Computational photography and low-level vision

Bill Freeman
Massachusetts Institute of Technology, USA

Abstract

First half: Learning low-level vision

Low-level vision involves measuring local evidence and propagating it over space. I'll review ways to represent and measure local evidence, and ways to propagate it over space.

This simple framework, in a non-parameteric, example-based approach, can be applied to many problems, including:

- estimating high resolution images from low-resolution ones
- estimating motion,
- interpreting line drawings,
- analyzing shading and reflectance.

I'll take us through some examples.

Second half: Photographing events over time

Artists and scientists have used photography to portray events occurring over a broad range of timescales, from pico-seconds to tera-seconds. I'll review some of the hardware and computational techniques used to create photographs over these different timescales.

The talk addresses work by artists, photographers, and computer scientists, including the artists Haas, Salavon, Funch, Jones, the photographers Marey, Muybridge, Haas, Prokudin-Gorskii, Doc Edgerton, and research from Microsoft, U. Washington, MIT, Harvard, Google, and Berkeley.

Describing events at the short time scales is mostly a hardware issue. In contrast, capturing events over longer time scales often involves tracking objects or removing changes in illumination or viewpoint, and thus is a computational problem. There is still much to be done to capture and convey events over long timescales well.

Syllabus: Computational Photography, Super Resolution, Registration, Motion Estimation, Analyzing Shading and Reflectance, Graphical Models