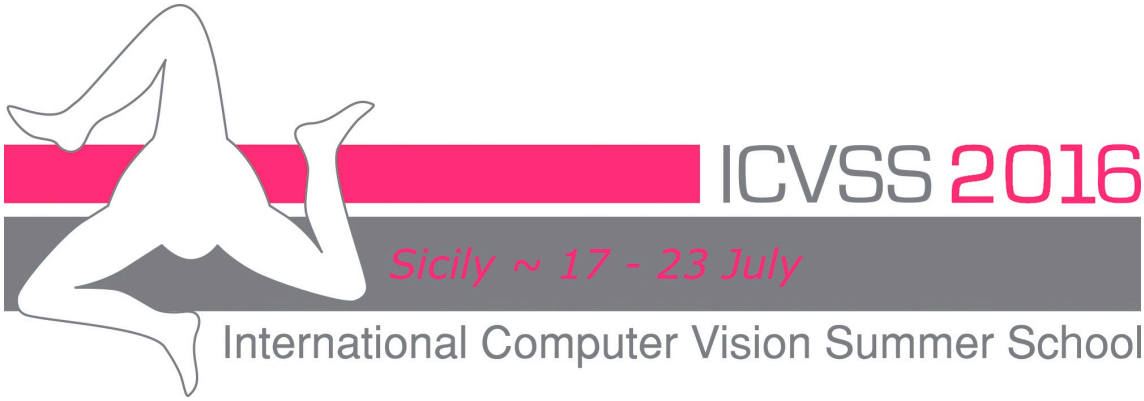

The Future and Generative Models: A Case Study of Human Bodies in Motion

Michael J. Black
Max Planck Institute for Intelligent Systems

Abstract

There are several technological revolutions going on simultaneously and, while deep learning is the most visible, others may be even more important. Deep learning has been enabled by big datasets of images and platforms that enable human labeling on a large scale. At the same time a revolution is taking place in computer graphics with open source gaming engines bringing high-quality rendering to everyone. Additionally, while we have seen the tiny beginnings of 3D scanning with devices like Kinect, the technology is poised to become widespread in consumer devices in the next few years. Likewise, 3D printing and virtual reality applications suggest that databases of 3D CAD models are poised to expand. The great leaps in deep learning have resulted from category-level labeling, which is relatively easy for humans. The labeling of metric properties of 3D scenes and objects will prove much harder. This argues for the sensible use of generative models of the 3D world that can be fit to relatively small amounts of training data and yet have have strong generalization ability. A key lesson of current deep learning methods is that simple models, that are easy to train, are often preferable to more powerful models that are hard to train. Future generative models will exploit this insight to enable end-to-end training. The future will also likely combine discriminative, bottom up, pattern recognition methods with generative,

top down, models. This promises a return to the early roots of computer vision but with new tools.

As a case study this talk will look at recent progress in 3D human pose and shape estimation from images, video, and RGB-D. The human body is complex and deformable. For many applications in computer vision, graphics, fashion, and medicine, having a realistic, low-dimensional, generative 3D model of the body is useful. Getting a good one, however, is difficult. This talk will review the history of our work on learning 3D models of the human body from 3D scans. It will try to answer "what" is a body model, "why" it is useful, and "how" to build one. It will summarize how to accurately align 3D meshes of bodies in arbitrary poses, how to build a statistical model of body shape and non-rigid pose variation, and how to fit such models to data including 3D scans, Kinect data, or mocap markers. The talk will also describe recent work on capturing and modeling the dynamics of soft tissue motion using our one-of-a-kind 4D body scanner. Finally, it will address the full problem of going from images to 3D shape and pose using a combination of deep learning and generative models.

Keywords
Generative Models, 3d Shape, 3D human pose and shape estimation from images, video, and RGB-D

Background material
Shape and from Kinect (ICCV'15): https://youtu.be/OimYQMgqnqM
SMPL body model (SIGGRAPH Asia '15): http://smpl.is.tue.mpg.de/
Dyna (SIGGRAPH '15): http://dyna.is.tue.mpg.de/
MoSh (SIGGRAPH Asia '14): http://mosh.is.tue.mpg.de/
FAUST (CVPR'14): https://youtu.be/usV_K_rmCxM

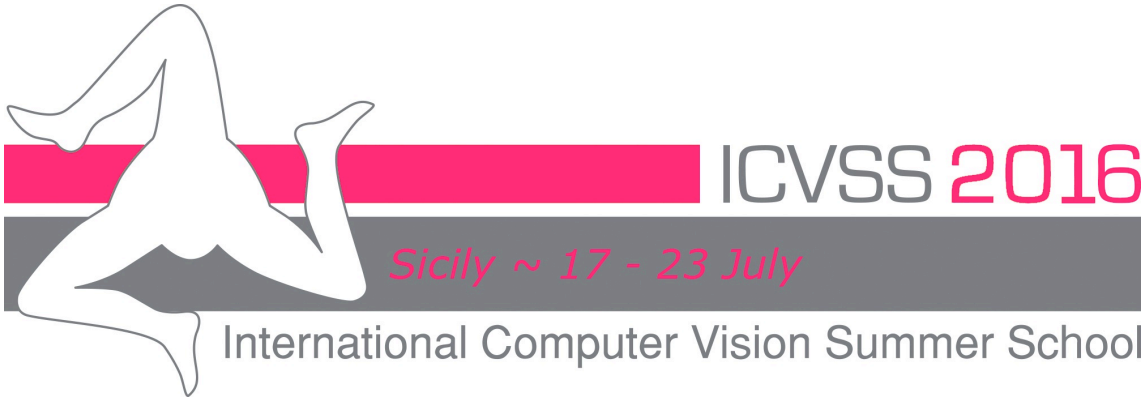

Biography of the Speaker

Michael Black received his B.Sc. from the University of British Columbia (1985), his M.S. from Stanford (1989), and his Ph.D. from Yale University (1992). After post-doctoral research at the University of Toronto, he worked at Xerox PARC as a member of research staff and an area manager. From 2000 to 2010 he was on the faculty of Brown University in the Department of Computer Science (Assoc. Prof. 2000-2004, Prof. 2004-2010). He is one of the founding directors at the Max Planck Institute for Intelligent Systems in Tübingen, Germany, where he leads the Perceiving Systems department. He is an Honorarprofessor at the University of Tuebingen, Visiting Professor at ETH Zürich, and Adjunct Professor (Research) at Brown University. His work has won several awards including the IEEE Computer Society Outstanding Paper Award (1991), Honorable Mention for the Marr Prize (1999 and 2005), the 2010 Koenderink Prize for Fundamental Contributions in Computer Vision, and the 2013 Helmholtz Prize for work that has stood the test of time. He is a foreign member of the Royal Swedish Academy of Sciences. He is also a co-founder and board member of Body Labs Inc.