Learning Detectors Quickly with Stationary Statistics

Jack Valmadre, Sridha Sridharan and Simon Lucey
Queensland University of Technology, Australia CSIRO, Australia Carnegie Mellon University, USA

Two recents methods avoid hard negative mining when training a classifier for detection

Discriminative decorrelation

(Hariharan et al. 2012)

The set of natural images exhibits stationarity

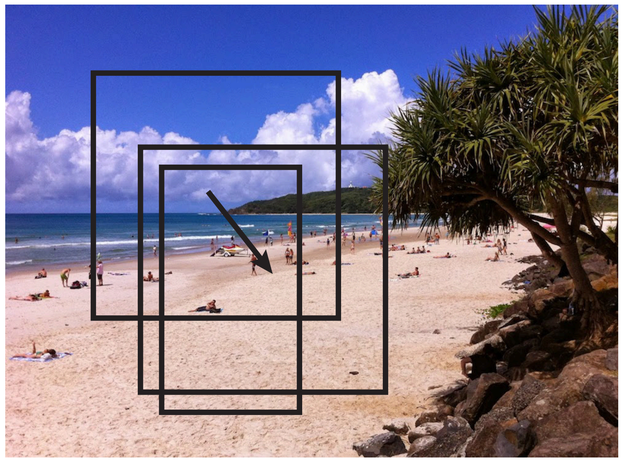

LDA with Toeplitz covariance matrix

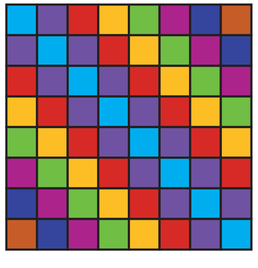

$$S_{ij} = g[j - i]$$

$$w = S^{-1}(\bar{x}_{\text{pos}} - \bar{x}_{\text{neg}})$$

Correlation filters

(Henriques et al. 2013, Boddeti et al. 2013, Kiani et al. 2013)

Ridge regression with circular shifts

Leads to circulant covariance matrix

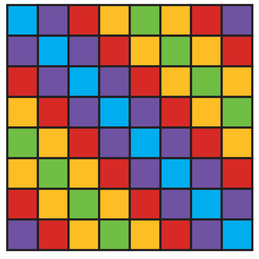

$$S_{ij} = h[j - i \bmod m]$$

$$w = \mathcal{F}^{-1}\{\text{diag}(\hat{h})^{-1}\hat{b}\}$$

Solved efficiently in the Fourier domain

Need to extract windows of desired size

The best of both worlds

Use FFT to compute Toeplitz covariance

Obtain circulant covariance for arbitrary window size from Toeplitz covariance:

$$h[\delta] = (1 - \theta)g[\delta] + \theta g[\delta - m] \qquad \theta = \delta/m$$

(Also the nearest circulant matrix)

Invert large block two-level Toeplitz matrices using conj. grad. with circulant pre-cond.

Heuristic

(Henriques et al. 2013)

Train with circulant covariance then crop outer pixels

Results

Pedestrian detection, HOG images, 31 channels

INRIA

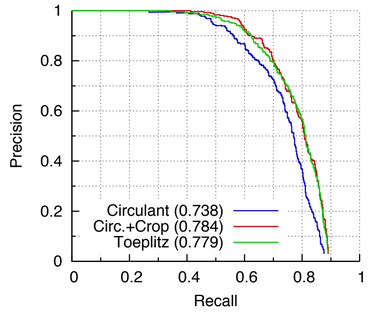

Caltech

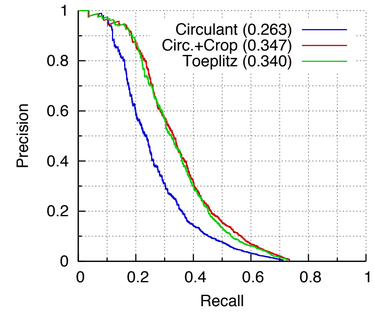

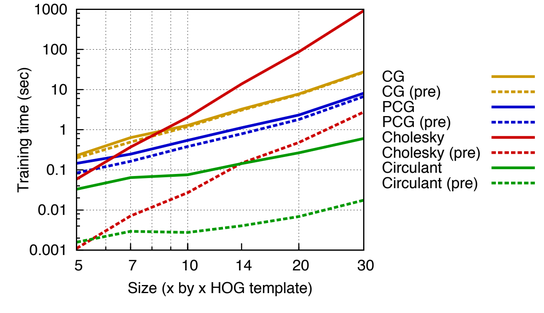

Time

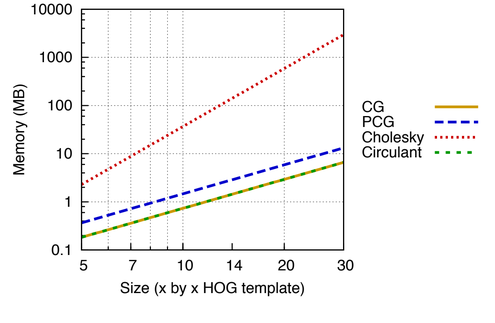

Memory