

Weakly supervised 3D shape analysis

Online repositories contain millions of 3D shapes, providing data for a wide range of data-driven 3D modeling interfaces. Such interfaces facilitate, accelerate and democratise computer-aided design and 3D content creation. By automatically learning design rules and structural principles from training data, these interfaces allow even novice and casual users to design complex and functional objects. While these repositories often provide tags, textual descriptions, and soft categorisation to facilitate text-based search, they are typically not annotated with the fine-grained component, correspondence and structural information needed to learn rich structural models.

This challenge is typically addressed in two contrasting ways, standard in the machine learning community: supervised methods, where the training data is meticulously annotated with all the needed information, and


unsupervised methods, which do not rely on annotation but rather on general geometric principles. Both these approaches have shortcomings: gathering large scale annotations is time-consuming and expensive, and reasoning without annotations is error-prone.

In our recent work, we have explored a third way of addressing this challenge: relying on easily gathered partial annotations of a large training set. This approach is known as weak supervision. We have proposed new algorithms for the following fundamental problems in shape analysis. In each case, the training annotations provide only a coarse and

incomplete indication of the information required in the result.

- Learning hierarchical structural models from shapes segmented into parts, but without training labels for these parts or hierarchies defined on them (SIGGRAPH 2017)
- Finding point correspondences between shapes, with labeled segmentations for the training shapes, but without training correspondences (ToG 2017)
- Learning how to complete a partial shape, with segmented training shapes, but without training labels for these segments (SIGGRAPH Asia 2017)
- Segmenting a shape into labeled parts, based on unsegmented training shapes annotated only with part presence / absence (ongoing)

Our work heavily uses machine learning techniques based on deep neural networks. The diverse geometry of 3D shapes, however, implies that they cannot easily be mapped to one-dimensional sequences as in text/speech processing, or 2D grids as in computer vision. Our work employs new architectures and adaptations of these networks, based on recursive, volumetric, view-based and point-based representations.



Going forward, we are excited to apply weakly supervised methods to new classes of problems in 3D shape processing. We believe such methods can effectively solve the problem of data scarcity in learning-based algorithms, and provide the foundation for new approaches to higher-level shape-based reasoning.

