מכון ויצמן למדע

THE WEIZMANN INSTITUTE OF SCIENCE

FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Vision and Robotics Seminar

Room 1 ,Ziskind Building
on Thursday, Jan 30, 2020
at 12:15

Oshri Halimi
Technion

Handling the Unknown with Non-Rigid Geometric Invariants

Abstract:

My goal is to demonstrate how geometric priors can replace the requirement for annotated 3D data. In this context I'll present two of my works. First, I will present a deep learning framework that estimates dense correspondence between articulated 3D shapes without using any ground truth labeling which I presented as an oral presentation at CVPR 2019: "Unsupervised Learning of Dense Shape Correspondence". We demonstrated that our method is applicable to full and partial 3D models as well as to realistic scans. The problem of incomplete 3D data can be encountered also in many other different scenarios. One such interesting problem of great practical importance is shape completion, to which I'll dedicate the second part of my lecture.
It is common to encounter situations where there is a considerable distinction between the scanned 3D model and the final rendered one. The distinction can be attributed to occlusions or directional view of the target, when the scanning device is localized in space. I will demonstrate that geometric priors can guide learning algorithms in the task of 3D model completion from partial observation. To this end, I'll present our recent work: "The Whole is Greater than the Sum of its Non-Rigid Parts". This famous declaration of Aristotle was adopted to explain human perception by the Gestalt psychology school of thought in the twentieth century. Here, we claim that observing part of an object which was previously acquired as a whole, one could deal with both partial matching and shape completion in a holistic manner. More specifically, given the geometry of a full, articulated object in a given pose, as well as a partial scan of the same object in a different pose, we address the problem of matching the part to the whole while simultaneously reconstructing the new pose from its partial observation. Our approach is data-driven, and takes the form of a Siamese autoencoder without the requirement of a consistent vertex labeling at inference time; as such, it can be used on unorganized point clouds as well as on triangle meshes. We demonstrate the practical effectiveness of our model in the applications of single-view deformable shape completion and dense shape correspondence, both on synthetic and real-world geometric data, where we outperform prior work on these tasks by a large margin.