



Figure 1: Log-average miss rate versus the runtime of each detector on 640×480 images from the Caltech Pedestrian Dataset. Run times of all detectors are normalized to the rate of a single modern machine, hence all times are directly comparable. (Note that the VJ implementation used did not utilize scale invariance and hence its slow speed). Legends are ordered by detector speed, reported in frames per second (fps). **(a)** Speed vs. performance for pedestrians 100 pixels and up (the ‘near scale’ setting). **(b)** Speed vs. performance for pedestrians 50 pixels and up (the ‘reasonable’ setting). While the slowest detector happens to also be the most accurate (MultiFtr+Motion), for pedestrians over 50 pixels the two fastest detectors, FPDW and Crosstalk, are also among the most accurate, respectively.