




Dissertation Defense

Dawei Yang

Learning to Generate 3D Training Data



Tuesday, June 30, 2020

10:00 am – 12:00 pm

bluejeans.com/5355936360/1024

ABSTRACT: Synthetic images rendered by graphics engines are a promising source for training deep neural networks, especially for vision and robotics tasks that involve perceiving 3D structures from RGB pixels. However, it is challenging to ensure that synthetic images can help train a deep network to perform well on real images. This is because graphics generation pipelines require numerous design decisions such as the selection of 3D shapes and the placement of the camera.

In this dissertation, we explore both supervised and unsupervised directions to automatically optimize those decisions. For supervised learning, we demonstrate how to optimize 3D parameters such that a trained network can generalize well to real images. We first show that we can construct a pure synthetic 3D shape dataset for deep learning to achieve state-of-the-art performance on a shape-from-shading benchmark. We further propose a hybrid gradient approach to accelerate the process. Our hybrid gradient is able to outperform classic black-box approaches on a selection of 3D perception tasks. For unsupervised learning, we propose a novelty metric for 3D parameter evolution based on deep autoregressive models. We show that without any extrinsic motivation, the novelty computed from autoregressive models can consistently encourage a random synthetic generator to produce more useful training data for 3D perception tasks.

Chairs: Profs. Jia Deng and David Fouhey