Novelty Assessment Report

Paper: Light of Normals: Unified Feature Representation for Universal Photometric Stereo
PDF URL: https://openreview.net/pdf?id=LRA5z3oXOI
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-30

Abstract

Universal photometric stereo (PS) is defined by two factors: it must (i) operate under arbitrary, unknown lighting conditions and (ii) avoid reliance on specific illumination models. Despite progress (e.g., SDM UniPS), two challenges remain. First, current encoders cannot guarantee that illumination and normal information are decoupled. To enforce decoupling, we introduce LINO UniPS with two key components: (i) Light Register Tokens with light alignment supervision to aggregate point, direction, and environment lights; (ii) Interleaved Attention Block featuring global cross-image attention that takes all lighting conditions together so the encoder can factor out lighting while retaining normal-related evidence. Second, high-frequency geometric details are easily lost. We address this with (i) a Wavelet-based Dual-branch Architecture and (ii) a Normal-gradient Perception Loss. These techniques yield a \textbf{unified} feature space in which lighting is explicitly represented by register tokens, while normal details are preserved via wavelet branch. We further introduce PS-Verse, a large-scale synthetic dataset graded by geometric complexity and lighting diversity, and adopt curriculum training from simple to complex scenes. Extensive experiments show new state-of-the-art results on public benchmarks (e.g., DiLiGenT, Luces), stronger generalization to real materials, and improved efficiency; ablations confirm that Light Register Tokens + Interleaved Attention Block drive better feature decoupling, while Wavelet-based Dual-branch Architecture + Normal-gradient Perception Loss recover finer details.

Core Task Landscape

This paper addresses: **Surface Normal Estimation from Multi-Light Images under Unknown Illumination**

A total of **25 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Neural Inverse Rendering with Volumetric Representations**
- **Surface-Based Photometric Stereo under Unknown Lighting**
- **Single-View Inverse Rendering and Scene-Level Reconstruction**
- **Multi-View Stereo with Illumination Fusion**
- **Calibration and System Design for Multi-Light Imaging**
- **General Illumination Models and Theoretical Frameworks**
- **Application-Specific Multi-Light Analysis**
- **Learning-Based Calibration and Joint Estimation**
- **Methodological Surveys and Comparative Studies**

Complete Taxonomy Tree

- Surface Normal Estimation from Multi-Light Images under Unknown Illumination Survey Taxonomy
- Neural Inverse Rendering with Volumetric Representations
  - Tensor-Factorized Neural Fields (1 papers)
  - [1] Tensoir: Tensorial inverse rendering (Haian Jin, 2023) View paper
  - Multi-View Photometric Stereo with Neural Radiance Fields (1 papers)
  - [5] Ps-nerf: Neural inverse rendering for multi-view photometric stereo (Yang Wen-qi, 2022) View paper
  - Relightable Avatar and Human Performance Capture (3 papers)
  - [6] Rana: Relightable articulated neural avatars (Umar Iqbal, 2023) View paper
  - [15] Capturing relightable human performances under general uncontrolled illumination (Guannan Li, 2013) View paper
  - [24] Relightable Neural Human Assets from Multi-view Gradient Illuminations (Taotao Zhou, 2022) View paper
- Surface-Based Photometric Stereo under Unknown Lighting
  - Universal Photometric Stereo with Unified Feature Representations ★ (1 papers)
  - [0] Light of Normals: Unified Feature Representation for Universal Photometric Stereo (Anon et al., 2026) View paper
  - Diffusion-Based Multi-Light Synthesis for Intrinsic Estimation (1 papers)
  - [2] Neural LightRig: Unlocking Accurate Object Normal and Material Estimation with Multi-Light Diffusion (Zexin He, 2024) View paper
  - Shadow-Aware and Reflectance-Agnostic Normal Estimation (3 papers)
  - [3] DeepShaRM: Multi-View Shape and Reflectance Map Recovery Under Unknown Lighting (Kohei Yamashita, 2023) View paper
  - [4] Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions (Takahiro Okabe, 2009) View paper
  - [10] Shadow-aware Uncalibrated Photometric Stereo Network (Yingming Wang, 2022) View paper
  - Sparse and Interpolation-Based Photometric Stereo (1 papers)

    - [7] SPLINE-Net: Sparse photometric stereo through lighting interpolation and normal estimation networks (Qian Zheng, 2019) View paper
    - Intensity Profile and Isotropic Reflectance Methods (2 papers)
    - [18] From intensity profile to surface normal: photometric stereo for unknown light sources and isotropic reflectances (Feng Lu, 2015) View paper
    - [20] Recovery of surface normals and reflectance from different lighting conditions (Carme JuliÃ, 2008) View paper
- Single-View Inverse Rendering and Scene-Level Reconstruction
    - Single-Image Outdoor Inverse Rendering (1 papers)
    - [11] Outdoor inverse rendering from a single image using multiview self-supervision (Ye Yu, 2021) View paper
    - Surface Factorization under Single Unknown Illumination (1 papers)
    - [13] NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination (Zhang Xiuming, 2021) View paper
- Multi-View Stereo with Illumination Fusion
    - Multiview-Photometric Fusion under Uncalibrated Illumination (1 papers)
    - [22] Fusing multiview and photometric stereo for 3d reconstruction under uncalibrated illumination (Chenglei Wu, 2010) View paper
    - Gaussian Splatting for Illumination-Inconsistent Reconstruction (1 papers)
    - [16] GS-I^{3}: Gaussian Splatting for Surface Reconstruction from Illumination-Inconsistent Images (Wang Tengfei, 2025) View paper
- Calibration and System Design for Multi-Light Imaging
    - Automatic Light Direction Detection and Correction (1 papers)
    - [8] Automating RTI: Automatic light direction detection and correcting non-uniform lighting for more accurate surface normals (Matthew McGuigan, 2020) View paper
    - Near-Light Photometric Stereo and Low-Light Reconstruction (1 papers)
    - [21] Indoor scene reconstruction using near-light photometric stereo (J. Liao, 2016) View paper
    - Virtual Instrumentation and Multi-Illumination Dome Systems (1 papers)
    - [17] Virtual Instrument for a Multi-illumination Dome System (M. A. Khanesar, 2025) View paper
- General Illumination Models and Theoretical Frameworks (1 papers)
    - [9] Photometric stereo with general, unknown lighting (R Basri, 2007) View paper
- Application-Specific Multi-Light Analysis
    - Surface Analysis and Fingerprint Recognition (1 papers)
    - [12] Surface analysis and fingerprint recognition from multi-light imaging collections (McGuigan, 2023) View paper
    - Fine-Grained Change Detection under Varied Illumination (1 papers)
    - [25] Fgcnet: Learning to Detect Fine-Grained Changes Under Varied Illuminations (Qian Zhang, n.d.) View paper
    - Matte-Model Fitting for Multi-Light Reflectance (1 papers)
    - [23] Guided Robust Matte-Model Fitting for Accelerating Multi-light Reflectance (Ruggero Pintus, 2017) View paper
- Learning-Based Calibration and Joint Estimation (1 papers)
    - [14] Learning to estimate surface normal via deep photometric stereo networks (Tao Zhang, 2020) View paper
- Methodological Surveys and Comparative Studies (1 papers)
    - [19] Advanced Techniques for Reconstructing Objects and Scenes From Variations in Lighting and Viewpoint (Lichy, 2024) View paper

Narrative

Core task: surface normal estimation from multi-light images under unknown illumination. The field divides into several major branches that reflect different modeling philosophies and application contexts. Neural inverse rendering with volumetric representations (e.g., Tensoir[1], PS-NeRF[5]) leverages implicit scene encodings to jointly recover geometry and material properties, often targeting full 3D reconstruction. Surface-based photometric stereo under unknown lighting focuses more directly on normal estimation by modeling reflectance and light transport without requiring calibrated sources, spanning classical formulations (General Unknown Lighting[9]) and modern learning-based approaches (Deep Photometric Stereo[14], SPLINE-Net[7]). Single-view inverse rendering and scene-level reconstruction methods (NeRFactor[13], Outdoor Inverse Rendering[11]) tackle broader decomposition problems, while multi-view stereo with illumination fusion (Fusing Multiview Photometric[22]) integrates geometric cues across viewpoints. Calibration and system design branches address hardware setups (Automating RTI[8], Virtual Multiillumination Dome[17]), and application-specific analyses explore domains like fingerprint imaging (Multilight Fingerprint[12]) or relightable human capture (Relightable Human Performances[15], Relightable Neural Human[24]).

A particularly active line of work centers on universal or unified feature representations that handle diverse materials and lighting conditions without per-scene calibration. Light of Normals[0] sits within this branch, emphasizing a unified framework that generalizes across surface types under unknown illumination. This contrasts with methods like DeepShaRM[3], which explicitly models shadow and interreflection effects, or RANA[6], which focuses on robust aggregation of multi-light observations. Meanwhile, approaches such as Shadow-aware Photometric[10] and Attached Shadow Coding[4] tackle specific challenges like cast shadows, and recent works (GS-I3[16], Neural LightRig[2]) explore hybrid representations blending volumetric and surface-based reasoning. The central tension across these directions involves balancing model expressiveness—capturing complex light transport and material variation—with the need for practical generalization when lighting is uncalibrated, a challenge that Light of Normals[0] addresses through its unified feature design.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

The original leaf focuses on learning unified feature representations that decouple lighting and normal information for universal photometric stereo under arbitrary unknown illumination. Sibling subtopics address complementary challenges: diffusion-based synthesis generates auxiliary lighting conditions, intensity profile methods exploit pixel-level observations across lights, shadow-aware techniques handle occlusions and complex reflectance, and sparse methods tackle limited lighting samples. All share the goal of robust normal estimation from multi-light images without calibrated illumination.

Similarities: - All categories target surface normal estimation from multi-light images under unknown or uncalibrated illumination conditions - Methods across categories avoid requiring specific illumination models or calibrated light sources - Common emphasis on handling general, real-world lighting scenarios rather than controlled laboratory setups - Shared challenge of disentangling lighting effects from surface geometry information

Differences: - Original leaf learns unified feature spaces for decoupling, while diffusion-based methods synthesize auxiliary lighting conditions using generative priors - Intensity profile methods operate on pixel-level observations across lights for isotropic materials,

whereas the original leaf aims for broader material generality through feature learning - Shadow-aware methods explicitly model occlusions and avoid reflectance disentanglement, contrasting with the original leaf's focus on unified representations that separate lighting and normals - Sparse methods address limited lighting samples through interpolation, while the original leaf assumes sufficient multi-light observations for feature learning

Suggested Search Directions: - Hybrid approaches combining unified feature learning with diffusion priors for enhanced generalization - Integration of shadow-aware modeling into unified feature spaces to handle occlusions - Extension of unified representations to handle both dense and sparse lighting scenarios - Comparative studies on feature decoupling versus explicit physical modeling (shadows, reflectance) for robustness

Sibling Subtopics

- **Diffusion-Based Multi-Light Synthesis for Intrinsic Estimation** (leaves: 1, papers: 1)
- Scope: Methods leveraging diffusion priors to generate auxiliary multi-lighting conditions for geometry and material recovery.
- Exclude: Excludes direct photometric stereo without diffusion models; see sibling categories.
- **Intensity Profile and Isotropic Reflectance Methods** (leaves: 1, papers: 2)
- Scope: Approaches using pixel intensity profiles across varying lights to infer normals for general isotropic materials.
- Exclude: Excludes anisotropic reflectance methods and shadow-based techniques; see sibling categories.
- **Shadow-Aware and Reflectance-Agnostic Normal Estimation** (leaves: 1, papers: 3)
- Scope: Techniques explicitly modeling shadows or avoiding reflectance disentanglement for robust normal estimation under complex illumination.
- Exclude: Excludes methods requiring reflectance modeling or ignoring shadow information; see sibling categories.
- **Sparse and Interpolation-Based Photometric Stereo** (leaves: 1, papers: 1)
- Scope: Methods handling sparse lighting observations through interpolation networks or lighting estimation from limited samples.
- Exclude: Excludes dense lighting methods and calibrated setups; see sibling categories.

Contributions Analysis

Overall novelty summary. The paper proposes LINO UniPS, a universal photometric stereo method that decouples lighting from normal information using Light Register Tokens and Interleaved Attention Blocks, while preserving high-frequency geometric details via a wavelet-based dual-branch architecture. It resides in the 'Universal Photometric Stereo with Unified Feature Representations' leaf, which currently contains no other papers in the taxonomy. This places the work in a relatively sparse research direction within the broader surface-based photometric stereo landscape, suggesting the specific combination of unified feature learning and explicit lighting decoupling under arbitrary unknown illumination is not yet densely populated.

The taxonomy reveals several neighboring directions that address related but distinct challenges. The sibling leaf 'Diffusion-Based Multi-Light Synthesis' explores auxiliary lighting generation, while 'Shadow-Aware and Reflectance-Agnostic Normal Estimation' focuses on explicit shadow modeling or avoiding reflectance disentanglement. Nearby branches include volumetric neural rendering approaches and single-view inverse rendering methods. The paper's emphasis on unified feature spaces that factor out lighting while retaining normal evidence distinguishes it from shadow-centric methods and positions it between classical surface-based photometric stereo and modern neural inverse rendering paradigms.

Among thirteen candidates examined, the wavelet-based architecture contribution shows one refutable candidate from ten examined, indicating some prior work on frequency-domain processing for geometric detail preservation. The Light Register Tokens and Interleaved Attention Block contribution examined zero candidates, suggesting either limited semantic overlap in the search or a genuinely novel architectural design. The PS-Verse dataset contribution examined three candidates with none refutable, though the limited search scope means comprehensive dataset novelty assessment remains uncertain. The analysis explicitly covers top-K semantic matches and does not claim exhaustive coverage of all relevant prior work.

Given the sparse taxonomy leaf and limited search scope of thirteen candidates, the work appears to occupy a relatively underexplored niche combining explicit lighting decoupling with wavelet-based detail preservation. However, the single refutable candidate for the wavelet contribution and the constrained search scale suggest caution in drawing definitive novelty conclusions. A broader literature review covering more candidates and adjacent research areas would strengthen confidence in the originality assessment.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Light Register Tokens with Light Alignment supervision and Interleaved Attention Block

Description: The authors propose specialized learnable tokens (Point, Direction, Env) supervised by explicit Light Alignment loss to aggregate global illumination information, combined with an Interleaved Attention Block featuring global cross-image attention. This design enables the encoder to decouple lighting from surface normals and produce a unified feature representation.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 2: Wavelet-based Dual-branch Architecture and Normal-gradient Perception Loss

Description: The authors introduce a dual-branch architecture using discrete wavelet transform to preserve high-frequency information during feature extraction, paired with a confidence-weighted loss that emphasizes errors in high-frequency regions. These components work together to recover fine-scale geometric details that are typically lost in conventional up/downsampling operations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Wavelet pyramid recurrent structure-preserving attention network for single image super-resolution

URL: View paper

Brief Assessment

Wavelet Pyramid Recurrent[28] focuses on single image super-resolution (SISR) for 2D images, not 3D surface reconstruction or photometric stereo. The application domains and technical objectives are fundamentally different.

---

2. Diffusion-fof: Single-view clothed human reconstruction via diffusion-based fourier occupancy field

URL: View paper

Brief Assessment

Diffusion-FOF[34] applies wavelet transforms to 3D human reconstruction from single images, not photometric stereo normal estimation. The technical context and application domain differ fundamentally from the original paper's multi-light normal recovery task.

---

3. Deep discrete wavelet transform network for photometric stereo

URL: View paper

Prior Art Analysis

Deep Discrete Wavelet[31] demonstrates prior work that uses discrete wavelet transform (DWT) in a dual-branch architecture for photometric stereo to preserve high-frequency geometric details. The candidate paper explicitly describes decomposing images into low-frequency and high-frequency components via DWT, extracting features from both spatial and frequency domains through separate CNN-based modules, and fusing them for surface normal reconstruction. This directly addresses the same technical challenge of preserving fine-scale geometric details during feature extraction that the original paper claims as novel.

Evidence

Evidence 1 - **Rationale**: Both papers propose using DWT-based architectures to preserve high-frequency information for photometric stereo, demonstrating that this approach existed prior to the original paper's submission. - **Original**: we adopt a wavelet-based dual-branch architecture and a normal-gradient perception loss, which together substantially improve the reconstruction of fine-grained geometric details. - **Candidate**: this paper presents a discrete wavelet transform-based photometric stereo network (dwtps-net) to handle the input photometric stereo images in both the spatial and frequency domains. in dwtps-net, we extract shading features from images and also decompose the images using discrete wavelet transform ...

Evidence 2 - **Rationale**: Both papers explicitly state their core contribution as using wavelet-based dual-branch architectures to preserve high-frequency details during feature extraction, showing Deep Discrete Wavelet[31] proposed this approach earlier. - **Original**: to overcome these limitations, we propose a two-pronged approach: a wavelet-based dual-branch architecture to preserve information during feature extraction, and a normal gradient perception loss to guide the model's focus towards these fine details during training. - **Candidate**: to address the aforementioned issues, we propose a discrete wavelet transform-based photometric stereo network (dwtps-net) to better simultaneously tackle highfrequency structural details and low-frequency global information. dwtps-net employs wavelet transform to extract the shading cue features in...

---

4. Neural wavelet-domain diffusion for 3d shape generation, inversion, and manipulation

URL: View paper

Brief Assessment

Wavelet Domain Diffusion[29] applies wavelets to 3D shape generation via implicit representations and diffusion models, not to photometric stereo surface reconstruction. The technical domains and objectives are fundamentally different.

---

5. Micro-macro Wavelet-based Gaussian Splatting for 3D Reconstruction from Unconstrained Images

URL: View paper

Brief Assessment

Micro-macro Wavelet[30] focuses on 3D Gaussian splatting for novel view synthesis from unconstrained images, not photometric stereo for surface normal reconstruction. The wavelet usage serves different purposes in different domains.

---

6. Neural wavelet-domain diffusion for 3d shape generation

URL: View paper

Brief Assessment

Neural Wavelet Diffusion[27] applies wavelets to 3D shape generation using implicit representations and diffusion models, not to photometric stereo or surface normal reconstruction from multi-light images. The technical domains and objectives are fundamentally different.

---

7. A Wavelet-based Stereo Matching Framework for Solving Frequency Convergence Inconsistency

URL: View paper

Brief Assessment

Wavelet Stereo Matching[35] addresses stereo matching for disparity estimation, not photometric stereo for surface normal reconstruction. The technical domains and objectives are fundamentally different.

---

8. Wavenerf: Wavelet-based generalizable neural radiance fields

URL: View paper

Brief Assessment

WaveNeRF[32] focuses on novel view synthesis using neural radiance fields with wavelet decomposition for multi-view stereo, not surface normal reconstruction from photometric stereo. The technical domains and objectives differ fundamentally.

---

9. Endowave: Rational-wavelet 4d gaussian splatting for endoscopic reconstruction

URL: View paper

Brief Assessment

Endowave[26] applies wavelet decomposition to 4D Gaussian splatting for endoscopic reconstruction, focusing on separating high/low-frequency components in surgical video. The original paper targets photometric stereo for surface normal estimation with a dual-branch encoder architecture. These are fundamentally different application domains and architectural contexts.

---

10. 3D-WAG: Hierarchical Wavelet-Guided Autoregressive Generation for High-Fidelity 3D Shapes

URL: View paper

Brief Assessment

3D-WAG[33] focuses on 3D shape generation using wavelet transforms for autoregressive modeling, not on surface reconstruction or photometric stereo tasks that preserve high-frequency geometric details during feature extraction.

Contribution 3: PS-Verse dataset with curriculum training strategy

Description: The authors construct a large-scale synthetic dataset containing 100,000 scenes with 17,805 textured 3D models, graded by geometric complexity (four levels plus normal mapping) and lighting diversity. They employ curriculum learning that progresses from simple to complex scenes, enhancing model robustness and generalization under challenging real-world lighting conditions.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Multi-scale and attention training of uncalibrated photometric stereo networks

URL: View paper

Brief Assessment

Multiscale Attention Photometric[36] does not describe any dataset construction or curriculum training strategy. The paper focuses on multi-scale and attention mechanisms for uncalibrated photometric stereo network architecture, without addressing synthetic dataset creation graded by complexity levels or curriculum learning approaches.

---

2. Refined Tensorial Radiance Field: Harnessing coordinate based networks for novel view synthesis from sparse inputs

URL: View paper

Brief Assessment

Refined Tensorial Radiance[38] focuses on novel view synthesis using multi-plane encoding and coordinate-based networks for sparse-input 3D/4D reconstruction. It does not address photometric stereo, surface complexity grading for PS datasets, or curriculum training strategies for PS tasks.

---

3. Task-Specific Near-Field Photometric Stereo for Measuring Metal Surface Texture

URL: View paper

Brief Assessment

Near-Field Metal Texture[37] focuses on task-specific training for machined metal surfaces using real datasets for specific material-machine combinations, not a large-scale synthetic dataset graded by geometric complexity with curriculum learning.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Light of Normals: Unified Feature Representation for Universal Photometric Stereo View paper
- [1] Tensoir: Tensorial inverse rendering View paper
- [2] Neural LightRig: Unlocking Accurate Object Normal and Material Estimation with Multi-Light Diffusion View paper
- [3] DeepShaRM: Multi-View Shape and Reflectance Map Recovery Under Unknown Lighting View paper
- [4] Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions View paper
- [5] Ps-nerf: Neural inverse rendering for multi-view photometric stereo View paper
- [6] Rana: Relightable articulated neural avatars View paper
- [7] SPLINE-Net: Sparse photometric stereo through lighting interpolation and normal estimation networks View paper
- [8] Automating RTI: Automatic light direction detection and correcting non-uniform lighting for more accurate surface normals View paper
- [9] Photometric stereo with general, unknown lighting View paper
- [10] Shadow-aware Uncalibrated Photometric Stereo Network View paper
- [11] Outdoor inverse rendering from a single image using multiview self-supervision View paper
- [12] Surface analysis and fingerprint recognition from multi-light imaging collections View paper
- [13] NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination View paper
- [14] Learning to estimate surface normal via deep photometric stereo networks View paper
- [15] Capturing relightable human performances under general uncontrolled illumination View paper
- [16] GS-I^{3}: Gaussian Splatting for Surface Reconstruction from Illumination-Inconsistent Images View paper
- [17] Virtual Instrument for a Multi-illumination Dome System View paper
- [18] From intensity profile to surface normal: photometric stereo for unknown light sources and isotropic reflectances View paper
- [19] Advanced Techniques for Reconstructing Objects and Scenes From Variations in Lighting and Viewpoint View paper
- [20] Recovery of surface normals and reflectance from different lighting conditions View paper
- [21] Indoor scene reconstruction using near-light photometric stereo View paper
- [22] Fusing multiview and photometric stereo for 3d reconstruction under uncalibrated illumination View paper
- [23] Guided Robust Matte-Model Fitting for Accelerating Multi-light Reflectance View paper
- [24] Relightable Neural Human Assets from Multi-view Gradient Illuminations View paper
- [25] Fgcnet: Learning to Detect Fine-Grained Changes Under Varied Illuminations View paper
- [26] Endowave: Rational-wavelet 4d gaussian splatting for endoscopic reconstruction View paper
- [27] Neural wavelet-domain diffusion for 3d shape generation View paper
- [28] Wavelet pyramid recurrent structure-preserving attention network for single image super-resolution View paper
- [29] Neural wavelet-domain diffusion for 3d shape generation, inversion, and manipulation View paper
- [30] Micro-macro Wavelet-based Gaussian Splatting for 3D Reconstruction from Unconstrained Images View paper
- [31] Deep discrete wavelet transform network for photometric stereo View paper
- [32] Wavenerf: Wavelet-based generalizable neural radiance fields View paper
- [33] 3D-WAG: Hierarchical Wavelet-Guided Autoregressive Generation for High-Fidelity 3D Shapes View paper
- [34] Diffusion-fof: Single-view clothed human reconstruction via diffusion-based fourier occupancy field View paper
- [35] A Wavelet-based Stereo Matching Framework for Solving Frequency Convergence Inconsistency View paper
- [36] Multi-scale and attention training of uncalibrated photometric stereo networks View paper
- [37] Task-Specific Near-Field Photometric Stereo for Measuring Metal Surface Texture View paper
- [38] Refined Tensorial Radiance Field: Harnessing coordinate based networks for novel view synthesis from sparse inputs View paper