depth with over-smooth details. Traditional method (TM) is ultra-fast but lacks robustness, producing noisy results. Passive RGB stereo baseline, MonSter, lags in both speed and quality.

Table 5. Runtime.

Methods	FPS
NSL-Mono/NSL-Stereo	18.94
NSL-Bino	15.38
TM	1098.89
ActiveStereoNet	80.39
MonSter	3.63

5 Conclusions

We present NSL, a neural decoding framework for single-shot structured light that performs correspondence matching in feature space rather than in the pixel domain. Our method significantly improves robustness and accuracy across both monocular and binocular structured light setups, outperforming traditional decoders and even RGB-based stereo methods. In future work, we aim to extend NSL to handle dynamic scenes, uncalibrated settings, and jointly learn pattern design and decoding in an end-to-end manner.

Acknowledgments

This research is an achievement of the Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Techonology). We gratefully acknowledge the support from Jiiov Technology for providing computing resources.

References

- Daniel G Aliaga and Yi Xu. 2008. Photogeometric structured light: A self-calibrating and multi-viewpoint framework for accurate 3d modeling. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1–8.
- Apple. 2025. Apple Face ID. <https://support.apple.com/en-us/102381> Accessed: 2025-04-08.
- Seung-Hwan Baek and Felix Heide. 2021. Polka lines: Learning structured illumination and reconstruction for active stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5757–5767.
- Jia-Ren Chang and Yong-Sheng Chen. 2018. Pyramid stereo matching network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5410–5418.
- Ziyang Chen, Wei Long, He Yao, Yongjun Zhang, Bingshu Wang, Yongbin Qin, and Jia Wu. 2024. Mocha-stereo: Motif channel attention network for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 27768–27777.
- Junda Cheng, Longliang Liu, Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Yong Deng, Jinliang Zang, Yurui Chen, Zhipeng Cai, and Xin Yang. 2025. MonSter: Marry Monodepth to Stereo Unleashes Power. *arXiv preprint arXiv:2501.08643* (2025).
- OpenCV Contributors. 2025. OpenCV. <https://opencv.org/> Accessed: 2025-04-10.
- Intel Corporation. 2025. Intel RealSense – Computer Vision and Depth Tracking Cameras. <https://www.intelrealsense.com/> Accessed: 2025-04-08.

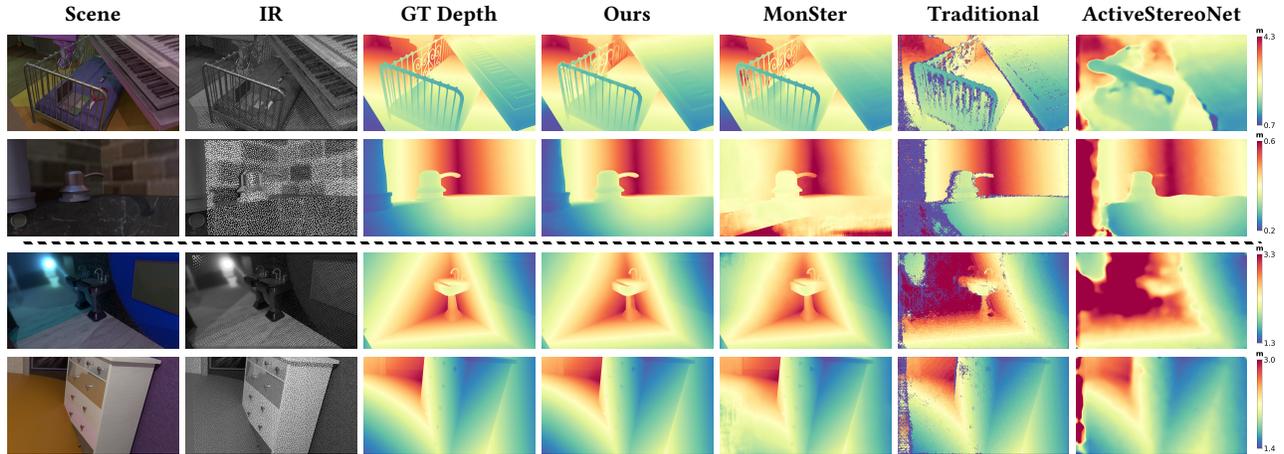


Fig. 6. **Qualitative results on our synthetic dataset.** Our method (top two rows: NSL-Mono(single IR+pattern); bottom two: NSL-Bino(dual IR+pattern)) is compared with MonSter (dual RGB), ActiveStereoNet (dual IR) and traditional methods (dual IR). It produces more continuous, high-fidelity depth with finer details.

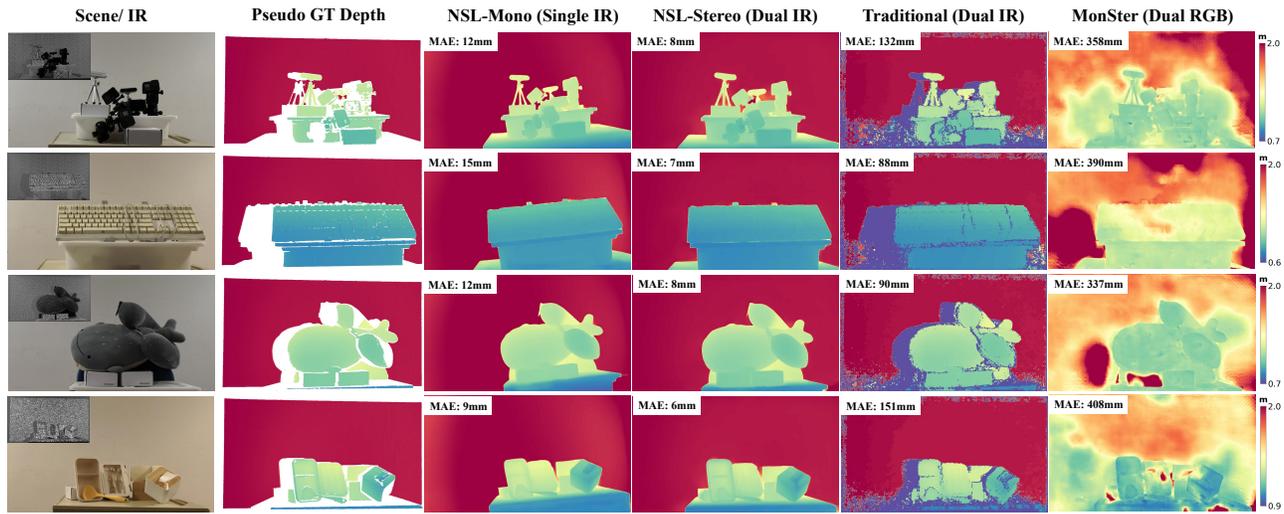


Fig. 7. **Quantitative results on real scenes.** Our method achieves 14mm (NSL-Mono) and 10mm (NSL-Stereo) accuracy in known-depth regions. It also delivers better details, smoother surfaces, and more complete depth in depth-unknown regions. The viewpoint difference between single IR and dual IR results stems from distinct epipolar rectifications: one for the camera-projector pair (monocular SL) and the other for the stereo IR camera setup (binocular SL).

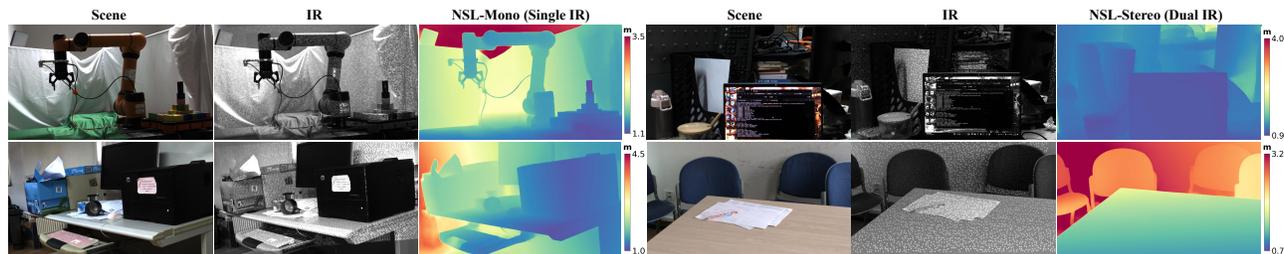


Fig. 8. **More qualitative results on real scenes captured with the device held by hand.** No pseudo ground truth available, unlike the fixed setup.



Fig. 9. **Point cloud quality comparison of random scene reconstruction under handheld capture.** The same denoising procedure is applied to all point clouds for fair visualization. (TM's point clouds are too noisy to be directly visualized). RGB images are warped onto the IR viewpoint.

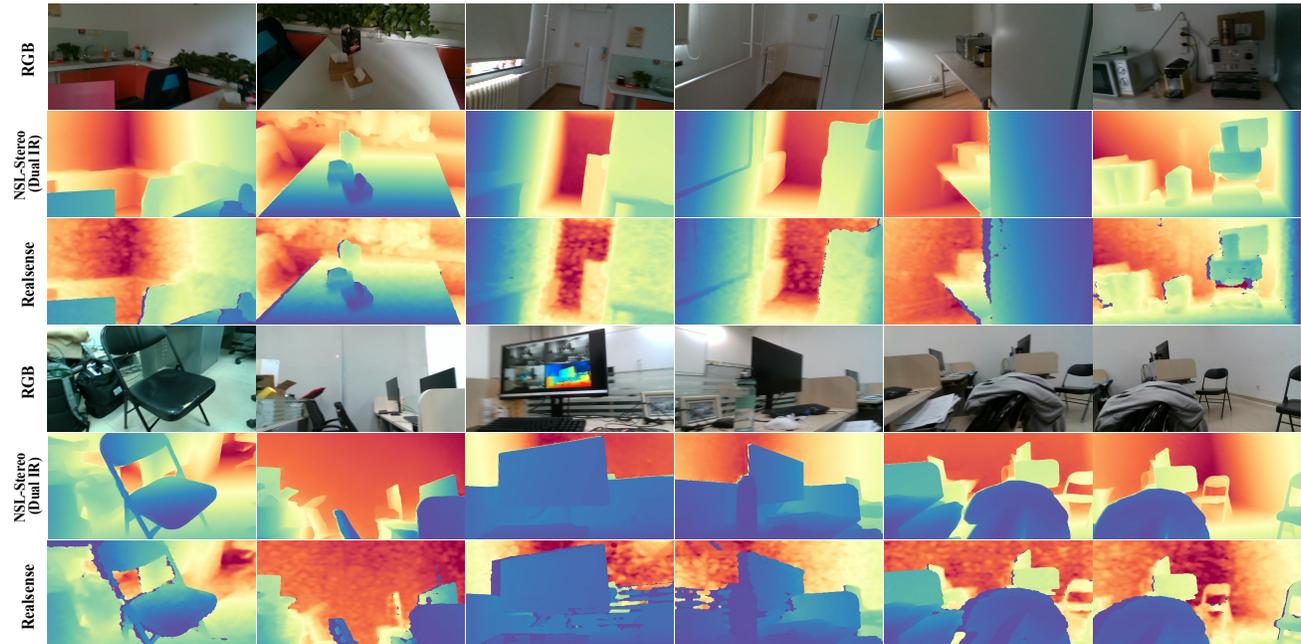


Fig. 10. **Qualitative results of NSL on RealSense dual IR data.** Trained only on synthetic data, NSL generalizes well to real IR and outperforms RealSense. Due to lower input quality, These depth maps are slightly inferior to those from our self-built device.

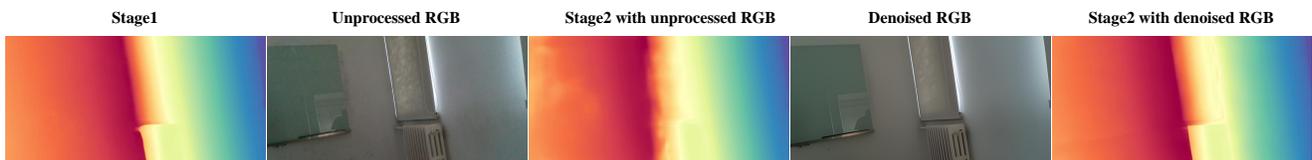


Fig. 11. The second-stage refinement may produces oversmoothed depth with noisy images from RealSense D435. Applying denoising and deblurring mitigates the issue. For clearer visualization, RGB images are used as inputs to the monocular depth refiner here.

- Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Jordi Salvador, Kiana Ehsani, Winson Han, Eric Kolve, Ali Farhadi, Aniruddha Kembhavi, and Roozbeh Mottaghi. 2022. ProcTHOR: Large-Scale Embodied AI Using Procedural Generation. In *NeurIPS*. Outstanding Paper Award.
- Sean Ryan Fanello, Christoph Rhemann, Vladimir Tankovich, Adarsh Kowdle, Sergio Orts Escolano, David Kim, and Shahram Izadi. 2016. Hyperdepth: Learning depth from structured light without matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5441–5450.
- Sean Ryan Fanello, Julien Valentin, Christoph Rhemann, Adarsh Kowdle, Vladimir Tankovich, Philip Davidson, and Shahram Izadi. 2017. Ultrastereo: Efficient learning-based matching for active stereo systems. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 6535–6544.
- Fixstars. 2025. libSGM. <https://github.com/fixstars/libSGM> Accessed: 2025-04-10.
- Jason Geng. 2011. Structured-light 3D surface imaging: a tutorial. *Advances in optics and photonics* 3, 2 (2011), 128–160.
- Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. 2017. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 270–279.
- Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. 2019. Group-wise correlation stereo network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 3273–3282.
- Caner Hazirbas, Sebastian Georg Soyer, Maximilian Christian Staab, Laura Leal-Taixé, and Daniel Cremers. 2019. Deep depth from focus. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*. Springer, 525–541.
- Shengyu Huang, Zan Gojic, Zian Wang, Francis Williams, Yoni Kasten, Sanja Fidler, Konrad Schindler, and Or Litany. 2023. Neural lidar fields for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 18236–18246.
- Junpeng Jing, Jiankun Li, Pengfei Xiong, Jiangyu Liu, Shuaicheng Liu, Yichen Guo, Xin Deng, Mai Xu, Lai Jiang, and Leonid Sigal. 2023. Uncertainty guided adaptive warping for robust and efficient stereo matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3318–3327.
- Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. 2015. Polarized 3d: High-quality depth sensing with polarization cues. In *Proceedings of the IEEE international conference on computer vision*. 3370–3378.
- Hiroshi Kawasaki, Ryo Furukawa, Ryusuke Sagawa, and Yasushi Yagi. 2008. Dynamic scene shape reconstruction using a single structured light pattern. In *2008 IEEE conference on computer vision and pattern recognition*. Ieee, 1–8.
- Thomas P Koninckx and Luc Van Gool. 2006. Real-time range acquisition by adaptive structured light. *IEEE transactions on pattern analysis and machine intelligence* 28, 3 (2006), 432–445.
- Aldo Laurentini. 1994. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on pattern analysis and machine intelligence* 16, 2 (1994), 150–162.
- JJ Le Moigne and Allen Mark Waxman. 1988. Structured light patterns for robot mobility. *IEEE Journal on Robotics and Automation* 4, 5 (1988), 541–548.
- Marc Levoy and Pat Hanrahan. 2023. Light field rendering. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 441–452.
- Kunhong Li, Longguang Wang, Ye Zhang, Kaiwen Xue, Shunbo Zhou, and Yulan Guo. 2024. LoS: Local structure-guided stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 19746–19756.
- Haotong Lin, Sida Peng, Jingxiao Chen, Songyou Peng, Jiaming Sun, Minghuan Liu, Hujun Bao, Jiashi Feng, Xiaowei Zhou, and Bingyi Kang. 2024. Prompting Depth Anything for 4K Resolution Accurate Metric Depth Estimation. *arXiv*.
- Lahav Lipson, Zachary Teed, and Jia Deng. 2021. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *2021 International Conference on 3D Vision (3DV)*. IEEE, 218–227.
- Manuel Martinez and Rainer Stiefelhagen. 2013. Kinect Unleashed: Getting Control over High Resolution Depth Maps.. In *MVA*. 247–250.
- Minoru Maruyama and Shigeru Abe. 1993. Range sensing by projecting multiple slits with random cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15, 6 (1993), 647–651.
- Microsoft. 2025. Azure Kinect DK – Develop AI Models. <https://azure.microsoft.com/en-us/products/kinect-dk> Accessed: 2025-04-08.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tanicik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- Parsa Mirdehghan, Wenzheng Chen, and Kiriakos N. Kutulakos. 2018. Optimal Structured Light à La Carte. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Trans. Mach. Learn. Res.* 2024 (2024). <https://openreview.net/forum?id=a68SUt6zFt>
- Jeffrey L Posdamer and Martin D Altschuler. 1982. Surface measurement by space-encoded projected beam systems. *Computer graphics and image processing* 18, 1 (1982), 1–17.
- René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*. 12179–12188.
- Gernot Riegler, Yiyi Liao, Simon Donne, Vladlen Koltun, and Andreas Geiger. 2019. Connecting the dots: Learning representations for active monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 7624–7633.
- Ryusuke Sagawa, Hiroshi Kawasaki, Shota Kiyota, and Ryo Furukawa. 2011. Dense one-shot 3D reconstruction by detecting continuous regions with parallel line projection. In *2011 International Conference on Computer Vision*. IEEE, 1911–1918.
- Joaquim Salvi, Sergio Fernandez, Tomislav Pribanic, and Xavier Llado. 2010. A state of the art in structured light patterns for surface profilometry. *Pattern recognition* 43, 8 (2010), 2666–2680.
- Daniel Scharstein and Richard Szeliski. 2003. High-accuracy stereo depth maps using structured light. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, Vol. 1. IEEE, I–I.
- Zhanghao Sun, Wei Ye, Jinhui Xiong, Gyeongmin Choe, Jialiang Wang, Shuo Chen, and Rakesh Ranjan. 2023. Consistent direct time-of-flight video depth super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5075–5085.
- Yuichi Taguchi, Amit Agrawal, and Oncel Tuzel. 2012. Motion-aware structured light using spatio-temporal decodable patterns. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*. Springer, 832–845.
- Michael W Tao, Pratul P Srinivasan, Jitendra Malik, Szymon Rusinkiewicz, and Ravi Ramamoorthi. 2015. Depth from shading, defocus, and correspondence using light-field angular coherence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1940–1948.
- Michal Tölgýessy, Martin Dekan, L’uboš Chovanec, and Peter Hubinský. 2021. Evaluation of the azure kinect and its comparison to kinect v1 and kinect v2. *Sensors* 21, 2 (2021), 413.
- Piet Vuylsteke and André Oosterlinck. 1990. Range image acquisition with a single binary-encoded light pattern. *IEEE transactions on pattern analysis and machine intelligence* 12, 2 (1990), 148–164.
- Peter M Will and Keith S Pennington. 1971. Grid coding: A preprocessing technique for robot and machine vision. *Artificial Intelligence* 2, 3-4 (1971), 319–329.
- Haofei Xu and Juyong Zhang. 2020. AaNet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1959–1968.
- Yuhua Xu, Xiaoli Yang, Yushan Yu, Wei Jia, Zhaobi Chu, and Yulan Guo. 2022. Depth estimation by combining binocular stereo and monocular structured-light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1746–1755.
- Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. 2024. Depth anything v2. *Advances in Neural Information Processing Systems* 37 (2024), 21875–21911.
- Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian D. Reid. 2018. Unsupervised Learning of Monocular Depth Estimation and Visual Odometry With Deep Feature Reconstruction. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. Computer Vision Foundation / IEEE Computer Society, 340–349.
- Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. 2019. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 185–194.
- Li Zhang, Brian Curless, and Steven M Seitz. 2002. Rapid shape acquisition using color structured light and multi-pass dynamic programming. In *Proceedings. First International Symposium on 3D Data Processing Visualization and Transmission*. IEEE, 24–36.
- Yinda Zhang, Sameh Khamis, Christoph Rhemann, Julien Valentin, Adarsh Kowdle, Vladimir Tankovich, Michael Schoenberger, Shahram Izadi, Thomas Funkhouser, and Sean Fanello. 2018. Activestereonet: End-to-end self-supervised learning for active stereo systems. In *Proceedings of the european conference on computer vision (ECCV)*. 784–801.