

# Single observation for light-field editing

Qiong-Hua Wang\*

## Abstract

Inferring explicit illumination cues from a single sparse observation fundamentally decouples scene representation from its captured environment. By shifting illumination from a static, baked-in attribute to a parametrically editable representation, the framework establishes a novel paradigm for scalable, photometrically consistent immersive displays under dynamic real-world conditions.

**Keywords:** Light field display, Immersive interaction, Neural illumination, One single observation view, Photometric consistency

The photorealistic integration of virtual content into real-world scenes has long been recognised as a central challenge for immersive display technologies, including virtual reality (VR), augmented reality (AR), and mixed reality (MR)<sup>1-3</sup>. The theoretical core of these systems is the plenoptic function, which provides a comprehensive description of the spatial and angular distributions of light rays<sup>4</sup>. Faithful reproduction of a dynamic light field requires a rigorous solution to the rendering equation, which accounts for the intricate coupling between incident illumination, scene geometry, and material reflectance<sup>5,6</sup>. This interplay ultimately determines photometric consistency, that is, whether virtual objects appear visually anchored within a scene or as perceptually detached overlays<sup>7</sup>.

However, the capture and reproduction of high-fidelity light fields under dynamic conditions present a formidable inverse problem. Conventional approaches typically rely on hardware-intensive acquisition setups such as light stages or dense multi-view camera arrays to adequately sample environmental illumination<sup>8,9</sup>. However, in many immersive scenarios, available observations are constrained to a sparse viewpoint, often corresponding to a single perspective from a wearable headset. Neural

radiance fields (NeRFs), while effective for view synthesis<sup>10</sup>, often exhibit baked-in illumination, where lighting is entangled with scene geometry and appearance<sup>11-13</sup>. This entanglement makes it difficult to adapt rendered content to changing real-world lighting, resulting in visual inconsistencies whenever the observer or illumination conditions change. These constraints reveal a fundamental mismatch between the assumptions of many illumination-dependent rendering strategies and the sensing realities of mixed-reality systems.


Although many existing approaches continue to rely on access to complete illumination information, such as panoramic environmental maps or densely sampled views, these assumptions rarely hold true in practical mixed-reality settings. The recent work of Hong et al. provides a significant contribution by shifting the focus from exhaustive illumination acquisition to a perceptually sufficient framework for estimating and editing illumination from a single observation<sup>14-18</sup>. Recognising that exact recovery of global illumination from a single view is an ill-posed problem, the authors demonstrated that the essential elements of illumination can be inferred from minimal visual input. As illustrated in Fig. 1, this method bypasses the intractable task of full-environment recovery by utilising minimal visual input to drive a generative process that updates the neural light fields in a controllable

Correspondence: Qiong-Hua Wang (qionghua@buaa.edu.cn)  
School of Instrumentation and Optoelectronic Engineering, Beihang University, Beijing 100191, China

© The Author(s) 2026



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.



**Fig. 1** Neural illumination editing pipeline for adaptive light field display.

and photometrically consistent manner. The conceptual workflow introduced by Hong et al. demonstrated a paradigm shift from physical reconstruction to generative synthesis. Based on a single sparse observation, the system inferred compact and explicit illumination cues (intensity, direction, and semantic descriptors). These cues act as physical priors to guide a generative network that dynamically reconstructs and updates the neural light field, ensuring that the virtual content exhibits shading and shadows that are photometrically consistent with the evolving real-world environment.

The methodological novelty of the system lies in its ability to bridge the gap between explicit optical physics and generative semantic control. Rather than considering illumination estimation as a numerical end, the authors employed a computational optical perception (COP) module to translate sparse signals into structured parametric representations, effectively using semantic descriptors, such as text-based prompts, to constrain the solution space for light field editing. By injecting these priors into a generative synthesis engine, the framework ensures that the updated scene appearance, including complex shadows and view-dependent highlights, responds coherently to illumination variations while preserving strict geometric consistency across views<sup>19</sup>. This represents a significant departure from purely image-based style transfer, which offers a physically grounded approach to neural relighting.

The significance of these results extends beyond improvements in visual quality. This study addresses a fundamental practical constraint in mixed-reality systems by demonstrating that high-fidelity, photometrically

consistent light-field editing can be achieved with an extremely sparse observation budget. This establishes a scalable pathway for adaptive light-field rendering without relying on specialised capture hardware or dense sensing. This capability is particularly relevant for real-world immersive displays, where both sensing and computation are inherently limited.

More broadly, this study is conducted at a time of growing demand for lightweight and energy-efficient computational imaging solutions. As immersive display hardware continues to evolve towards compact, glasses-free 3D platforms<sup>20,21</sup>, the demand for algorithms that maximise the perceptual impact per compute unit has become increasingly critical. The hybrid paradigm demonstrated herein, which integrates physical illumination priors with the inference capacity of generative models, is likely to inspire a new wave of research in responsive digital light fields. This approach holds particular promise for applications ranging from automotive heads-up displays to mobile AR, where maintaining visual coherence with the physical world is essential for both user immersion and comfort.

#### **Data availability**

All data are available from the corresponding authors upon reasonable request.

#### **Conflict of interest**

The author declares no competing interests.

Received: 19 March 2026 Revised: 03 April 2026 Accepted: 07 April 2026

Published online: 09 June 2026

## References

1. Xiong, J. H. et al. Augmented reality and virtual reality displays: emerging technologies and future perspectives. *Light: Science & Applications* **10**, 216 (2021).
2. Kim, T. S. et al. Future trends of display technology: micro-LEDs toward transparent, free-form, and near-eye displays. *Light: Science & Applications* **14**, 335 (2025).
3. Kress, B. C. & Pace, M. Holographic optics in planar optical systems for next generation small form factor mixed reality headsets. *Light: Advanced Manufacturing* **3**, 42 (2022).
4. Levoy, M. & Hanrahan, P. Light field rendering. Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques. New Orleans: ACM, 31-42 (1996).
5. Kajiya, J. T. The rendering equation. *ACM SIGGRAPH Computer Graphics* **20**, 143-150 (1986).
6. Cook, R. L. & Torrance, K. E. A reflectance model for computer graphics. *ACM Transactions on Graphics* **1**, 7-24 (1982).
7. Mania, K. et al. The effect of visual and interaction fidelity on spatial cognition in immersive virtual environments. *IEEE Transactions on Visualization and Computer Graphics* **12**, 396-404 (2006).
8. Debevec, P. The light stages and their applications to photoreal digital actors. Proceedings of the SIGGRAPH Asia. Singapore, Singapore: SIGGRAPH, 2012, 1-6.
9. Gardner, A. et al. Linear light source reflectometry. *ACM Transactions on Graphics* **22**, 749-758 (2003).
10. Mildenhall, B. et al. NeRF: representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**, 99-106 (2022).
11. Zhang, X. M. et al. NeRFactor: neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics* **40**, 237 (2021).
12. Srinivasan, P. P. et al. NeRV: neural reflectance and visibility fields for relighting and view synthesis. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. Nashville, TN, USA: IEEE, 2021, 7491-7500.
13. Verbin, D. et al. Ref-NeRF: structured view-dependent appearance for neural radiance fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **47**, 9426-9437 (2025).
14. Hong, X. Y. et al. Single-view neural illumination estimation and editing for dynamic light field display. *Light: Science & Applications* **15**, 147 (2026).
15. Ramamoorthi, R. & Hanrahan, P. A signal-processing framework for inverse rendering. Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques. Los Angeles: ACM, 117-128 (2001).
16. Barbastathis, G., Ozcan, A. & Situ, G. H. On the use of deep learning for computational imaging. *Optica* **6**, 921-943 (2019).
17. Mengu, D. et al. At the intersection of optics and deep learning: statistical inference, computing, and inverse design. *Advances in Optics and Photonics* **14**, 209-290 (2022).
18. Suo, J. L. et al. Computational imaging and artificial intelligence: the next revolution of mobile vision. *Proceedings of the IEEE* **111**, 1607-1639 (2023).
19. Zeng, C. et al. Relighting neural radiance fields with shadow and highlight hints. Proceedings of the ACM SIGGRAPH 2023 Conference. Los Angeles, CA, USA: ACM, 2023, 73.
20. Park, J. H. & Lee, B. Holographic techniques for augmented reality and virtual reality near-eye displays. *Light: Advanced Manufacturing* **3**, 137-150 (2022).
21. Solomashenko, A. B. et al. Industrial applications of AR headsets: a review of the devices and experience. *Light: Advanced Manufacturing* **6**, 23 (2025).