Toward Real-World 3D Rendering

□ 권의영, 오지형 / 중앙대학교

Abstract

Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have achieved notable success in photorealistic 3D reconstruction and novel view synthesis but rely heavily on high-quality multi-view images, which limits their robustness under real-world degradations such as noise, blur, low-resolution (LR), and weather artifacts. To address this, 3D Low-Level Vision (3D LLV) extends classical 2D restoration tasks such as deblurring and weather degradation removal into the 3D domain. This survey formalizes the problem of degradation-aware rendering and outlines key challenges related to spatio-temporal consistency and ill-posed optimization. It categorizes recent approaches that integrate LLV into neural rendering frameworks and examines their applicability to domains including autonomous driving, AR/VR, and robotics. By reviewing representative methods, datasets, and evaluation protocols, this work identifies 3D LLV as a fundamental direction for robust 3D scene reconstruction under real-world conditions.

I. Introduction

Vision is a fundamental human sense, critical for perceiving and interacting with the 3D world, and underpins applications such as the metaverse, AR/VR, and robotics. The growing demand for accurate 3D perception has driven progress in learning-based 3D representations, with Neural Radiance Fields (NeRF) [1] and 3D Gaussian Splatting (3DGS) [2] achieving strong results in photorealistic reconstruction and novel view synthesis. These methods jointly model appearance and geometry using multi-view images and camera

※ Note: This article is adapted from the author's previously submitted survey paper currently under review by an IEEE journal (preprint available at arXiv:2506.16262). The content has been paraphrased and condensed, and certain figures and mathematical e-pressions may resemble those in the original manuscript.

poses, enabling faithful rendering from unseen viewpoints. However, most approaches assume access to clean, high-resolution (HR) multi-view inputs [3]-[6], which is rarely the case in real-world environments where degradations such as noise, blur, low-resolution (LR), and weather artifacts are common [7], [8].

To address this, 3D Low-Level Vision (3D LLV) has emerged by extending classical tasks like deblurring and weather removal into the 3D domain [9]. In 2D LLV, degradations are modeled as $y = H_x + n$, where y is the observed degraded image, x is the clean (latent) image, H is the degradation operator (e.g., blur, downsampling), and n is the additive noise.

In contrast, 3D LLV aims to recover a scene representation s from degraded N multi-view images $\{y_i\}_{i=1}^{N}$ and corresponding camera parameters $\{c_i\}_{i=1}^{N}$. This is formalized by a degradation-aware rendering function R_D which maps the 3D scene representation s and the camera pose c_i to a degraded image estimate $\hat{y}_i$:

$$\hat{y}_i = R_D(s, c_i). \tag{1}$$

As shown in Fig. 1(a), the training phase supervises the rendered image $\hat{y}_i$ by comparing it with the degraded observation y_i. During inference, a clean rendering function R_c, which takes the shared 3D scene representation s and the camera pose c_i as input, is used to generate high-quality outputs, as illustrated in Fig. 1(b).

Two degradation strategies are considered in this formulation. Sequence-based modeling $R_{D_{seq}}$, a clean image is first rendered and then degraded using a view-specific degradation operator H_i, which may include blur, noise, or downsampling:

$$R_{D_{seq}}(s, c_i) = H_i R_c(s, c_i). \tag{2}$$

Composition-based modeling $R_{D_{comp}}$, clean and degraded renderings are fused using a differentiable composition function C:

$$R_{D_{comp}}(s, c_i) = C(R_c(s, c_i),\ R_D(s, c_i)). \tag{3}$$

This formulation enables flexible modeling of complex corruption by disentangling content and degradation. Moreover, since only degraded observations are available, the model must maintain

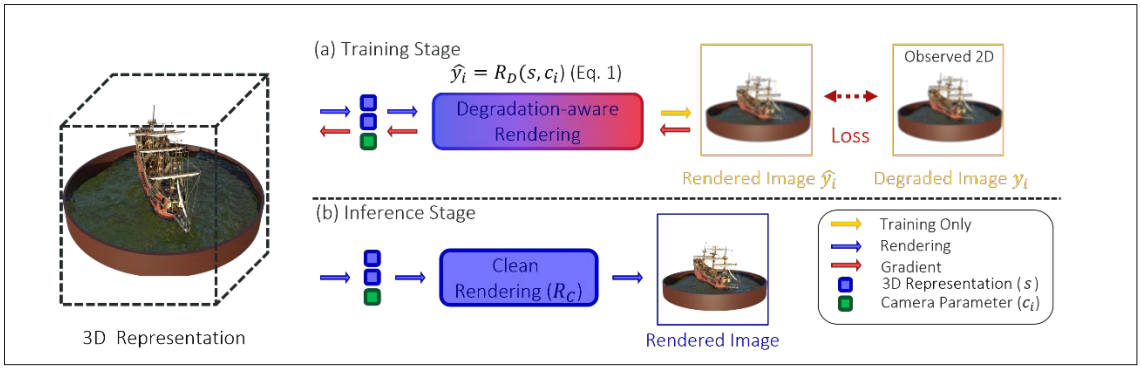

<Figure 1> Degradation-aware rendering pipeline

geometric and temporal consistency across views, making the problem more ill-posed than in 2D. As a result, 3D LLV shifts from pixel-level restoration to degradation-aware rendering that integrates 3D structure and degradation priors. This enables robust reconstruction under real-world degradations, unlike conventional neural rendering pipelines that rely on clean inputs. Such capability is essential in domains like autonomous driving [10], AR/VR [11], and robotics [12], where perception must remain reliable under noise, compression, or visibility issues. These scenarios require 3D LLV frameworks that ensure consistent and accurate scene reconstruction under adverse conditions. This survey formalizes the 3D LLV task and surveys recent approaches that integrate LLV techniques into neural rendering. Section 2 defines the problem and outlines key challenges in 3D LLV. Section 3 reviews deblurring and weather removal in the 3D domain. Section 4 introduces representative datasets and evaluation protocols. Section 5 concludes the survey.

II. Problem Definition and Challenges

NeRF [1] and 3D Gaussian Splatting (3DGS) [2] have achieved impressive results in 3D reconstruction using high-quality images. However, real-world images often suffer from degradations that significantly impair reconstruction accuracy and spatiotemporal consistency [13].

Multi-view pose estimation methods such as SfM [14]-[16] and COLMAP [17], [18] are particularly vulnerable to such degradations. Blur, LR, or illumination changes hinder feature detection and matching, leading to inaccurate poses that degrade NeRF and 3DGS initialization [3].

Understanding how different degradation types affect this process is crucial for enhancing model robustness. Fig. 2 summarizes three major degradations: motion blur, defocus blur, and weather-induced artifacts. These factors disrupt feature alignment, pose estimation, and multi-view consistency, degrading 3D reconstruction quality. To address this, recent 3D LLV approaches incorporate deblurring and weather degradation removal into neural rendering.

Deblurring restores high-frequency structures lost to motion or defocus blur [19]-[21], both of which obscure geometry and alignment.

Weather degradation removal tackles visibility loss from haze, rain, or snow [22]-[24], where scattering and absorption reduce contrast and introduce lighting inconsistencies. Such degradations often co-occur, posing severe challenges to neural rendering pipelines that assume clean observations.

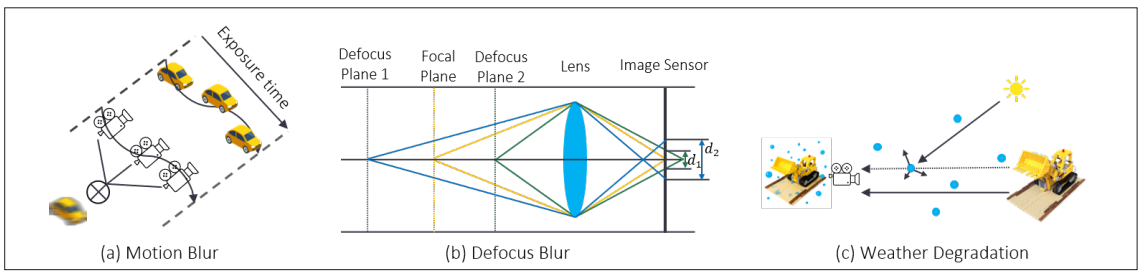

<Figure 2> Degradation Factors in Real-World 3D Scene Capture

III. Low-Level Vision for Robust Rendering

This section introduces representative 3D LLV methods for addressing various degradation conditions. The methods are presented in the order of deblurring and weather degradation removal. These methods are organized as shown in Fig. 3.

1. Deblurring in 3D LLV

To mitigate blur-induced failures in 3D reconstruction, recent methods adopt diverse deblurring strategies depending on the blur type, including trajectory-based, event-based, defocus-based, and motion-defocus hybrid approaches, as illustrated in Fig. 4.

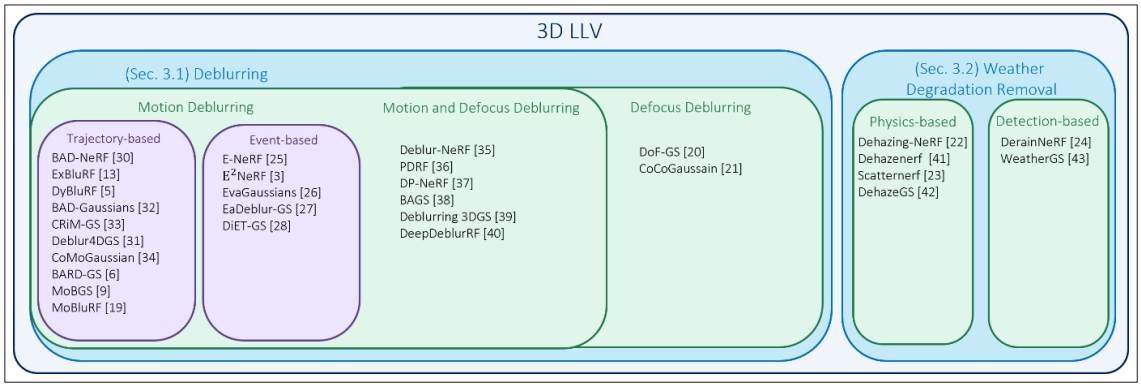

<Figure 3> Taxonomy

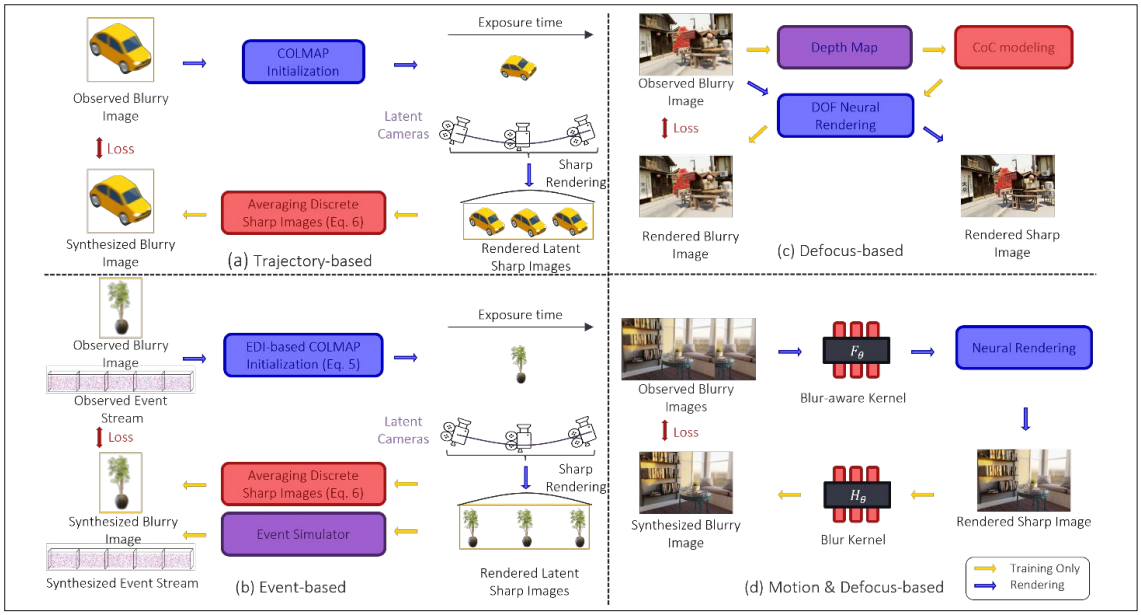

<Figure 4> Deblurring Strategies in Neural Rendering

1) Motion Deblurring

To address this, event-based approaches leverage the high temporal resolution of event cameras. E-NeRF [25] learns by comparing brightness changes derived from events and predicted images. E^2NeRF [3], EvaGaussians [26], EaDeblur-GS [27], and DiET-GS [28] utilize the Event-based Double Integral (EDI) model [29] to initialize camera poses under motion blur and recover latent sharp images. EDI formulates the blurred image B as follows:

$$B = I(s) \cdot \frac{1}{\tau}\int_{s-\tau/2}^{s+\tau/2} \exp\big(c \cdot E(t)\big)\, dt \qquad (4)$$

where I is the sharp image at center exposure time s, τ is the exposure duration, E(t) is the accumulated event stream, and c is the contrast threshold. A sharp image at arbitrary time t is given by:

$$I(t)=I(s)\cdot\exp\big(c \cdot E(t)\big) \qquad (5)$$

The reconstructed sharp image sequence is used to initialize camera poses and generate point clouds using COLMAP. Blurred images are then synthesized by averaging rendered sharp views along the estimated camera trajectory:

$$B = \widehat{\frac{1}{n}\sum_{i=1}^{n} I(t_i)} \qquad (6)$$

An additional event loss is introduced to enforce consistency between the predicted and actual brightness changes. DiET-GS [28] further enhances rendering quality through periodic consistency regularization and diffusion priors.

Trajectory-based 3D deblurring methods model the continuous motion of the camera or object during exposure and synthesize motion blur by averaging sharp images rendered along the estimated path. This follows the same formulation as Eq. (6), where a blurred image is approximated by temporally averaging multiple sharp images at sampled time steps t_i. BAD-NeRF [30], DyBluRF [5], Deblur4DGS [31], and BAD-Gaussians [32] define the camera trajectory by interpolating between two learnable poses in SE(3), jointly optimizing the scene representation and motion path. ExBluRF [13] extends the trajectory modeling approach by employing higher-order Bézier curves for smoother motion representation. CRiM-GS [33] and CoMoGaussian [34] introduce neural ODEs to represent temporally continuous motion with greater flexibility.

Meanwhile, video-based methods such as MoBluRF [19], DyBluRF [5], BARD-GS [6], and MoBGS [9] utilize masks to separate dynamic and static regions, optimizing each independently.

2) Defocus Deblurring

Neural rendering typically assumes that all scene points are in focus, relying on sharp and depth-consistent views for accurate reconstruction [20]. However, real-world imaging with shallow depth-of-field (DOF) and large apertures often induces defocus blur, where out-of-focus points form finite-sized blur circles, known as Circles of Confusion (CoC) [20, 21]. The CoC radius, derived from the thin-lens model, increases as object distance deviates from the focal distance, resulting in spatially varying blurs that severely degrades geometric consistency near depth discontinuities [20]. To address this, recent works incorporate physically grounded defocus modeling into neural rendering by estimating CoC from depth

and camera parameters, simulating depth-aware blur during training, and optionally include optimization of focus and aperture settings [20, 21]. DOF-GS [20] introduces learnable per-view focal distance and aperture factors, applying CoC-based blur kernels to 2D-projected Gaussians and supervising sharp region localization using an all-in-focus image. CoCoGaussian [21] renders multiple CoC Gaussians in a circular pattern around each base Gaussian and blends them through weighted summation to synthesize defocused images.

3) Joint Motion and Defocus Deblurring

In real-world scenarios, motion blur and defocus blur frequently co-occur due to fast camera movement and shallow depth of field (DOF), violating the assumptions of sharpness and view consistency in neural rendering pipelines and resulting in degraded reconstruction quality [36, 39]. To address this challenge, several approaches have incorporated joint modeling of motion and defocus blur within the rendering framework. Deblur-NeRF [35] simulates compound blur by aggregating rays sampled from nearby positions and jointly optimizes both ray offsets and the NeRF parameters. PDRF [36] follows a coarse-to-fine strategy by estimating blur fields and LR geometry, then refining with HR voxel representations. DP-NeRF [37] applies rigid-body transformations in SE(3) to simulate spatially consistent motion blur while preserving view alignment. BAGS [38] predicts per-pixel degradation masks and blur kernels using a 2D prior network to guide the optimization of 3D GS. Deblurring 3DGS [39] encodes blur by deforming Gaussian covariance matrices, allowing for anisotropic spread without explicit kernel modeling. DeepDeblurRF [40] integrates a pre-trained 2D deblurring network into radiance field optimization, using enhanced 2D outputs as supervisory signals.

2. Weather Degradation Removal in 3D LLV

NeRF and 3DGS are highly sensitive to input quality, and adverse weather such as rain, haze, or snow often disrupts multi-view consistency and degrades structural reconstruction [23, 30]. To address this, existing approaches can be categorized into physics-

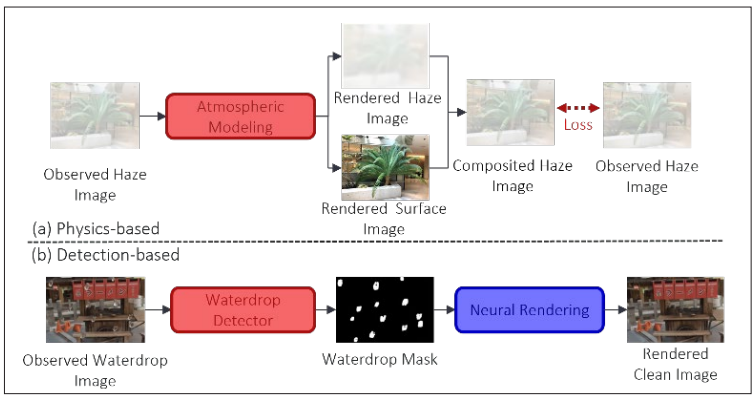

<Figure 5> Weather Degradation Removal Strategies in Neural Rendering

based and detection-based methods, as illustrated in Fig. 5.

1) Physics-based Methods

These methods explicitly incorporate physical haze models into volumetric rendering to disentangle weather effects from scene radiance [22, 23, 41, 42]. Koschmieder-based approaches, such as DehazeNeRF [41] and ScatterNeRF [23], model the observed radiance as a weighted combination of the scene radiance attenuated by atmospheric scattering and the airlight component dependent on scene depth.

DehazeNeRF generalizes this using the Radiative Transfer Equation, decomposing accumulated radiance into surface and haze components with physically grounded priors such as the Dark Channel Prior (DCP) [51]. ScatterNeRF simplifies modeling by explicitly separating haze and scene volumes and rendering them via two MLPs, applying entropy-based regularization to enforce volumetric separation.

ASM-based methods, including Dehazing-NeRF [22] and DehazeGS [42], estimate atmospheric parameters and integrate them into the rendering process. Dehazing-NeRF adopts a dual-branch network to jointly learn haze parameters and scene radiance, regularized by contrast and consistency losses. DehazeGS predicts transmittance from Gaussian depth using 1D convolutions. The final blending of latent Gaussians with global airlight is guided by pseudo-depth supervision and regularized using the DCP and the Bright Channel Prior (BCP) [53].

2) Detection-based Methods

These methods focus on localized degradations such as raindrops and snowflakes, which cause occlusion artifacts across views. DerainNeRF [24] uses a pretrained raindrop detector to generate binary masks that exclude occluded pixels from NeRF training. View-consistent occlusions are further refined using mean attention maps. WeatherGS [43] separates weather effects into particle-based and lens-based degradations. It detects each via specialized filters and applies the resulting masks in training, ensuring that corrupted regions do not interfere with scene learning.

IV. Dataset and Metrics

1. Dataset

We organize datasets based on a wide range of publicly available datasets commonly used in 3D reconstruction, neural rendering, and novel view synthesis.

NeRF-synthetic [1] contains eight synthetic Blender scenes with 100-300 images and full supervision. The scenes feature clean geometry, consistent lighting, and 360° or hemispherical viewpoints.

BlendedMVS [44] provides 113 scenes that mix real backgrounds with rendered objects, with 150-200 HR images per scene. Depth maps and camera poses reconstructed via COLMAP are included.

NSVF Synthetic [45] includes eight synthetic scenes with complex geometry and lighting, each with 400 images rendered with global illumination. Ground-truth depth and pose annotations are provided.

Deblur-NeRF [35] introduces 31 dynamic Blender-rendered scenes with 27 to 53 blurred multi-view images. Motion blur is simulated via pose interpola-

tion and defocus via depth-of-field rendering.

Mip-NeRF 360 [46] captures nine indoor and outdoor real scenes with hundreds of images taken along 360° camera trajectories. The dataset emphasizes unbounded scene modeling and background consistency.

2. Metrics

We summarize representative metrics used to assess the visual quality and temporal consistency of 3D LLV outputs.

Full-reference metrics, including PSNR [47], SSIM [48], and LPIPS [49], are widely used when ground-truth images are available. PSNR measures pixel-wise fidelity, SSIM captures structural similarity in luminance and texture, and LPIPS quantifies perceptual distance using deep features from pretrained networks. Temporal consistency is evaluated using temporal Optical Flow (tOF) [50], which measures flow discrepancies between consecutive frames to assess motion coherence in dynamic scenes. Together, these metrics provide a comprehensive framework for evaluating spatial fidelity, perceptual realism, and temporal stability in 3D LLV.

V. Conclusion

This survey formalizes the emerging field of 3D LLV, which incorporates super-resolution, deblurring, and weather degradation removal into neural rendering. These capabilities enable robust 3D reconstruction under real-world conditions. We categorized recent methods, summarized key datasets and metrics, and highlighted challenges related to spatiotemporal consistency and degradation modeling. 3D LLV offers a promising direction for enhancing the reliability of neural rendering in practical applications such as autonomous driving, AR/VR, and robotics.

참 고 문 헌

[1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” Communications of the ACM, 2021.

[2] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering.” ACM Transactions on Graphics, 2023.

[3] Y. Qi, L. Zhu, Y. Zhang, and J. Li, “E2nerf: Event enhanced neural radiance fields from blurry images,” in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023.

[4] M. Cannici and D. Scaramuzza, “Mitigating motion blur in neural radiance fields with events and frames,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.

[5] H. Sun, X. Li, L. Shen, X. Ye, K. Xian, and Z. Cao, “Dyblurf: Dynamic neural radiance fields from blurry monocular video,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024.

[6] Y. Lu, Y. Zhou, D. Liu, T. Liang, and Y. Yin, “Bard-gs: Blur-aware reconstruction of dynamic scenes via gaussian splatting,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2025.

참 고 문 헌

[7] K. Zhou, W. Li, Y. Wang, T. Hu, N. Jiang, X. Han, and J. Lu, "Nerflix: High-quality neural view synthesis by learning a degradation-driven inter-viewpoint mixer," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023.

[8] M. Li, M. Lu, X. Li, and S. Zhang, "Rustnerf: Robust neural radiance field with low-quality images," arXiv preprint arXiv:2401.03257, 2024.

[9] M.-Q. V. Bui, J. Park, J. L. G. Bello, J. Moon, J. Oh, and M. Kim, "Mobgs: Motion deblurring dynamic 3d gaussian splatting for blurry monocular video," arXiv preprint arXiv:2504.15122, 2025.

[10] M. Khan, H. Fazlali, D. Sharma, T. Cao, D. Bai, Y. Ren, and B. Liu, "Autosplat: Constrained gaussian splatting for autonomous driving scene reconstruction," in Proceedings of the IEEE International Conference on Robotics and Automation, 2025.

[11] L. Xu, V. Agrawal, W. Laney, T. Garcia, A. Bansal, C. Kim, S. Rota Bulò, L. Porzi, P. Kontschieder, A. Božič et al., "Vr-nerf: High-fidelity virtualized walkable spaces," in Proceedings of the ACM SIGGRAPH Asia 2023 Conference, 2023.

[12] Q. Dai, Y. Zhu, Y. Geng, C. Ruan, J. Zhang, and H. Wang, "Graspnerf: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf," in Proceedings of the IEEE International Conference on Robotics and Automation, 2023.

[13] D. Lee, J. Oh, J. Rim, S. Cho, and K. M. Lee, "Exblurf: Efficient radiance fields for extreme motion blurred images," in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023.

[14] S. Agarwal, Y. Furukawa, N. Snavely, I. Simon, B. Curless, S. M. Seitz, and R. Szeliski, "Building rome in a day," Communications of the ACM, 2011.

[15] D. Crandall, A. Owens, N. Snavely, and D. Huttenlocher, "Discrete-continuous optimization for large-scale structure from motion," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2011.

[16] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: exploring photo collections in 3d," in ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques, 2006.

[17] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2016.

[18] J. L. Schonberger, E. Zheng, J.-M. Frahm, and M. Pollefeys, "Pixelwise view selection for unstructured multi-view stereo," in Proceedings of the European Conference on Computer Vision, 2016.

[19] M.-Q. V. Bui, J. Park, J. Oh, and M. Kim, "Moblurf: Motion deblurring neural radiance fields for blurry monocular video," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2025.

[20] Y. Wang, P. Chakravarthula, and B. Chen, "Dof-gs: Adjustable depth-of-field 3d gaussian splatting for refocusing, defocus rendering and blur removal," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2025.

[21] J. Lee, S. Cho, T. Kim, H.-D. Jang, M. Lee, G. Cha, D. Wee, D. Lee, and S. Lee, "Cocogaussian: Leveraging circle of confusion for gaussian splatting from defocused images," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2025.

[22] T. Li, L. Li, W. Wang, and Z. Feng, "Dehazing-nerf: neural radiance fields from hazy images," arXiv preprint arXiv:2304.11448, 2023.

[23] A. Ramazzina, M. Bijelic, S. Walz, A. Sanvito, D. Scheuble, and F. Heide, "Scatternerf: Seeing through fog with physically-based inverse neural rendering," in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023.

[24] Y. Li, J. Wu, L. Zhao, and P. Liu, "Derainnerf: 3d scene estimation with adhesive waterdrop removal," in IEEE International Conference on Robotics and Automation, 2024.

[25] S. Klenk, L. Koestler, D. Scaramuzza, and D. Cremers, "E-nerf: Neural radiance fields from a moving event camera," IEEE Robotics and Automation Letters, 2023.

[26] W. Yu, C. Feng, J. Tang, J. Yang, Z. Tang, X. Jia, Y. Yang, L. Yuan, and Y. Tian, "Evagaussians: Event stream assisted gaussian splatting from blurry images," arXiv preprint arXiv:2405.20224, 2024.

[27] Y. Weng, Z. Shen, R. Chen, Q. Wang, and J. Wang, "Eadeblur-gs: Event assisted 3d deblur reconstruction with gaussian splatting," arXiv preprint arXiv:2407.13520, 2024.

[28] S. Lee and G. H. Lee, "Diet-gs: Diffusion prior and event stream-assisted motion deblurring 3d gaussian splatting," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2025.

[29] L. Pan, C. Scheerlinck, X. Yu, R. Hartley, M. Liu, and Y. Dai, "Bringing a blurry frame alive at high frame-rate with an event camera," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019.

[30] P. Wang, L. Zhao, R. Ma, and P. Liu, "Bad-nerf: Bundle adjusted deblur neural radiance fields," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023.

참 고 문 헌

[31] R. Wu, Z. Zhang, M. Chen, X. Fan, Z. Yan, and W. Zuo, “Deblur4dgs: 4d gaussian splatting from blurry monocular video,” arXiv preprint arXiv:2412.06424, 2024.
[32] L. Zhao, P. Wang, and P. Liu, “Bad-gaussians: Bundle adjusted deblur gaussian splatting,” in Proceedings of the European Conference on Computer Vision, 2024.
[33] J. Lee, D. Kim, D. Lee, S. Cho, and S. Lee, “Crim-gs: Continuous rigid motion-aware gaussian splatting from motion blur images,” arXiv preprint arXiv:2407.03923, 2024.
[34] J. Lee, D. Kim, D. Lee, S. Cho, M. Lee, W. Lee, T. Kim, D. Wee, and S. Lee, “Comogaussian: Continuous motion-aware gaussian splatting from motion-blurred images,” arXiv preprint arXiv:2503.0533, 2025.
[35] L. Ma, X. Li, J. Liao, Q. Zhang, X. Wang, J. Wang, and P. V. Sander, “Deblur-nerf: Neural radiance fields from blurry images,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
[36] C. Peng and R. Chellappa, “Pdrf: Progressively deblurring radiance field for fast and robust scene reconstruction from blurry images,” arXiv preprint arXiv:2208.08049, 2022.
[37] D. Lee, M. Lee, C. Shin, and S. Lee, “Dp-nerf: Deblurred neural radiance field with physical scene priors,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023.
[38] C. Peng, Y. Tang, Y. Zhou, N. Wang, X. Liu, D. Li, and R. Chellappa, “Bags: Blur agnostic gaussian splatting through multi-scale kernel modeling,” in Proceedings of the European Conference on Computer Vision, 2024.
[39] B. Lee, H. Lee, X. Sun, U. Ali, and E. Park, “Deblurring 3d gaussian splatting,” in Proceedings of the European Conference on Computer Vision, 2024.
[40] H. Choi, H. Yang, J. Han, and S. Cho, “Exploiting deblurring networks for radiance fields,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2025.
[41] W.-T. Chen, W. Yifan, S.-Y. Kuo, and G. Wetzstein, “Dehazenerf: Multi-image haze removal and 3d shape reconstruction using neural radiance fields,” in International Conference on 3D Vision, 2024.
[42] J. Yu, Y. Wang, Z. Lu, J. Guo, Y. Li, H. Qin, and X. Zhang, “Dehazegs: Seeing through fog with 3d gaussian splatting,” arXiv preprint arXiv:2501.03659, 2025.
[43] C. Qian, Y. Guo, W. Li, and G. Markkula, “Weathergs: 3d scene reconstruction in adverse weather conditions via gaussian splatting,” IEEE International Conference on Robotics and Automation, 2025.
[44] Y. Yao, Z. Luo, S. Li, J. Zhang, Y. Ren, L. Zhou, T. Fang, and L. Quan, “Blendedmvs: A large-scale dataset for generalized multi-view stereo networks,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020.
[45] L. Liu, J. Gu, K. Zaw Lin, T.-S. Chua, and C. Theobalt, “Neural sparse voxel fields,” Advances in Neural Information Processing Systems, 2020.
[46] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, “Mip-nerf 360: Unbounded anti-aliased neural radiance fields,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.
[47] Q. Huynh-Thu and M. Ghanbari, “Scope of validity of psnr in image/video quality assessment,” Electronics letters, 2008.
[48] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” IEEE Transactions on Image Processing, 2004.
[49] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018.
[50] M. Chu, Y. Xie, J. Mayer, L. Leal-Taixé, and N. Thuerey, “Learning temporal coherence via self-supervision for gan-based video generation,” ACM Transactions on Graphics, 2020.
[51] K. He, J. Sun, and X. Tang, “Single image haze removal using dark channel prior,” IEEE Transactions on Pattern Analysis and Machine Intelligence, 2010.
[52] S. Sun and X. Guo, “Image enhancement using bright channel prior,” in International Conference on Industrial Informatics Computing Technology, Intelligent Technology, Industrial Information Integration, 2016.
[53] 피규어 출처: W. Kwon, J. Sung, M. Jeon, C. Eom, and J. Oh, “R3eVision: A Survey on Robust Rendering, Restoration, and Enhancement for 3D Low-Level Vision”, arXiv, 2025.

저 자 소 개

권 의 영

- 2025년 : 동아대학교 컴퓨터공학과 학사
- 2025년 ~ 현재 : 중앙대학교 첨단영상대학원 메타버스융합학과 석사과정
- 주관심분야 : 3D Vision, Low-Level Vision

오 지 형

- 2017년 : KAIST 전기공학부 학사
- 2019년 : KAIST 전기공학부 석사
- 2023년 : KAIST 전기공학부 박사
- 2023년 ~ 현재 : 중앙대학교 첨단영상대학원 조교수
- 2024년 ~ 현재 : 인쇼츠 (주) CAIO
- 주관심분야 : 3D Vision, Low-Level Vision, Generative AI