Towards a Perceptual Similarity Metric for 3D Shapes

Alexander Wang, Zhouqi Gong, Yuchen Zhou, Daniel Ritchie

Measuring the similarity between different shapes has various applications in computer graphics and computer vision, with tools like Chamfer distance and light field distance being common similarity metrics used for these applications. However, these purely mathematical formulas can be difficult to interpret to humans, to the point that it is possible for two shapes to appear different to humans while also having high similarity being reported by these metrics. As such, we argue that a similarity metric for 3D shapes should be based around human perceptual judgments. While there exists work into developing such metrics for 2D shapes, to the best of our knowledge, no work has been published for such a similarity metric for 3D shapes. Thus, we introduce an implementation of a perceptual similarity metric for 3D shapes. There are two primary components to this project: the first, to gather human perceptual judgments to determine similarity, and the second, to train a model based on the gathered data to best match the participants' responses. The first part, which was completed in the past, involved taking a base object from ModelNet40 and applying two distinct sets of transformations to it. 10,000 of these triplets of original object and two modifications were generated and respondents on Amazon Mechanical Turk were asked to judge which of the two modifications were more similar to the original object. In the second part, we chose to use a deep learning approach by using a modified Multi-View CNN as our model. Through tuning a variety of hyperparameters including norm regularization, margins used for loss, and the weight decay constant, testing different deep learning architectures, and generating additional data through data augmentation, we were able to develop an effective similarity metric that is able to respond the same way as human judgments 82.7% of the time.