Textile & Leather Review

ISSN 2623-6281| www.tlr-journal.com | 10.31881/TLR

CVD-Friendly 3D Content Generation from Multi-View Images

Xinghong Hu, Zhenyu Xiao

How to cite: Hu X, Xiao Z. CVD-Friendly 3D Content Generation from Multi-View Images. Textile & Leather Review. 2026; 9:1809-1825. https://doi.org/10.31881/TLR.2026.1809

How to link: https://doi.org/10.31881/TLR.2026.1809

Published: 25 April 2026

CVD-Friendly 3D Content Generation from Multi-View Images

Xinghong Hu[1], Zhenyu Xiao[2*]

[1]Shenzhen Polytechnic University, Shenzhen 518055, Guangdong, China

[2]Tulingjishi (ShenZhen) Technology Co., Ltd., Shenzhen 518083, Guangdong, China

*zhenyuxiao@outlook.com

Article

https://doi.org/10.31881/TLR.2026.1809

Published 25 April 2026

ABSTRACT

Color vision deficiency (CVD) affects a significant portion of the population and poses persistent challenges to visual content perception, particularly in immersive 3D environments. While existing color enhancement methods primarily focus on 2D images, they lack explicit 3D scene representations and fail to ensure color consistency across viewpoints. This limitation reduces their applicability to 3D-aware tasks, such as the digital visualization of complex textile patterns with fine-grained color variations. In this paper, we present a fast, CVD-friendly 3D content generation framework that integrates perceptual color enhancement with efficient multi-view 3D reconstruction. Our approach first applies a conditional GAN to enhance color discriminability for CVD observers on a per-view basis. These enhanced views are then integrated across viewpoints and used as supervision to reconstruct a colorblind-friendly 3D scene using 3D Gaussian Splatting. A multi-view color consistency constraint is introduced to suppress view-dependent color drift. Importantly, the proposed method preserves the original Gaussian optimization pipeline while achieving perceptually enhanced rendering without compromising reconstruction efficiency. This framework ensures color discriminability and spatial consistency, making it suitable for accessibility-oriented applications in 3D textile display.

KEYWORDS

color vision deficiency, 3D gaussian splatting, multi-view 3D reconstruction, digital textiles

INTRODUCTION

Color vision deficiency (CVD) affects approximately 8% of males and 0.5% of females worldwide, causing difficulties in distinguishing certain color pairs and reducing accessibility of visual content. To mitigate this issue, a variety of color perception enhancement methods have been proposed to improve color discriminability for CVD observers. In particular, for the digital textile and fashion industries, ensuring that users with CVD can accurately perceive the color variations of fabric textures, yarn-level patterns, and woven structures is essential for an inclusive user experience. Early representative works formulate CVD compensation as an optimization problem that minimizes perceptual color differences between original images and their simulated CVD counterparts [1]. To better preserve visual naturalness, subsequent approaches introduce regularization constraints or personalized color mappings to control color modifications [2]. More recently, generative models have been explored to enable flexible and data-driven color enhancement tailored to individual CVD characteristics [3].

Despite their effectiveness, existing CVD-oriented color enhancement methods are predominantly designed for 2D images and videos. These approaches lack explicit 3D scene representations and therefore cannot guarantee color consistency across viewpoints. When applied independently to multi-view inputs, 2D enhancement often introduces view-dependent color variations, which become amplified in downstream 3D reconstruction and rendering tasks. As a result, such methods are not suitable for 3D-aware applications, such as novel view synthesis or virtual showroom displays for textiles, where consistent color perception across viewpoints is critical for evaluating fabric appearance.

In parallel, recent advances in neural scene representations have enabled high-quality 3D reconstruction from multi-view images. Neural Radiance Fields (NeRF) [4] provide a powerful implicit representation for view synthesis, while 3D Gaussian Splatting (3DGS) [5] achieves comparable visual quality with significantly improved rendering efficiency. These representations have been further extended for appearance manipulation and editing, including text- or reference-guided methods built upon NeRF [6–8] and editing-oriented extensions of 3DGS [9]. However, existing 3D content manipulation methods primarily focus on artistic or semantic editing and do not consider perceptual accessibility for color vision–deficient observers.

In this work, we address this gap by proposing a fast CVD-friendly 3D content generation framework that integrates perceptual color enhancement directly into the 3D reconstruction process. By combining view-wise CVD-oriented enhancement with multi-view–consistent Gaussian-based reconstruction, our method enables perceptually discriminable and view-consistent novel-view rendering for CVD users. At the same time, it preserves the original 3D Gaussian optimization pipeline and its high computational efficiency. Beyond general 3D scenes, our framework is particularly relevant to digital textile visualization scenarios, where subtle color differences and view-consistent appearance are crucial for fabric inspection and presentation.

RELATED WORK

Color Perception Enhancement for CVD

Color perception enhancement for individuals with color vision deficiency (CVD) has been predominantly studied in the context of 2D images. Traditional approaches formulate the problem as an optimization task that minimizes the perceptual color difference between the original image and its recolored counterpart simulation under a CVD model [10]. To better preserve visual naturalness, subsequent methods introduce regularization constraints that enforce consistency between the original and recolored images [11]. More recent work explores personalized color enhancement by decoupling and explicitly controlling color representations in the latent space of generative models [12].

Despite their effectiveness in improving color discriminability for static images, these methods operate exclusively in the 2D image domain and lack an explicit 3D scene representation. As a result, they cannot guarantee color consistency across viewpoints, nor can they be directly applied to 3D-aware tasks such as novel view synthesis, spatially consistent re-rendering, or downstream applications that require a unified 3D representation. This limitation prevents existing 2D CVD enhancement methods from supporting derived 3D applications, motivating the exploration of CVD-aware color perception enhancement within a 3D scene reconstruction framework.

3D Content Manipulation from 2D Frames

3D content manipulation from 2D frames has been extensively studied in computer vision and graphics, with existing approaches generally falling into two categories: explicit 3D reconstruction methods and implicit neural scene representations.

Early explicit modeling approaches reconstruct geometric surfaces from multi-view images and apply texture mapping to obtain colored 3D models. These models enable content manipulation by modifying texture maps or material parameters and are commonly used in digital textile modeling and fabric appearance editing [13,14].

In recent years, implicit neural representations, particularly Neural Radiance Fields (NeRF) [4], have emerged as a powerful alternative for reconstructing and rendering 3D scenes directly from 2D observations. Building upon NeRF, several works have explored appearance manipulation and re-rendering. PaletteNeRF [15] and RecolorNeRF [16] further explore color re-rendering by constraining the NeRF appearance using predefined color palettes or templates. Despite their effectiveness, these methods typically rely on fixed textual prompts, reference images, or color priors, limiting their applicability to perceptually adaptive manipulation tasks. Moreover, they are not designed to optimize perceptual discriminability under color vision deficiency, nor do they incorporate CVD simulation models into their optimization objectives. Adapting such methods to CVD-oriented tasks would require substantial reformulation of their color supervision and loss design. Therefore, in this work, we compare against a post-processing enhancement baseline that shares the same CVD-aware enhancement module but differs in whether multi-view consistency is enforced during 3D reconstruction.

Beyond volumetric NeRF-based representations, recent work has introduced explicit point-based neural primitives for efficient 3D scene modeling. 3D Gaussian Splatting (3DGS) [5] represents a scene as a set of anisotropic 3D Gaussians with learnable spatial, opacity, and color attributes, achieving high-quality novel view synthesis with real-time rendering performance. Subsequent studies have extended 3DGS to improve scalability and editability. Editing-oriented methods such as GaussianEditor [9] integrate semantic or instruction-based constraints into the Gaussian parameter space, enabling text-guided manipulation of 3D appearance. DreamGaussian [17] further combine diffusion models with Gaussian splatting to support text-driven 3D generation and re-rendering.

Overall, existing NeRF- and 3DGS-based 3D content manipulation methods lack mechanisms to selectively enhance color discriminability for viewers with varying types and severities of color vision deficiency (CVD). In contrast, our work aims to leverage 3D reconstruction while explicitly incorporating CVD-aware color perception models to recover lost chromatic contrast, enabling perceptually consistent and view-consistent color enhancement for CVD users from multi-view 2D observations.

METHODS

Given a set of multi-view RGB images $\boldsymbol{I} = \{I_j\}_{j=1}^{N}$ captured from calibrated cameras with known projection matrices $\{P_j\}$, our goal is to reconstruct a 3D scene representation whose rendered novel views are perceptually discriminable for colorblind observers, while maintaining geometric fidelity and high computational efficiency. To this end, we propose a two-stage framework that progressively integrates view-wise colorblind enhancement with efficient 3D reconstruction.

View-wise Colorblind Enhancement via Conditional GAN

In the first stage, we perform colorblind-oriented appearance enhancement independently for each input view. Specifically, a conditional generative adversarial network (cGAN) is employed to map each RGB image I_j to a colorblind-friendly image $\tilde{I}_j$. The generator G_ϑ is defined as:

$$\tilde{I}_j = G_\theta(I_j) \tag{1}$$

Unlike conventional image-to-image translation methods that aim to reconstruct realistic colors, our objective focuses on perceptual enhancement under color vision deficiency. Let $S(\cdot)$ denote a differentiable colorblind simulation operator, which is defined as:

$$S(I_{RGB}) = M \cdot I_{RGB} \tag{2}$$

where M is a personalized projection matrix from normal vision to CVD vision in RGB color space. We apply this operator to both the input and generated images to obtain

$$I_j^{cb} = S(I_j), \quad \tilde{I}_j^{cb} = S(\tilde{I}_j) \tag{3}$$

The superscript $(\cdot)^{cb}$ indicates that the image is represented in the simulated CVD color space.

To encourage enhanced color discriminability under CVD simulation, we define a color distinguishability loss L_{cd} based on structural similarity. For a local window centered at pixel location p, let $\mu_x(p)$ and $\mu_y(p)$ denote the mean intensity values of the input image and the simulated enhanced image, respectively, and let $\sigma_x(p)$ and $\sigma_y(p)$ denote the corresponding standard deviations. The cross-covariance between the two local windows is denoted by $\sigma_{xy}(p)$. Using these definitions, the color distinguishability loss is formulated as:

$$L_{cd}=-\frac{1}{N}\sum_{p}\left[\frac{\frac{2}{n-1}\sum_{i=1}^{n}\left(I_j\left(p_i\right)-\mu_{I_j}\right)\left(\tilde{I}_j^{cb}\left(p_i\right)-\mu_{\tilde{I}_j^{cb}}\right)}{\sigma^2_{I_j(p_i)}+\sigma^2_{\tilde{I}_j^{cb}(p_i)}}\right] \tag{4}$$

All local statistics are computed over a fixed-size window centered at p. This formulation explicitly encourages increased contrast and structural separability in the simulated CVD space while preserving local image structure.

Together with the standard adversarial loss L_{GAN} and the L1 loss, the overall optimization objective of Stage I is given by

$$L_{stageI} = L_{GAN} + \lambda_{cd}\, L_{cd} + \lambda_{L1}\, L_{L1} \tag{5}$$

This stage produces view-wise enhanced images that are perceptually optimized for colorblind observers while remaining structurally faithful to the original inputs. Paired data are obtained following [18].

Efficient 3D Reconstruction with Colorblind-aware Supervision

In the second stage, we reconstruct a colorblind-friendly 3D scene representation, termed ***CVDGS***, using the enhanced multi-view images $\{\tilde{I}_j\}$ as supervision. We adopt a 3DGS-based framework as the reconstruction backbone, which represents the scene as a set of 3D Gaussian primitives:

$$G = \{\mu_k, \Sigma_k, c_k, \alpha_k\}_{k=1}^{M} \tag{6}$$

where μ_k and Σ_k denote the mean and covariance of the k_{th} Gaussian, c_k represents its color attribute, and α_k denotes its opacity.

Given a camera view *j*, the CVDGS model renders an image $\hat{I}_j$ by splatting all visible Gaussians onto the image plane. The reconstruction objective minimizes the photometric error between the rendered image and the enhanced supervision:

$$L_{rec} = \Sigma_j \, ||\hat{I}_j - \tilde{I}_j||_1 \quad (7)$$

To achieve high training efficiency, CVDGS regulates the number and effectiveness of Gaussian primitives through multi-view consistent densification and pruning. Specifically, the contribution of each Gaussian is evaluated based on its aggregated reconstruction error across multiple views:

$$e_k = \Sigma_j \, w_{j,k} \cdot ||\hat{I}_j - \tilde{I}_j|| \quad (8)$$

where $w_{j,k}$ denotes the visibility weight of the *k*-th Gaussian in view *j*. Gaussian primitives are densified only when they consistently contribute to high reconstruction error across multiple views, while Gaussians with negligible multi-view contribution are pruned. This strategy avoids excessive growth of Gaussian primitives and significantly reduces optimization overhead without sacrificing rendering quality.

Although Stage I enhances each view independently, such unconstrained processing may introduce appearance inconsistency across viewpoints, which is detrimental to 3D reconstruction. Therefore, during the second stage, we explicitly enforce multi-view consistency among the enhanced images.

Using known camera parameters, we establish correspondences between pixels p_i and p_j from views *i* and *j* that project to the same 3D point. For these corresponding pixels, we impose a multi-view consistency loss that penalizes color discrepancies:

$$L_{mv} = \Sigma_{i,j} \, \Sigma_{p_i,p_j} \, ||\hat{I}_i(p_i) - \hat{I}_j(p_j)||_1 \quad (9)$$

By minimizing this loss, the enhanced appearance of the same physical surface is aligned across different viewpoints. This constraint effectively suppresses view-dependent color drift introduced by independent GAN inference and ensures that the enhanced images are coherent under viewpoint changes, which is critical for stable optimization of a shared 3D scene representation.

We note that, due to the view-wise nature of Stage I, it is possible that the same original color may be mapped to different enhanced colors in adjacent views. While such contradictions cannot be fully avoided at the image level, they are mitigated during Stage II through joint multi-view optimization. Specifically, enhanced color observations that occur more frequently across views provide stronger and more consistent supervision signals for the shared Gaussian color attributes. During optimization, these dominant color mappings are preserved, while sporadic or view-specific deviations are suppressed. As a result, the reconstructed 3D representation converges toward perceptually consistent colors that are statistically supported across multiple views, rather than reflecting isolated enhancement artifacts from individual frames.

EXPERIMENTS

We evaluate the proposed color vision deficiency (CVD)-aware reconstruction framework on selected scenes from Mip-NeRF 360 [19], which provides high-resolution video sequences with accurate camera poses.

Instead of using full scenes, we construct a targeted evaluation subset by extracting frames that exhibit color confusion areas under CVD simulation. Specifically, each video is first processed to identify frames with potential color discriminability difficulties using the Color Contrast Preservation Ratio (CCPR), which is computed between the simulated and original images. Frames with CCPR values below a predefined threshold (0.8) are considered challenging for color vision–deficient observers. If more than 60 consecutive frames exhibit CCPR values below the threshold, indicating persistent color confusion over time, the entire continuous frame segment is selected and included in the evaluation dataset.

For all seven selected sequences, we preserve the original camera parameters and image resolutions to ensure fair comparisons across methods. Color vision deficiency simulation follows the Machado model [20] for the most severe form of Deuteranopia, unless otherwise specified. While the proposed framework is applicable to other CVD types and severities by adjusting the projection matrix M a comprehensive evaluation across all variants is beyond the scope of this paper and is left for future work.

cGAN Architecture and Training

The generator follows a U-Net–based encoder–decoder architecture with skip connections to preserve spatial structure. The input is an RGB image, and the output is a distinguishability-enhanced RGB image of the same resolution. The discriminator adopts a PatchGAN design, which classifies local image patches as real or fake, encouraging high-frequency consistency without enforcing global color shifts. We set *c* and λ_{L1} = 100 in all experiments. The cGAN is trained using the Adam optimizer, and the learning rate is set to 2×10^{-4} and decayed linearly after half of the total training iterations. Training is performed for 100 epochs with a batch size of 1, using randomly sampled views from the training split. The generator weights are shared across all views and scenes.

Once trained, the cGAN serves as a reusable color enhancement module and can generate paired training data for Gaussian-based reconstruction across arbitrary scenes. For each input frame, the cGAN produces a corresponding distinguishability-enhanced image in approximately 0.02 seconds, enabling efficient preprocessing without introducing noticeable overhead.

CVDGS Training

The 3D Gaussian parameters, including position, covariance, opacity, and color, are optimized using Adam with learning rates of 1.6×10^{-4} for position, 1.0×10^{-3} for covariance, 2.5×10^{-3} for color, and 5.0×10^{-2} for opacity. An exponential decay schedule is applied to the position learning rate. Gaussian densification is performed every 100 iterations during the early training phase, where new Gaussians are spawned in regions with high multi-view reprojection error, and low-contributing Gaussians are removed based on opacity and screen-space size thresholds. Gaussians with opacity lower than 0.005 are pruned, and those whose projected size exceeds a predefined threshold are split to better capture fine structures. The densification process is terminated after 3,000 iterations, after which only optimization continues. The total number of training iterations is 5,000, and the reconstruction converges significantly faster than standard 3D Gaussian Splatting methods. All experiments are conducted on a single NVIDIA RTX 3090 GPU, with the full reconstruction process completing within approximately 100–110 seconds per scene. Importantly, the enhanced images are treated as standard RGB inputs, and all geometric optimization, densification, and pruning steps remain unchanged. This design ensures that performance improvements stem from perceptual color enhancement rather than modifications to the 3D reconstruction pipeline itself.

RESULTS

Quantitative Results

We compare the proposed method with a post-processing baseline: CVDGS with cGAN-based image enhancement (CVDGS + cGAN Post). All methods are evaluated on the CCPR-filtered Mip-NeRF 360 subset under both the normal RGB space and the simulated CVD space. Color discriminability is evaluated using SSIM and CCPR computed between original frames and recolored frames.

To evaluate multi-view color consistency in 3D space, we introduce the Relative Warping Error (RWE), which measures whether the same 3D point is rendered with consistent color across different viewpoints.

Specifically, RWE is defined as the ratio between the appearance variation of corresponding pixels rendered from the same 3D point under different views and the appearance variation observed in the original input images:

$$RWE = Average\left(\Sigma_{i,j} \Sigma_{p_i,p_j} \frac{\|\hat{I}_i(p_i) - \hat{I}_j(p_j)\|_1}{\|I_i(p_i) - I_j(p_j)\|_1}\right) \tag{10}$$

Table 1 shows that both enhancement-based methods substantially improve CCPR and SSIM compared to the input, demonstrating the effectiveness of cGAN-based color enhancement. Specifically, CVDGS + cGAN Post achieves the highest CCPR (0.842) and SSIM (0.760), benefiting from aggressive view-wise color contrast enhancement.

In contrast, our method attains comparable CCPR (0.839) and SSIM (0.758) while significantly outperforming the post-processing baseline in terms of RWE. The RWE of our method is reduced from 0.087 to 0.033, indicating that the same 3D points are rendered with substantially more consistent colors across viewpoints. This result highlights an important trade-off between image-level color enhancement and multi-view color consistency. While post-processing enhancement can slightly increase CCPR on individual views, it introduces view-dependent color variations that lead to higher warping error in 3D rendering. By integrating color enhancement into a multi-view–consistent reconstruction framework, our method achieves a more balanced improvement, preserving both perceptual color discriminability and stable 3D appearance across viewpoints.

Table 1. Quantitative comparison between state-of-the-art methods and our method

	Input	CVDGS + cGAN Post	Ours
CCPR	0.713	0.842	0.839
SSIM	0.625	0.760	0.758
RWE	-	0.087	0.033

Qualitative Results

Figure 1 shows qualitative comparisons of novel-view renderings under simulated Deuteranopia. CVDGS + cGAN Post suffers from noticeable color flickering and inconsistency across viewpoints, particularly in regions labeled with the red rectangle. These artifacts stem from independent view-wise enhancement and are amplified during novel-view synthesis.

In contrast, our method produces color-consistent and perceptually distinguishable renderings across all viewpoints. Regions that are ambiguous in the post-processing baselines become clearly separable, while structural details and fine geometry remain intact. Importantly, the enhanced colors are coherently integrated into the 3D representation, avoiding view-dependent artifacts.

The comparison highlights a clear distinction between post-processing enhancement and the proposed integrated approach. Although both baselines employ the same cGAN architecture, applying color enhancement after reconstruction fails to enforce multi-view consistency, resulting in unstable color appearance in 3D renderings.

By contrast, our method incorporates color enhancement prior to and during 3D optimization, allowing the Gaussian representation to jointly encode geometry and CVD-aware color information. This design leads to consistent color contrast improvement without compromising reconstruction quality or runtime efficiency.

Overall, the results demonstrate that multi-view–aware integration of perceptual color enhancement is crucial for achieving stable and effective CVD-friendly 3D reconstruction, especially in scenes with persistent color ambiguity.

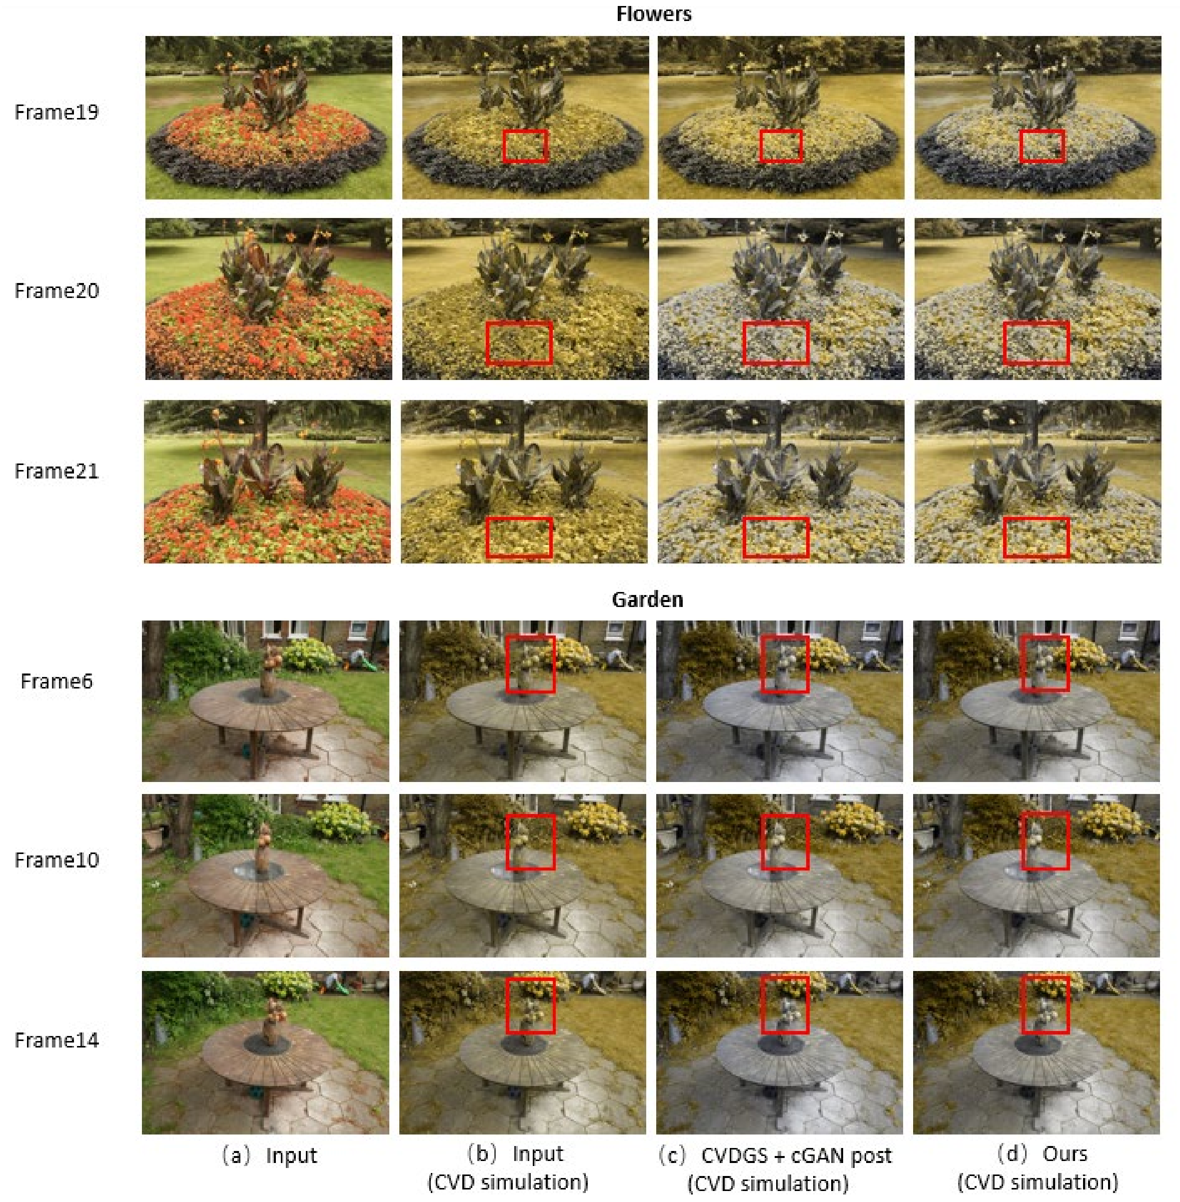

Figure 1. Qualitative comparisons between our method and post-processing approach

To verify that the proposed framework does not introduce unnecessary color changes under normal vision, we conduct an additional evaluation on all remaining frames that are not selected by the CCPR-based filtering.

On these frames, we compare the enhanced images with the original inputs using SSIM and PSNR computed in the standard RGB space.

Across all scenes, the enhanced images achieve an average SSIM of 0.962 and an average PSNR of 31.8 dB with respect to the original images. These high similarity scores indicate that, for scenes where color vision–deficient observers do not experience significant perceptual difficulty, the proposed method preserves natural appearance and avoids over-recoloring.

LIMITATION

The proposed multi-view consistency loss primarily targets surfaces with approximately view-independent (diffuse) appearance, which is a common assumption in Gaussian-based scene representations. For materials exhibiting strong specular reflections or transparency, color appearance may vary inherently across viewpoints. In such cases, enforcing strict color consistency may be ill-posed.

We note that our evaluation relies on the Machado model for simulating color vision deficiency. While this model is physiologically motivated and widely adopted, real-world CVD perception varies significantly across individuals. As a result, simulated evaluation should be interpreted as an approximation rather than a definitive measure of subjective perception. Incorporating user-specific perceptual feedback or alternative simulation models is an important direction for future work.

CONCLUSION

In this paper, we present a color vision deficiency (CVD)-aware 3D content generation framework that integrates perceptual color enhancement with fast multi-view Gaussian-based reconstruction. By focusing on temporally stable regions with persistent color confusion and incorporating a cGAN-based enhancement strategy in a multi-view–consistent manner, our approach enables color-discriminable novel-view rendering for color vision–deficient observers without compromising reconstruction quality or efficiency.

Through quantitative experiments, we demonstrate that simply applying color enhancement as a post-processing step fails to ensure consistent perceptual improvement across viewpoints. In contrast, our method achieves stable and significant gains in color contrast preservation under simulated Deuteranopia while maintaining comparable geometric fidelity and runtime performance to CVDGS. The improved consistency and discriminability validated in our results suggest that this framework can be effectively extended to the digital presentation of complex textile scenes, providing a reliable tool for colorblind accessible 3D fabric visualization.

The results highlight the importance of integrating perceptual objectives into the reconstruction pipeline rather than treating them as independent image-level operations. We believe this work takes a meaningful step toward perception-aware 3D scene reconstruction and opens new possibilities for accessibility-oriented rendering and visualization in immersive environments.

Author Contributions

Hu X. designed, collected and analyzed the data, and drafted the manuscript. Xiao Z. conducted the study, critically revised the manuscript for important intellectual content, and gave final approval of the version to be published. Hu X. and Xiao Z. participated fully in the work, take public responsibility for appropriate portions of the content, and agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

Conflicts of Interest

The authors declare no conflict of interest.

Funding

This research was supported by the Shenzhen Science and Technology Program (Grant No. RCBS20221008093249079), the Shenzhen Polytechnic University-level Scientific Research Project (No. 6022310013K), and the Shenzhen Polytechnic High-level Talent Research Start-up Project (No. 6022310050k).

REFERENCES

[1] Hu XH, Liu XT, Zhang ZM, Xia MH, Li CZ, Wong TT. Colorblind-shareable videos by synthesizing temporal-coherent polynomial coefficients. ACM Transactions on Graphics. 2019; 38(6):1–12. doi: 10.1145/3355089.335653

[2] Wang Y, Li D, Hu M, Cai L. Non-local recoloring algorithm for color vision deficiencies with naturalness and detail preserving. In: Zhai G, Zhou J, Yang H, An P, Yang X, editor. Communications in Computer and Information Science. Proceedings of the International Forum on Digital TV and Wireless Multimedia Communications; 19-20 September 2019; Shanghai, China. Singapore, Springer; 2019. p. 23–34. doi: 10.1007/978-981-15-3341-9_3

[3] Jiang S, Liu D, Li D, Xu C. Personalized image generation for color vision deficiency population. Proceedings of the IEEE/CVF International Conference on Computer Vision; 1-6 October 2023; Paris, France. New York, NY, USA. IEEE; 2023. p. 22514-22523. doi: 10.1109/ICCV51070.2023.02063

[4] Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R. NeRF: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM. 2021; 65(1):99–106. doi: 10.1145/3503250

[5] Kerbl B, Kopanas G, Leimkühler T, Drettakis G. 3D Gaussian Splatting for real-time radiance field rendering. ACM Transactions on Graphics. 2023; 42(4):1–14. doi: 10.1145/3592433

[6] Chen Y, Yuan Q, Li Z, Liu Y, Wang W, Xie C, et al. UPST-NeRF: Universal photorealistic style transfer of neural radiance fields for 3D scenes. IEEE Transactions on Visualization and Computer Graphics. 2025; 31(4):2045-2057. doi: 10.1109/TVCG.2024.3378692

[7] Fan Z, Jiang Y, Wang P, Gong X, Xu D, Wang Z. Unified implicit neural stylization. In: Avidan S, Brostow G, Cissé M, Farinella GM, Hassner T, editor. Lecture Notes in Computer Science. Proceedings of the European Conference on Computer Vision; 23-27 October 2022; Tel Aviv, Israel. Cham, Switzerland, Springer; 2022. p. 636–654. doi: 10.1007/978-3-031-19784-0_37

[8] Zhang K, Kolkin N, Bi S, Luan F, Xu Z, Shechtman E, et al. ARF: Artistic radiance fields. In: Avidan S, Brostow G, Cissé M, Farinella GM, Hassner T, editor. Lecture Notes in Computer Science. Proceedings of the European Conference on Computer Vision; 23-27 October 2022; Tel Aviv, Israel. Cham, Switzerland, Springer; 2022. p. 717–733. doi: 10.1007/978-3-031-19821-2_41

[9] Chen Y, Liu H, et al. GaussianEditor: Editing 3D Gaussians with semantic control. ArXiv:2311.16037. https://doi.org/10.48550/arXiv.2311.16037.

[10] Huang X, Zhang Z, Li X, Wong TT. Deep visual sharing with colorblind. IEEE Transactions on Computational Imaging. 2019; 5(4):649–659. doi: 10.1109/TCI.2019.2908291

[11] Zhu Z, Toyoura M, Go K, Kashiwagi K, Fujishiro I, Wong TT, et al. Personalized image recoloring for color vision deficiency compensation. IEEE Transactions on Multimedia. 2021; 24:1721–1734. doi: 10.1109/TMM.2021.3070108

[12] Rigos A, Chatzistamatis S, Tsekouras GE. A systematic methodology to modify color images for dichromatic human color vision and its application in art paintings. International Journal of Advanced Trends in Computer Science and Engineering. 2020; 9(4):5015–5025. doi: 10.30534/ijatcse/2020/119942020

[13] Seitz SM, Curless B, Diebel J, Scharstein D, Szeliski R. A comparison and evaluation of multi-view stereo reconstruction algorithms. Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06); 17-22 June 2006; New York, NY, USA. New York, NY, USA, IEEE; 2006. p.807–820. doi: 10.1109/CVPR.2006.19

[14] Furukawa Y, Ponce J. Accurate, dense, and robust multiview stereopsis. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2010; 32(8):1362–1376. doi: 10.1109/TPAMI.2009.161

[15] Kuang Z, Luan F, Bi S, Shu Z, Wetzstein G, Sunkavalli K. PaletteNeRF: Palette-based appearance editing of neural radiance fields. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 17-24 June 2023; Vancouver, BC, Canada. New York, NY, USA, IEEE; 2023. p. 20691-20700. doi: 10.1109/CVPR52729.2023.01982

[16] Gong B, Wang Y, Han X, Dou Q. RecolorNeRF: Layer-decomposed radiance fields for efficient color editing of 3D scenes. Proceedings of the ACM International Conference on Multimedia; 29 October-3 November 2023; Ottawa, ON, Canada; New York, NY, USA, Association for Computing Machinery; 2023. p. 8004–8015. doi: 10.1145/3581783.3611957

[17] Tang J, Ren J, Zhou H, Liu Z, Zeng G. DreamGaussian: Generative Gaussian splatting for efficient 3D content creation. arXiv preprint. 2023; arXiv:2309.16653. doi: 10.48550/arXiv.2309.16653

[18] Barron JT, Mildenhall B, Verbin D, Srinivasan PP, Hedman P. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 18-24 June 2022; New Orleans, LA, USA; New York, NY, USA, IEEE; 2022. p. 5470–5479. doi: 10.1109/CVPR52688.2022.00539

[19] Lu C, Xu L, Jia J. Contrast preserving decolorization with perception-based quality metrics. International Journal of Computer Vision. 2014; 110(2):222–239. doi: 10.1007/s11263-014-0732-6

[20] Machado GM, Oliveira MM, Fernandes LAF. A physiologically-based model for simulation of color vision deficiency. IEEE Transactions on Visualization and Computer Graphics. 2009; 15(6):1291–1298. doi: 10.1109/TVCG.2009.113