Lighting-Aware Radiance Field Inpainting

(Masters 2 internship)

George Drettakis, GRAPHDECO, Inria Sophia Antipolis (France)

http://team.inria.fr/graphdeco
George.Drettakis@inria.fr

http://www-sop.inria.fr/members/George.Drettakis/

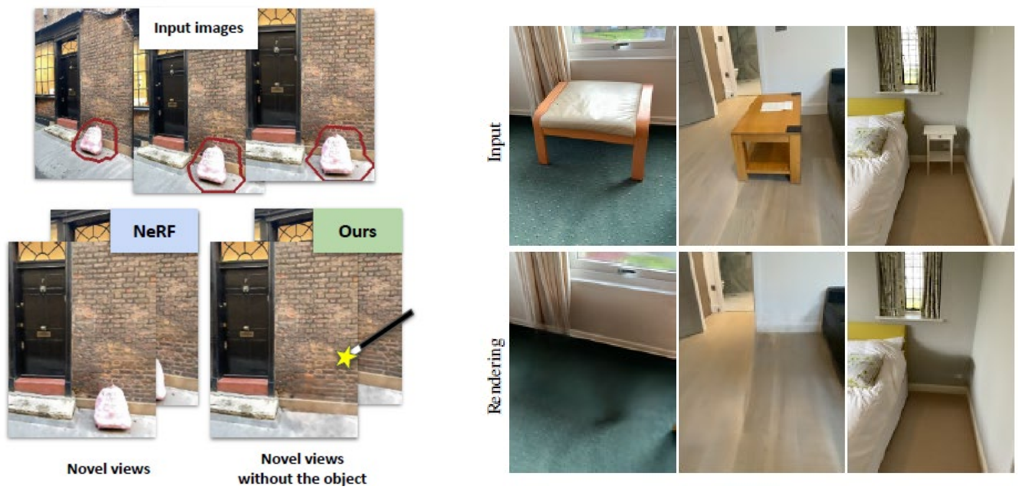

(a) *(b)*

Figure 1: Images from [Weder23]: (a) multiple input images are used to create a NeRF and then an object can be removed; the content behind the object is inpainted. (b) Despite successful inpainting overall, shadows from the removed objects can remain -- this is particularly visible in the leftmost scene.

Context and goal

Neural Radiance Fields (NeRFs) [Xie22] and Radiance Fields (RF) in general have received widespread interest in recent years, and are now being applied in many different application domains from e-commerce and building inspection to visual effects and games. Editing such representations remains complicated; one of the main problems is that when we remove an object from the scene, we need to *inpaint* the remaining part of the RF so it remains realistic (see Fig. 1 (a) above). Initial attempts to solve this problem [Weder23,Mirzaei23] achieve convincing results albeit with some visual artifacts. Such artifacts are often due to the change in lighting that occurs after removing an objects, e.g., “left over” shadows in the inpainted result (Fig. 1 (b)).

In this internship we will follow a similar strategy as [Weder23,Mirzaei23] exploiting powerful 2D inpainting solutions and imposing 3D consistency. The novelty of the approach we propose will reside in taking into account lighting and shadows due to objects being removed, and correctly handling the corresponding changes in lighting.

Approach

In this internship we will work with our recent 3D Gaussian Splatting (3DGS) framework [Kerbl23], that allows us to reason more intuitively about 3D compared to traditional NeRFs. We will start by adapting a method similar to [Weder23,Mirzaei23] to 3DGS. We will first run 3DGS on the unaltered scene, and then use one of the modern segmentation tools (eg [Kirillov23]) to segment out a given object, developing a multi-view refinement strategy to propagate the mask across views. We expect this to work well, but specific regularizations may be required to get the desired quality of this step.

Once a reliable multi-view mask is available, we will apply one of the many state of the art single view inpainting methods to each of the input views, similar to [Weder23,Mirzaei23], and develop a new approach to optimizing a multi-view consistent RF. To take into account shadowing and other lighting effects implied from removing the selected object, we will estimate a coarse version of the lighting environment (eg using radiance collection methods similar to [Philip21] or using a simplified version of [Song19]) and use physically-based rendering methods to identify regions in shadow that change with scene editing. It is possible that we will require strong priors on shape (e.g., planar surface priors in interior scenes) to perform high-quality inpainting of geometry. If so, we will investigate strategies based either on uncertainty estimation of neural fields [Goli23] and/or using large-scale ground truth scans [Yeshwanth23]. We will evaluate our results on the readily available datasets from [Weder23,Mirzaei23] that contain real scenes with and without objects.

Work environment and requirement

The internship will take place at Inria Sophia Antipolis in the GRAPHDECO group (http://team.inria.fr/graphdeco). Inria will provide a monthly stipend of around 1100 euros for EU citizens in their final year of masters, and 400 euros for other candidates.

Candidates should be passionate about computer graphics and neural rendering methods, and preferably have strong programming and mathematical skills as well as knowledge in computer graphics, geometry processing and machine learning, with experience in C++, real-time rendering techniques, path-tracing (knowledge of mitsuba1/2/3 is a plus), OpenGL and GLSL on the graphics side. Some knowledge of machine learning is a plus.

References

[Goli23] Goli, Lily, Cody Reading, Silvia Selllán, Alec Jacobson, and Andrea Tagliasacchi. "Bayes' Rays: Uncertainty Quantification for Neural Radiance Fields." *arXiv preprint arXiv:2309.03185* (2023).

[Kerbl23] Kerbl, Bernhard, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. "3d gaussian splatting for real-time radiance field rendering." *ACM Transactions on Graphics (ToG)* 42, no. 4 (2023): 1-14. https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/

[Kirillov23] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y. and Dollár, P., 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.

[Mirzaei23] Mirzaei, A., Aumentado-Armstrong, T., Derpanis, K. G., Kelly, J., Brubaker, M. A., Gilitschenski, I., & Levinshtein, A. (2023). SPIn-NeRF: Multiview segmentation and perceptual inpainting with neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 20669-20679) https://openaccess.thecvf.com/content/CVPR2023/papers/Mirzaei_SPIn-NeRF_Multiview_Segmentation_and_Perceptual_Inpainting_With_Neural_Radiance_Fields_CVPR_2023_paper.pdf

[Song19] Song, S., & Funkhouser, T. (2019). Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 6918-6926).

[Weder23] Weder, S., Garcia-Hernando, G., Monszpart, A., Pollefeys, M., Brostow, G.J., Firman, M. and Vicente, S., 2023. Removing objects from neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 16528-16538). https://openaccess.thecvf.com/content/CVPR2023/papers/Weder_Removing_Objects_From_Neural_Radiance_Fields_CVPR_2023_paper.pdf

[Xie22] Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V. and Sridhar, S., 2022, May. Neural fields in visual computing and beyond. In Computer Graphics Forum (Vol. 41, No. 2, pp. 641-676).

[Yeshwanth23] Yeshwanth, Chandan, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. "ScanNet++: A High-Fidelity Dataset of 3D Indoor Scenes." In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12-22. 2023.