

Rapid Deformable Object Detection using Dual-Tree Branch and Bound

Iasonas Kokkinos

Center for Visual Computing
Ecole Centrale Paris

Galen Team
INRIA-Saclay



Deformable Part Model score:

$$S(x) = \sum_{p=1}^P \max_{x_p} \left[U_p(x_p) - (x_p - x - \mu_p)^T \begin{bmatrix} c_{p,1} & 0 \\ 0 & c_{p,2} \end{bmatrix} (x_p - x - \mu_p) \right]$$

$\arg \max S(x)$

Felzenszwalb, Girshick, McAllester, Ramanan, PAMI 2010



Max-Product/Dynamic Programming $O(PN^2)$

Generalized Distance Transforms $O(PN)$

$$N = |\{x\}|$$

Approximate solutions:

Y. Chen et al. Rapid inference on a novel and/or graph for object detection, NIPS 2007

P. F. Felzenszwalb, R. B. Girshick, and D. A. McAllester. Cascade object detection with DPMs CVPR 2010

I. Kokkinos and A. Yuille. HOP: Hierarchical Object Parsing, CVPR, 2009

M. Pedersoli, A. Vedaldi, and J. Gonzalez. A coarse-to-fine approach for object detection, CVPR 2011

B. Sapp, A. Toshev, and B. Taskar. Cascaded models for articulated pose estimation, ECCV, 2010

Rapid Deformable Object Detection using Dual-Tree Branch and Bound


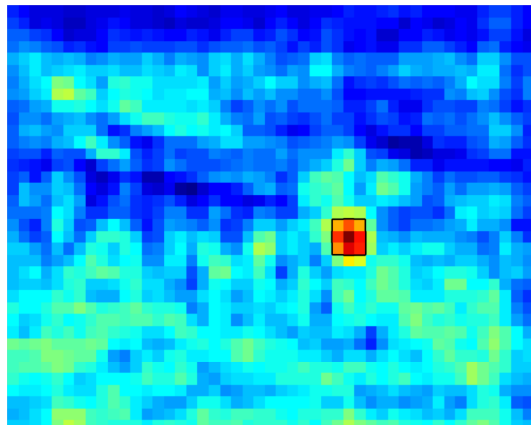
Branch and Bound: exact solution

Best-case: $O(P \log N)$


Input & Detection result




Detector score $S(x)$



BB for $\arg \max_x S(x)$



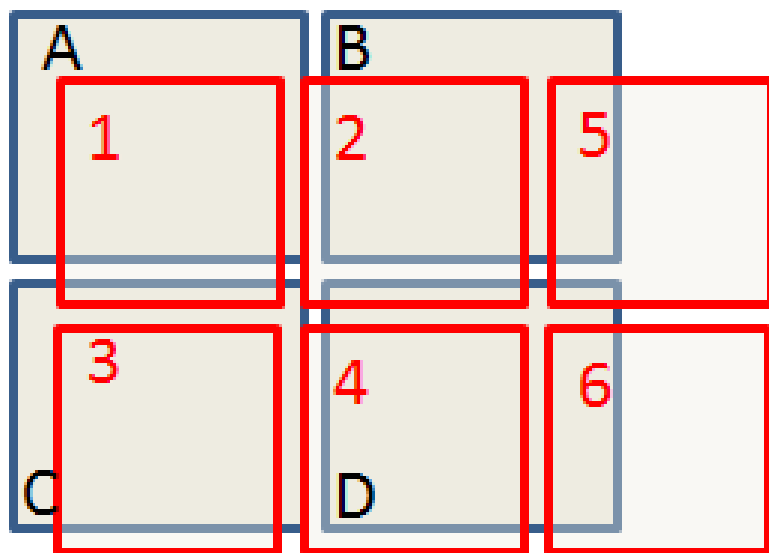
BB for $S(x) \geq -1$




Rapid Deformable Object Detection using Dual-Tree Branch and Bound

Dual Tree Branch and Bound: Dual Trees [1] + ESS [2] for DPMs

Dual Trees: bound $\max_{x \in B_A} \max_{x_p \in B_1} - (x_p - x)^T \begin{bmatrix} c_1 & 0 \\ 0 & c_2 \end{bmatrix} (x_p - x)$



Upper and lower bounds




High: objects in A & parts in 1
 Low: objects in A & parts in 6

[1] A.G. Gray and A.W. Moore. Nonparametric density estimation: Toward computational tractability. ICDM 2003

[2] C. Lampert, M. Blaschko, and T. Hofmann. Beyond sliding windows- efficient subwindow search. CVPR 2008

Rapid Deformable Object Detection using Dual-Tree Branch and Bound

Comparisons with GDT on 3000 images (precomputed unary terms)



Single object: speedup increases with threshold

Multiple-objects, 1-best: 100-fold speedup for > 50% of images

Current bottleneck: unary term computation – *amenable to bounding*

Code available from
<http://vision.mas.ecp.fr/Personnel/iasonas/>