Consistent Neural Rendering of Large Scale Landscapes

Guillaume Cordonnier, George Drettakis, Inria Sophia Antipolis (France)
guillaume.cordonnier@inf.ethz.ch, George.Drettakis@inria.fr
https://people.inf.ethz.ch/~gcordonn/
http://www-sop.inria.fr/members/George.Drettakis/

Context and goal

Landscapes, in particular high mountain scenes, are a common background in many computer graphics applications. Various methods aim to synthesize high-level descriptions of such landscapes as a set of maps (e.g., elevation, presence of rocks, the density of vegetation) [1]. Exploring these landscapes from a first-person viewpoint requires the addition of geometrical elements (trees, rocks) and detailed textures. Typical methods for detail enhancement usually focus on a specific detail type [2] and hence do not ensure consistency between the various elements. A final rendering step is often subcontracted to heavyweight offline renderers.

Generating a high-quality rendering directly from landscape maps yields several challenges. Landscapes cover large spatial ranges, and therefore features at several scales are visible from a given viewpoint. Details and texture should have visual consistency (e.g., loose rocks objects on a rocky-textured cliff), but also a physical and geological consistency (e.g., screes usually stabilize on gentle 30% slopes). Authoring the distribution of these details requires time and expertise, especially to ensure an apparent consistency.

Our goal is to side-step this pipeline and develop a learning-based method that would directly output a plausible rendering of a set of input landscape maps from a given camera. Special care needs to be taken so that the algorithm robustly handles user interventions while preserving the consistency of the generated landscapes.

Approach

In a first step, we will consider only the elevation and convert it into a set of geometrical features, such as slope, orientation, or curvature. We will then project these features to display them from the camera location and explore models similar to pix2pix [3] to convert the projected feature to a high quality rendering in image space. We will extend this approach to enforce temporal consistency from camera movement and support high resolutions. We will also explore the use of differentiable rendering to learn the features instead of hardcoding them, for example using a cascaded network architecture, conditioning the features by the image-space rendering..

We will then add semantic maps representing landscape characteristics, such as grass, forests, screes, or cliffs. We will develop a “brush,” enabling the user to control the distribution of these elements [4]. We will build a specific architecture designed to ensure that the landscape elements respect the user’s input and are also consistent with physical laws. For that, we will use tools specific to simulation-based landscape generation[5].

We will train the model with a dataset reconstructed from a scan of a real mountain range, and a dataset generated from synthetic data. We assume that the synthetic dataset will provide less plausible results but gives us access to the ground truth of the various landscape layers. Except for elevation, these layers are not readily accessible from the scanned dataset, and therefore a possible extension includes a non-supervised classification of landscape elements[6].

Work environment and requirement

The internship will take place at Inria Sophia Antipolis. Inria will provide a monthly stipend of around 1100 euros for EU citizen in their final year of master, and 400 euros for other candidates.

Candidates should have strong programming and mathematical skills as well as knowledge in computer graphics, geometry processing and machine learning. Experience with physical simulation is a plus.

References

[1] A review of digital terrain modeling
E Galin, E Guérin, A Peytavie, G Cordonnier, MP Cani, B Benes, J Gain
Computer Graphics Forum 38 (2), 553-577 (Eurographics 2019)
https://hal.archives-ouvertes.fr/hal-02097510/file/A%20Review%20of%20Digital%20Terrain%20Modeling.pdf

[2] Co-Abstraction of Shape Collections
Mehmet Ersin Yumer, Levent Burak Kara
ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2012)
http://www.meyumer.com/pdfs/yumer-kara-SIGASIA2012.pdf

[3] Image-to-Image Translation with Conditional Adversarial Networks
Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, Alexei A. Efros
IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017
https://phillipi.github.io/pix2pix/

[4] Semantic Image Synthesis with Spatially-Adaptive Normalization
Taesung Park, Ming-Yu Liu, Ting-Chun Wang, Jun-Yan Zhu
IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2019
https://nvlabs.github.io/SPADE/

[5] Authoring landscapes by combining ecosystem and terrain erosion simulation
Guillaume Cordonnier, Eric Galin, James Gain, Bedrich Benes, Eric Guérin, Adrien Peytavie, Marie-Paule Cani
ACM Transactions on Graphics (Proc. SIGGRAPH 2017)
https://nvlabs.github.io/SPADE/

[6] Multi-Scale Label-Map Extraction for Texture Synthesis
Y. D. Lockerman, B. Sauvage, R. Allègre, J. M. Dischler, J. Dorsey, H. Rushmeier
ACM Transactions on Graphics (Proc. SIGGRAPH 2016)
https://nvlabs.github.io/SPADE/