Abstract

Image data has been the staple requirement in today's computer vision ecosystem for a long time and expensive to acquire. There is great need for vast and diverse real world data to train the increasingly powerful deep neural networks to meet the challenges of today. A solution to many such challenges is to generate synthetic image data. However, the approach for this solution is rife with issues in quality, depth or semantic information to emulate a real image. Often these methods also require greater computationally intensive add-ons, and special post-processing techniques. Domain Adaptation facilitates this requirement and the same neural networks that require data can assist in acquiring synthetically produced data in a reliable manner.

The purpose of this master's thesis was to arrive at a domain adaptation technique that translate synthetic datasets into real world data using the concept of Neural Style Transfer. One of the recent approaches in Style Transfer is the High Resolution Network for Photorealistic Style Transfer. The transfer network has outperformed many current networks due to its ability to generate photorealism in any image and retaining semantic information crucial to perform object recognition on a real image. To achieve this goal, datasets have been generated out of synthetic images that aim to be transformed into real images capable of performing accurately on an SSD object detector.

The thesis concluded with a synthetic graphics rendered image dataset used as baseline to compare with stylized datasets to evaluate accuracy. The results yielded some open questions about the potential of perspective images styled with omnidirectional datasets as one dataset styled on perspective images produced results close to the baseline.