Hybrid Rendering

Masters 2 Internship (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"). The pose of the cameras is estimated with Structure from Motion, and a coarse 3D model computed with Multi-view stereo (MVS); IBR reprojects the input photos into novel views, allowing free-viewpoint navigation.

In recent work, we have developed novel algorithms that greatly improve the quality of IBR [Hedman 16, Hedman 18], compared to baseline solutions such as Unstructured Lumigraph (ULR) [Buehler 01] or Bayesian Rendering [Ortiz-Cayon 15]. In the most recent solution [Hedman 18], we use a deep learning approach to significantly improve image quality (see Figure 1).

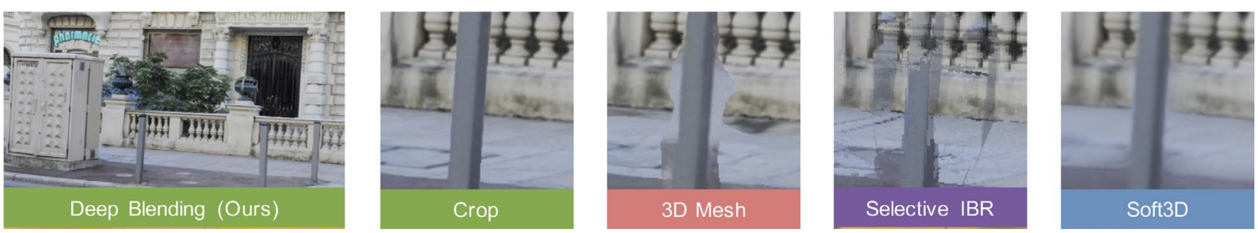

Figure 1: Left & first Crop: Our most recent result [Hedman 18], 3D Mesh: Textured Rendering; Selective IBR [Ortiz-Cayon 15], Soft3D [Penner17]

Each solution has different speed/quality tradeoffs, typically involving increasing computational cost as the quality improves. However, the quality of textured MVS meshes is constantly progressing, both in geometry and the extracted texture [Waechter 14]. For many regions in any given novel view, the textured mesh is of sufficient quality; in other parts ULR is sufficient, and only regions with low quality MVS reconstruction or complex materials really need more sophisticated solutions such as [Hedman 18]. Our goal is to develop a hybrid rendering algorithm, that will optimize for speed and quality by choosing the cheapest possible algorithmic solution for a given region of the novel view.

Approach

To propose this novel algorithm, we need to address three main challenges. 1) Identifying which algorithm provides a given quality for a given image region (and possibly which parameters to use), 2) designing and implementing a common data structure to accommodate the different representations used by each algorithm that will facilitate online algorithm selection during rendering and 3) designing and implementing the hybrid rendering algorithm.

Finding the cheapest algorithm with sufficient quality. We will start with a Bayesian formulation, inspired by [Ortiz-Cayon 15], using an image-region representation of the input images (e.g., superpixels), and a leave-one-out error metric (but possibly using Deep Features as an error metric) to choose between textured mesh, ULR [Buehler 01] and [Hedman16, 18] (possibly just the heuristic blend). In addition to choosing which algorithm is sufficient for a given region, we will investigate ways of estimating the best parameters to use (e.g., number of images to blend, blending strategy etc.). An optional second approach will be to use deep learning to choose the algorithm and the parameters, inspired by [Hedman 18].

Common Spatio-directional representation. Each algorithm (i.e., standard mesh rendering, [Buelher 01] and [Hedman 01]) use different representations: i.e., textured global mesh, the global mesh with calibrated cameras [Buehler 01] and per-view meshes [Hedman 16, 18]. In addition, we need to map image-regions used for identification to these meshes. We will develop a new data structure, possibly based on the voxel approach of [Hedman 16, 18], mapping the result of the quality estimation to the different meshes in a unified manner.

Hybrid Rendering Algorithm. The final part of this internship will be to develop an efficient GPU-based rendering algorithm that maximizes quality while minimizing rendering cost, by rendering parts of the novel view with each different algorithm. Additional blending passes will be required to avoid visible artifacts due to the differences between algorithms.

The intern will develop this research in the existing C++/OpenGL software framework of the group that includes existing implementations of [Buehler 01, Ortiz-Cayon 15, Hedman 16 and 18].

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. Successful Masters internships may lead to a Ph.D. in the context of the ERC FUNGRAPH project (http://fungraph.inria.fr)

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.

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/

[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/

[Waechter 14] M. WAECHTER, N. MOEHRLE, et M. GOESELE. Let there be color! Large-scale texturing of 3D reconstructions. In : *European Conference on Computer Vision*. Springer, Cham, 2014. p. 836-850.

https://www.gcc.tu-darmstadt.de/media/gcc/papers/Waechter-2014-LTB.pdf