



Trading off Invariance and Discriminative Power in Low-Level representations

[Stefano Soatto](#)

UCLA, USA

Abstract

Local features have become a staple of most recognition systems, whether of individual objects or object categories. Why are local features useful? Why do we want to break down images into little pieces? And what kind of "signature" do we want to use to describe these pieces? The image of an object depends on its properties (shape, reflectance), but also on other accidents of the image formation process: Occlusions, viewpoint, illumination, intrinsic variability.

Local features are designed to trade off invariance to these accidents, also known as "nuisance factors", while at the same time maintaining as much "information" about the object as possible. For the most part, this tradeoff is set in an ad-hoc fashion, using common sense, engineering practice, inspiration from the primate visual system. Is there a systematic way to design features? Or to set their parameters? How does a choice of local feature play in the design of the classifier downstream? What are the limitations, and the strenghts, of existing local features? In this lecture, we will address some of these questions, drawing from recent literature that the students are supposed to have read prior to the commencement of the school.

Reading list/Syllabus:

- D. G. Lowe. *Distinctive image features from scale-invariant keypoints*. *IJCV*, 2(60):91–110, 2004.
- K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. *A comparison of affine region detectors*. *IJCV*, 1(60):63–86, 2004.
- P. Moreels and P. Perona. *Evaluation of features detectors and descriptors based on 3D objects*. *In Proc. ICCV*, 2005.
- S. Roth and M. J. Black. *On the spatial statistics of optical flow*. *In Proc. ICCV*, 2005.