Per-View Neural Textures for Image-Based Rendering

Masters 2 Internship (4-6 months)

George Drettakis, Inria Sophia Antipolis (France)
George.Drettakis@inria.fr

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

Context and goal

Traditional rendering in computer graphics consists of taking a geometric model with textures, materials and lights and performing lighting simulation (or some approximation thereof); an alternative with very different tradeoffs is Image-Based Rendering (IBR) [Buehler 01, Chaurasia 13], which allows free-viewpoint 3D navigation of a real scene, captured with a set of photographs ("multi-view dataset"). We first estimate a coarse 3D model and camera poses using computer vision techniques; IBR then reprojects the input photos into novel views, allowing free-viewpoint navigation. We have developed novel algorithms that greatly improve the quality of IBR [Hedman 16, Hedman 18]. In the most recent solution [Hedman 18], we use a deep learning approach to significantly improve image quality.

A different approach is Deferred Neural Rendering (DNR) that learns novel-view synthesis in an end-to-end manner [Thies 19]. In particular, this approach learns a feature vector in each texel of a UV-map built over a reconstructed object, allowing high-quality novel view renderings including non-diffuse effects and in cases with imprecise geometry. However, a very large number of images are required to learn these feature vectors to allow good quality rendering. It is also unclear to which extent the texture atlas must be continuous to allow the method to work in general scenes, and how robust the method is to occlusion boundaries. In this internship, we will evaluate these issues, and propose a new solution using additional data, in the spirit of [Hedman 16, 18].

Approach

We will start by evaluating the dependence of DNR on the number of input images and the type of materials in the scene, and then move on to the dependence of occlusion boundary complexity. All of these tests will be performed using synthetic scenes, to allow flexibility and quantitative evaluation of the results. The internship will include a significant system-building effort (e.g disk caching, sample cropping etc.), in particular to handle the large data requirements for training. We will focus on these three aspects, i.e., image density, material properties and occlusion boundaries.

We will then focus on developing a view-dependent UV-map approach to learning features. We will use the per-view meshes and UV-maps of [Hedman 18], and learn equivalent features to those in DNR for these meshes. Rendering will then be performed in a hybrid manner, by first estimating a novel view depth map using the per-view meshes, then using a network-based renderer similar to DNR. We expect this strategy to significantly diminish the number of images required for training, and to allow treatment of scenes with complex occlusion boundary configurations.

Depending the advancement, we will also investigate the treatment of non-diffuse materials. Treating this problem presents different challenges, since the density and placement of input cameras is critical in this context. We will investigate theoretical arguments using Fourier analysis (in the spirit of [Milednhall 19]), coverage and reconstruction of the reflected radiance field, to determine the required density and placement of cameras to correctly learn to reconstruct non-diffuse reflective behavior.

Work environment and requirements

The internship will take place at Inria Sophia Antipolis, in the beautiful French Riviera. Inria will provide a monthly stipend between 450 and 1100€ depending on the situation of the candidate. The intern will work closely with the Ph.D. students in the group.

Candidates should have strong programming and mathematical skills as well as knowledge in computer graphics (a 4th year or higher graphics course is desirable), computer vision, geometry processing and machine learning. Successful Masters internships may lead to a Ph.D. in the context of the ERC FUNGRAPH project (http://fungraph.inria.fr)

References

[Buehler 01] C. BUEHLER, M. BOSSE, L. MCMILLAN, S. GORTLER, and M. COHEN. "Unstructured lumigraph rendering." In *Proceedings SIGGRAPH 2001*, pp. 425-432. ACM, 2001. https://www.ics.uci.edu/~gopi/ICS280Win02/UnstructuredLumigraph.pdf

[Chaurasia 13] G. CHAURASIA, S. DUCHENE, O. SORKINE-HORNUNG, & G. DRETTAKIS. (2013). Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, *32*(3), 30.

http://www-sop.inria.fr/reves/Basilic/2013/CDSD13/

[Hedman 16] P. HEDMAN, T. RITSCHEL, G. DRETTAKIS, G. BROSTOW, Scalable Inside-Out Image-Based Rendering, ACM Transactions on Graphics 35 (SIGGRAPH Asia), 6, December 2016 http://www-sop.inria.fr/reves/Basilic/2016/HRDB16/

[Hedman 18] P. HEDMAN, J. PHILIP, T. PRICE, J.-M. FRAHM, G. DRETTAKIS, G. BROSTOW, Deep Blending for Image-Based Rendering, ACM Transactions on Graphics 37 (SIGGRAPH Asia), 6, Dec. 2018 http://www-sop.inria.fr/reves/Basilic/2018/HPPFDB18/

https://repo-sam.inria.fr/fungraph/deep-blending/

[Milednhall 19] B. Mildenhall, P. P. Srinivasan, R. Ortiz-Cayon, N. K. Kalantari, R. Ramamoorthi, R. Ng, and A. Kar, "Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines," ACM Trans. Graph., vol. 38, no. 4, pp. 29:1--29:14, July 2019.

[Ortiz-Cayon 15] ORTIZ-CAYON, Rodrigo, DJELOUAH, Abdelaziz, et DRETTAKIS, George. A bayesian approach for selective image-based rendering using superpixels. In : *International Conference on 3D Vision-3DV*. 2015.

http://www-sop.inria.fr/reves/Basilic/2015/ODD15/

[Penner 17] E. PENNER and L. ZHANG. 2017. Soft 3D reconstruction for view synthesis. ACM Transactions on Graphics (TOG) 36, 6 (2017), 235

https://ericpenner.github.io/soft3d/

[Thies 19] Thies, Justus and Zollhofer, Michael and Niessner, Matthias, Deferred Neural Rendering: Image Synthesis using Neural Textures, ACM Transactions on Graphics 2019 (TOG), 2019