Tutorial 7: Photometric 3D Reconstruction

Modern 3D computer vision methods, represented by multi-view stereo and structure-from-motion, have achieved faithful 3D reconstruction from a set of images. But are the reconstruction quality and density really sufficient for your purpose? Despite requiring more controlled setups than multi-view stereo, photometric approaches have proven to be invaluable tools in applications such as Hollywood movies, industrial quality inspection, etc. since they can reconstruct fine surface details at superior quality. This tutorial covers a thread of ``photometric'' approaches to high-fidelity 3D reconstruction, which enables truly dense 3D estimation at the level of pixel-level details from shading observations.

The tutorial will mainly cover photometric stereo techniques that take as input a set of images observed under different illumination conditions from a fixed viewpoint to compute the shape in the form of surface normals with the same high resolution as the 2D image. While conventional photometric stereo methods make various assumptions over reflectance and illumination, they are being relaxed in modern methods by powerful convex optimization and machine learning approaches so as to be practical in diverse scenarios. In addition, newer datasets for evaluation and convenient light source calibration techniques have been introduced to make photometric stereo methods more accessible and practical. This tutorial will walk through these recent advances.

By attending this tutorial, you will learn (1) relevant background on photometric methods, (2) basic photometric stereo, (3) calibrated, uncalibrated, and semi-calibrated settings in photometric stereo, (4) new robust estimation techniques in photometric stereo, (5) machine learning approaches for photometric stereo, (6) dataset and evaluation methods, (7) light source calibration techniques, (8) 3D shape reconstruction from surface normal, and (9) implementation details with Python. The tutorial will be as self-contained as possible. Expected target audience are students, researchers, and practitioners, who may not be familiar with photometric 3D reconstruction techniques but are interested in high-quality 3D acquisition from images.

The tutorial will be delivered in three parts:

Part I: Introduction to Photometric Stereo
- Basic concepts about shape, illumination, and reflectance
- Conventional Lambertian least-squares photometric stereo
- Calibrated and uncalibrated photometric stereo, and semi-calibrated photometric stereo

Part II: Recent advances in Photometric Stereo
- Robust photometric stereo via sparse regression
- Photometric Stereo meets deep learning
- Benchmark datasets and evaluation

Part III: Photometric Stereo for Practitioners
- Light source calibration
- Surface normal to depth map
- Implementation details in Python

Speakers:

Michael Waechter
Osaka University, Japan

Boxin Shi
Peking University, China

Yasuyuki Matsushita,
Osaka University, Japan