Alana White

Capstone Abstract

5/12/2023

Domain-specific Language for Evaluating Likelihood of 3D Shape Labelings

In order to evaluate the likelihood of a proposed labeling of 3D shape regions of a mesh being accurate, we propose a Domain-specific Language (DSL) for evaluating properties of 3D regions. This DSL can evaluate properties of individual regions (e.g. its aspect ratio, whether it touches the ground), as well as relationships between regions (e.g. if a component provides support in a gravitational direction, relative position, connection). Using this DSL, a user can write rules that evaluate the likelihood of whether a region's labeling is correctly assigned (one such rule might be that a chair leg should be skinny, below a chair seat, and touching the ground). We then combine this with a Large Language Model (LLMs), such as GPT-4, to write rules such that we extract common sense knowledge from the LLM. We then use these rules to evaluate machine-learning based labeling proposals from the Neurally-Guided Shape Parser (NGSP) [1] to find aid in finding the most likely labelings for regions of a 3D shape.

This is still an ongoing research project in the Visual Computing Lab. Following my departure from Brown, development will continue, particularly in prompt engineering to produce code from LLMs that can more accurately and easily interface with the DSL.

References

[1] R. K. Jones, A. Habib, R. Hanocka and D. Ritchie, "The Neurally-Guided Shape Parser: Grammar-based Labeling of 3D Shape Regions with Approximate Inference," 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 2022, pp. 11604-11613, doi: 10.1109/CVPR52688.2022.01132.