View Synthesis in Light Field Volume Rendering Using Convolutional Neural Networks

Seán Martin, Master of Science in Computer Science
University of Dublin, Trinity College, 2018

Supervisor: Dr. Michael Manzke

In this dissertation, view synthesis for light field volume rendering using deep learning is investigated. Given a single volume rendered central view along with the camera parameters and an estimated depth map from the rendering pipeline, a structured grid of 8×8 views of the volume are synthesised. Firstly, this involves estimating a per pixel depth map during volume ray casting. This depth map is converted to a disparity map using the camera parameters for the purposes of image warping. Preprocessing is performed on the disparity map and reference view, such as image warping and normalisation, in order to be fed to a Convolutional Neural Network. Deep learning is applied with the prospect of accounting for inaccuracies in the depth map, non-Lambertian effects, and occlusions to improve the visual coherency of synthesised views over pure disparity based image warping. Multiple residual Convolutional Neural Network architectures are experimented with, and in particular, the results of three dimensional convolutions are compared with two dimensional convolutions. The experiments performed reveal that the networks all learnt to improve the visual quality of views that are far away from the central reference view, but failed to improve views that were close to the reference. In addition, convolutional neural networks are fast enough for view synthesis at interactive rates, and geometrical image warping is the bottleneck for light field synthesis in real-time. Overall, the methods presented have limitations, but show promising results for future research and expansion. Source code for this dissertation is available online at `https://github.com/flabby99/LF_View_Synthesis` and a video demonstration is available at `https://youtu.be/AQ4ec7Bgn1s`.