

Procedural Generation of Materials for Deep Learning

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

Computer Graphics images are created from scenes composed of geometry, materials and lights. Creating these assets is time-consuming and expensive; material design is one of the most difficult components of this process. Material acquisition is a very interesting alternative to manual material design and varies from the usage of complex acquisition hardware to more recent methods that attempt to capture a material using a single photo [Aittala 13, Li 17, Deschaintre 18 & 19]. The most recent such methods [Li 17, 18] including our own [Deschaintre 18 & 19] use deep learning to infer the material properties, i.e., parameters of the Spatially Varying Bidirectional Reflectance Distribution Function (SVBRDF).

An important challenge in this process is finding the training data used for deep learning. In our previous work [Deschaintre 18 & 19] we used artist-created procedural materials, that are very realistic, but often have limitations in their usage license, and are limited to specific materials that the artists have chosen to create.

Our goal is to overcome this difficulty by designing an efficient and effective approach to generating procedural materials in a way that achieves the same level of success (or even better) than the previous methods using artist-created materials.

Approach

The key element of deep learning methods for material capture is that the network should learn the connection between SVBRDF parameters and appearance. For that, we do not necessarily need materials designed by artists, but simply many relevant examples of plausible materials.

In this internship, we will develop a procedural generation system creating sets of normal, diffuse, roughness and specular maps (i.e., the parameters of the SVBRDF) based on different rules. Such rules could include, continuity for normals, options for correlation or not between maps, number of different materials per map, etc.. Such maps will not necessarily be “realistic looking”. Once put together, it is straightforward to train the network of [Deschaintre 18].

The first goal will be to answer the question on whether this new data generation system yields similar results to hand-created materials while providing a lot of additional flexibility, i.e., decorrelation of parameter maps, adding some random anisotropy in the generation, or depthmap or displacement maps, not relying on propriety, easy to add new rules.

If the initial results are equivalent to hand-crafted materials, we will investigate the "smallest required set" of material variations to achieve the desired result, which will allow much faster training. We will also investigate the effect of different parameters in the generation process and attempt to improve the results for cases the previous method could not handle well.

If on the other hand the results are not equivalent, we will attempt to determine which aspects of hand-crafted materials are important, by progressively mixing procedural and hand-crafted exemplars for training, and systematically investigating different aspects of the two approaches for training (e.g., material classes, different mixtures etc).

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

- [Aittala 13] Aittala, Miika, Tim Weyrich, and Jaakko Lehtinen. "Practical SVBRDF capture in the frequency domain." *ACM Trans. Graph.* 32.4 (2013): 110-1.
- [Li 17] Li, Xiao, Yue Dong, Pieter Peers, and Xin Tong. "Modeling surface appearance from a single photograph using self-augmented convolutional neural networks." *ACM Transactions on Graphics (TOG)* 36, no. 4 (2017): 45.
- [Li 18] Li, Z., Sunkavalli, K. and Chandraker, M., 2018. Materials for masses: SVBRDF acquisition with a single mobile phone image. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 72-87).
- [Deschaintre 18] Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, Adrien Bousseau, Single-Image SVBRDF Capture with a Rendering-Aware Deep Network, *ACM Transactions on Graphics (SIGGRAPH Conference Proceedings)*, Volume 37, Number 128, pages 15, Aug 2018 <http://www-sop.inria.fr/reves/Basilic/2018/DADDB18/>
- [Deschaintre 19] Valentin Deschaintre, Miika Aittala, Frédo Durand, George Drettakis, Adrien Bousseau, Flexible SVBRDF Capture with a Multi-Image Deep Network, *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)*, Volume 38, Number 4, July 2019 <http://www-sop.inria.fr/reves/Basilic/2019/DADDB19/>