

Pedestrian Detection at 100 Frames Per Second

R. Benenson, M. Mathias, R. Timofte, and L. V. Gool (CVPR 2012)



N models, 1 image scale
(a) Naive approach



1 model, N image scales
(b) Traditional approach



1 model, N/K image scales
(c) FPDW approach



N/K models, 1 image scale
(d) Our approach

- ▶ Fast and high quality pedestrian detection
- ▶ Two new algorithmic speed-ups:
 - ▶ Exploiting geometric context extracted from stereo images
 - ▶ Efficiently handling different scales
- ▶ Object detection without image resizing using stixels
- ▶ Similar to Viola and Jones: scale the features not the images, applied to HOG-like features
- ▶ Detections at 50 fps (135 fps on CPU+GPU)
- ▶ Evaluated on INRIA Persons and Bahnhof sequence