Computer Vision and 3D Reconstruction

Dr George Vogiatzis, Aston University

Computer Vision, the field that is concerned with enabling machines to “see”, has an illustrious 50 year history within Engineering. While the field now encompasses a multitude of different aspects of “seeing”, (e.g. Object detection, recognition, image search) at its very heart has always been the problem of 3D reconstruction, I.e. the task of understanding the 3D world by building a mathematical representation of that world from sequences of images. The applications range widely, from cultural heritage preservation, to reverse engineering and 3D special effects for films and games. In the past few years, 3D reconstruction technology has matured to the level of commercial applications, with a number of systems becoming available for use even by non-expert users. And while all this success has generated excitement for the potential of the technology, it has also highlighted how far we still are from our original goal of understanding the 3D world from visual data. While existing systems are very good with generating accurate low-level geometry and appearance representations, we have not been so successful with higher level semantic interpretation which arguably is closer to what it means to “understand” a 3D scene. Unfortunately this is crucial for a multitude of applications that require deeper structural modelling that goes beyond simple 3D primitives.

In this talk I will give a birds’ eye view of the whole domain describing some of the advances that led to today’s commercial technologies. At the same I will highlight the challenges of semantic interpretation and how Deep Learning holds the promise of bridging the gap between low level 3D reconstruction and scene understanding. As a case study of that approach we will use some of our recent work on architectural façade parsing.