


Ten Years of Pedestrian Detection, What Have We Learned?

R. Benenson, M. Omran, J. Hosang, and B. Schiele (ECCV 2014)



- ▶ Aim is to review progress over the last decade of pedestrian detection, & try to quantify which ideas had the most impact on final detection quality
- ▶ Evaluates on Caltech-USA, INRIA and KITTI datasets for comparing methods
- ▶ Conclusions:
 - ▶ There is no conclusive empirical evidence indicating that whether non-linear kernels provide meaningful gains over linear kernel
 - ▶ The 3 families of pedestrian detectors (DPMs, decision forests, deep networks) are based on different learning techniques, their results are surprisingly close
 - ▶ Multi-scale models provide a simple and generic extension to existing detectors. Despite consistent improvements, their contribution to the final quality is minor
 - ▶ Most of the progress can be attributed to the improvement in features alone
 - ▶ Combining the detector ingredients found to work well (better features, optical flow, and context) shows that these ingredients are mostly complementary