

Fast Bounding Box Estimation based Face Detection

Bala Subburaman Venkatesh^{1,2} and Sébastien Marcel¹

¹Idiap Research Institute, 1920, Martigny, Switzerland

²École Polytechnique Fédérale de Lausanne (EPFL), 1015, Lausanne, Switzerland

{venkatesh.bala,marcel}@idiap.ch

<http://www.idiap.ch>, <http://www.epfl.ch>

Abstract. The sliding window approach is the most widely used technique to detect an object from an image. In the past few years, classifiers have been improved in many ways to increase the scanning speed. Apart from the classifier design (such as cascade), the scanning speed also depends on number of different factors (such as grid spacing, and scale at which the image is searched). When the scanning grid spacing is larger than the tolerance of the trained classifier it suffers from low detections. In this paper we present a technique to reduce the number of miss detections while increasing the grid spacing when using the sliding window approach for object detection. This is achieved by using a small patch to predict the bounding box of an object within a local search area. To achieve speed it is necessary that the bounding box prediction is comparable or better than the time it takes in average for the object classifier to reject a subwindow. We use simple features and a decision tree as it proved to be efficient for our application. We analyze the effect of patch size on bounding box estimation and also evaluate our approach on benchmark face database (CMU+MIT). We also report our results on the new FDDB dataset [1]. Experimental evaluation shows better detection rate and speed with our proposed approach for larger grid spacing when compared to standard scanning technique.

1 Introduction

The sliding window approach is the most common technique used for object detection [2–4]. A classifier is evaluated at every location, and an object is detected when the classifier response is above a preset threshold. Many systems need face processing tasks (detection, tracking, recognition), and needing them to run in real-time with out loosing much of individual performance has become a challenging task. Cascades introduced by Viola et al. [4] speed up the detection by rejecting the background quickly and spending more time on object like regions. Although cascades were introduced, scanning with fine grid spacing is still computationally expensive. To increase the scanning speed one approach is to train a classifier with perturbed training data to handle small shifts in the object location. Another simple approach is to increase the grid spacing (decreases

the number of subwindows being evaluated). Unfortunately, as the grid spacing is increased the number of detection decreases rapidly. In this paper we focus on increasing the detection rate and speed of the sliding window approach by building a bounding box estimator with high performance (speed and accuracy). To reach that objective we propose a method to predict the bounding box of an object, using a simple yet effective binary test and a decision tree. We show that this method reduces the miss detections while increasing the scanning grid spacing. In this work we don't intend to increase the performance of the main face classifier, but rather try to improve the detection rate for larger grid spacing which also has an effect on scanning speed.

This paper is organized as follows. In next Section we describe the related work, and in Section 3, we describe our approach on how we increase the detection rate and speed by using bounding box estimation. In Section 4, we show our experiment results and finally, conclusion and future work are given in Section 5.

2 Related work

Object detection has been approached in many different ways in the literature. Either parts of the object or the whole object have to be classified in some way. The main idea of detecting parts rather than whole image object is to reduce the variability in appearance of the object. Bounding box estimation can also be posed as an object part identification problem. This has been done using different features but the most popular ones are scale-invariant feature transform (SIFT) [5] and Ferns [6].

SIFT descriptors have been popular as it has proved to be robust to illumination changes, but the computation of SIFT descriptors turns out to be costly. Feature matching is done with Nearest Neighbor approach, but to reduce the complexity, an approximate algorithm called the Best Bin First (BBF) algorithm proposed by Beis et al. [7] is used. In [6] an alternative feature called Ferns was introduced which showed comparable or better performance than SIFT features for patch identification. Ferns consists a set of binary features, and the binary feature is obtained by comparing the intensity value of two pixels. A Semi-Naive Bayesian classifier is used to identify a patch, where the class posterior probabilities are modeled with hundreds of binary features.

In [8] the bounding box and pose are estimated before giving the hypothesized window to a pose specific Support Vector Machine (SVM) classifier. Histogram of SIFT like features, and a Naive Bayesian approach is used to learn different poses and bounding box. The bounding box estimation is carried out in two steps. First a fixed sized window is used to infer the aspect ratio, and then the area is estimated. This approach maximizes the overlap between the estimated location of the box and the ground truth. In [9], a component based face detector is described. Their system consists of a two-level hierarchy of SVM classifiers. On the first level, components of face are independently detected and the second level, the geometrical configuration of the detected components are checked

with face model. While their technique might perform better, but since many component classifiers are evaluated the speed could be an issue.

Object detection using generalized Hough transform has also gained in popularity. One of the model proposed by Leibe et al. [10], Implicit Shape model (ISM), consists of class specific codebook of local appearance from the object category and spatial probability distribution of where the codebook entry may be found on object. During recognition this information is used to perform a generalized Hough transform in a probabilistic framework. However as pointed out by Jürgen et al. [11], codebook-based Hough transform comes at a significant computational price, and the authors have suggested using random forest to directly learn a mapping between the appearance of an image patch and its Hough vote, more precisely a probabilistic vote about the position of an object centroid. In shape regression machine [12], parameters such as translation, scale and rotation are estimated using regression approach and was applied on medical images with improved speed in detection. We can see that to speed up the detection we need to estimate the location of the object as quickly and accurately as possible.

Our work is inspired by [11] in the way that the object centers are estimated by using Hough forest. Each leaf node in the tree gives a probabilistic votes of the object centroid. The hypothesis which is tested at each node is a simple comparison of values (can be intensity or gradient in x and y direction) at two locations (in some sense similar to [6]). In [11], a forest with 15 trees is used which takes considerable amount of time if a dense scan is performed. Therefore we decided to use only a single decision tree for estimating the offset (bounding box) of the face. The advantage of using a decision tree is that the number of tests that needs to be performed grows only logarithmically thus saving computational time at runtime.

3 Proposed approach

In this section we first analyze the standard sliding window approach, and then describe how to reduce the miss detections with larger grid spacing by using a bounding box estimator. We then describe how a decision tree is learnt for this task.

3.1 Analysis of standard sliding window technique

The standard sliding window technique with regular grid scan is shown in Figure 1(a), where a classifier C_{object} is placed on the scanning grid and checks if it is an object or not. We start by formulating the chance of hit \mathcal{H}_c , as the chances for the target object to be within the classifier detection range, with respect to the scanning grid interval (w_s, h_s) , and to the translation tolerance (w_t, h_t) of the classifier C_{object} , (see Figure 1(a)).

$$\mathcal{H}_c \approx \frac{w_t h_t}{w_s h_s} \quad (1)$$

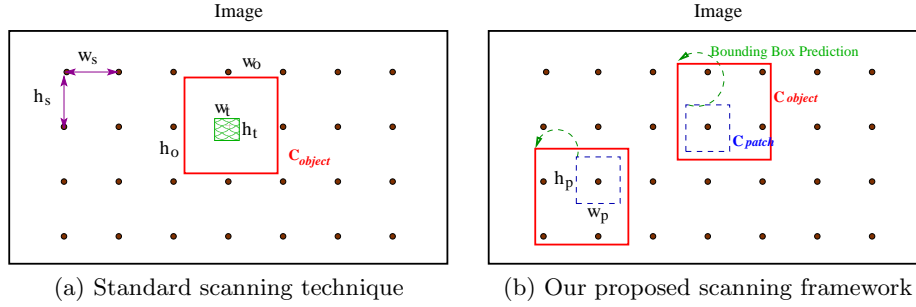


Fig. 1. Standard scanning technique vs our proposed scanning framework. The dots represent the scanning grid with interval (w_s, h_s) , target object size (w_o, h_o) , translation tolerance (w_t, h_t) of target object classifier C_{object} , target patch size (w_p, h_p) , and target patch classifier C_{patch} . The classifier C_{patch} predicts the bounding box for C_{object} in our approach.

As an example, let's assume that the object present in the image is of the same size as the classifier is trained with, if $w_t = h_t = 3$ and $w_s = h_s = 6$ then the chance of getting a hit \mathcal{H}_c is 0.25, which is very low. For \mathcal{H}_c greater than 1, means that the classifier has more chances to detect the object. As we decrease w_s and h_s (a finer search), \mathcal{H}_c increases, while scanning speed decreases (slower). Our goal is to increase \mathcal{H}_c without decreasing too much of the scanning speed (thus making it faster), which is described below.

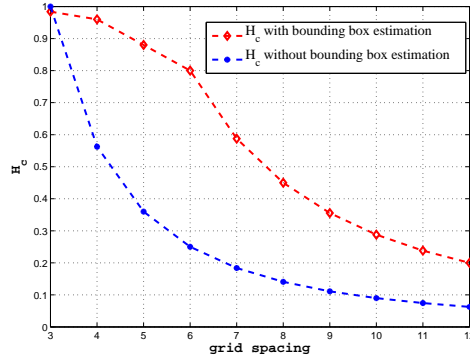


Fig. 2. Estimated chance of hit \mathcal{H}_c with and without bounding box estimation with respect to scanning grid spacing.

3.2 Chance of hit with our approach

In this subsection we explain how our method increases the probability of hit. Figure 1(b), shows the proposed scanning framework. The classifier C_{patch} is evaluated on a regular grid, while the main classifier C_{object} is placed on location predicted by C_{patch} . Assuming that we have a classifier C_{patch} that predicts the patch location correctly within the translation tolerance (w_t, h_t) of the classifier C_{object} , with prediction rate d_p , then the chance of hit can be approximately given by:

$$\mathcal{H}_c \approx d_p \mathcal{H}_p \quad (2)$$

$$\mathcal{H}_p = \frac{(w_o - w_p + 1)(h_o - h_p + 1)}{w_s h_s} \quad (3)$$

where \mathcal{H}_p is the chance of hit for the patch, (w_p, h_p) is the patch width and height, and (w_o, h_o) is the object width and height, with constraints $w_p < w_o$ and $h_p < h_o$ (see Figure 1(b)). For example if, $w_p = h_p = 14$, $w_o = h_o = 19$, $w_s = h_s = 6$, and $d_p = 0.8$ (this value is taken from our experiment results), we get $\mathcal{H}_c = 0.8$, which is 55% greater than standard scanning approach. The smaller the patch size is, the more the spacing between the grid can be, for a increase in scanning speed. Unfortunately at the same time estimating the bounding box becomes complex as individual patch will contain less and less information for distinguishing one from another. Figure 2 shows the chance of hit with and without bounding box estimation. To generate the plot we have used the same values as given in the example.

3.3 Bounding box estimation

The key idea for our algorithm lies in estimating the bounding box with high performance (speed and accuracy). We intend to use decision tree as it proved to be simple and efficient for our task. The decision tree is trained in a supervised manner. The training data, binary test and tree construction, and data stored in leaf node is described below.

Training data. For an object of size $w_o \times h_o$, and patch size of $w_p \times h_p$, we can have $\frac{(w_o - w_p + 1)(h_o - h_p + 1)}{2}$ number of overlapping patches (see Figure 3(a)). We represent a set of patches by $\{\mathcal{P}_i = (\mathcal{I}_i, \mathbf{d}_i)\}$, where \mathcal{I}_i is the appearance of the patch and \mathbf{d}_i is the offset of the patch. The offset vector \mathbf{d}_i is a 2D vector representing (x, y) shifts from the object center or from a fixed point in the object.

Binary test. In a decision tree \mathcal{T} , a test has to be performed at a node. We first consider a simple binary test introduced in [11, 6], which is given by:

$$t_f(\mathcal{I}) = \begin{cases} 1 & \text{if } \mathcal{I}(x, y) \leq \mathcal{I}(x_t, y_t) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

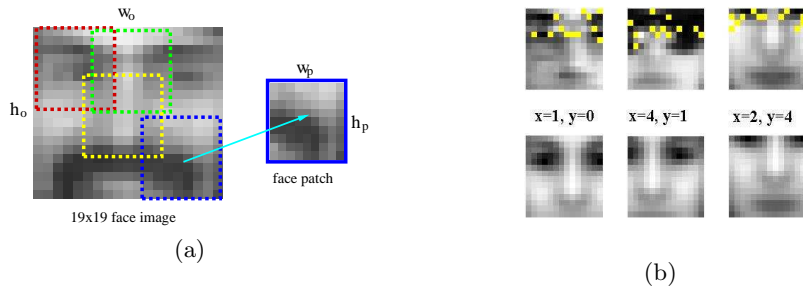


Fig. 3. (a) Examples of some overlapping face patches, All patches lie within the face region. (b) In this figure each column corresponds to a leaf node. The first row shows the average of the test images that arrives at a leaf node. The pixel locations that are evaluated ($t_{\mu f}$) at the nodes of a tree are also indicated. The second row provides the estimated offset value of the patches reaching a leaf node. The third row shows the average image of all the test patches having the offset value corresponding to the leaf node. The patch size of 14x14 is shown here.

where (x, y) and (x', y') are two locations in the patch \mathcal{I} . We also propose a new test which is given by:

$$t_{\mu f}(\mathcal{I}) = \begin{cases} 1 & \text{if } \mathcal{I}(x, y) \leq \text{avg}(\mathcal{I}) \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where $\text{avg}(\mathcal{I})$ is the average of the pixel values in the patch \mathcal{I} . Our test requires only half the number of pixel access compared to the previous test, but requires an integral image to quickly calculate the average value.

Tree construction During training, each non-leaf node picks the binary test that splits the training samples in an optimal way. We use the offset uncertainty as in [11] which is defined as:

$$U(A) = \sum_{i \in A} (\mathbf{d}_i - \mathbf{d}_A)^2 \quad (6)$$

where \mathbf{d}_A is the mean offset vector over all object patches in the set $A = \{\mathcal{P}_i = (\mathcal{I}_i, \mathbf{d}_i)\}$. A binary test t^* is chosen to minimize the following expression:

$$t^* = \arg \min_{t=1, \dots, T} (U(A_L) + U(A_R)) \quad (7)$$

where T is the number of possible binary tests, and A_L and A_R are the subset of training samples reaching the left node and the right node respectively. Each leaf node l in the constructed tree stores a single offset vector (x_l, y_l) . If A_l is the subset of training examples that arrives at the leaf node l , then (x_l, y_l) is

given by:

$$x_l = \frac{1}{|A_l|} \sum_{k \in A_l} x_k \quad y_l = \frac{1}{|A_l|} \sum_{k \in A_l} y_k \quad (8)$$

Similarly to [11], we use two stopping criteria for the construction of the tree: the maximum depth of the tree and the minimum number of samples at a node. If a node has this minimum number of samples, we add an additional constraint which checks the variance of the offset vectors with a specified threshold. This way we have a better estimate of the offset at the leaf nodes. At runtime, a patch is given to the tree and the estimated offset value at the leaf node is used to place the main object classifier for subsequent detection. For illustration, we show in Figure 3(b) the pixel locations (x, y) associated to the tests t_{μ_f} at each node in the tree from the root to different leaves. In these examples, we basically observe that the tree learns the shifts near the eye location.

4 Experiments

We first evaluate the performance of bounding box estimation for different patch size. We then compare our proposed scanning framework with standard scanning technique with respect to detection rate, false alarm rate and scanning speed on benchmark and FDDB face database.

4.1 Evaluation of bounding box (bbx) estimation

We evaluate the performance of bounding box estimation for different patch sizes (w_p, h_p) and for two different types of binary test t_f and t_{μ_f} . We first describe the training and testing face dataset and parameters set for training the decision tree. We obtain approximately 35,000 cropped face images (19x19) (faces are scaled and cropped with respect to eye location) from standard face database (BANCA, BIOID, Purdue, and XM2VTS). A subset of 15,000 face images are used for training, 10,000 are used for validation and the rest 10,000 are used for testing. The dataset which we used for this evaluation have well defined eye locations and can assume that the patches have good groundtruth offset values.

To build a decision tree we set the maximum depth, minimum number of samples and variance threshold. To obtain the best parameter of the tree cross validation approach could be used which will be done in our future work. Here we set the parameters based on our experience to quickly test our hypothesis. The depth of the tree is varied from 12 to 15 depending on the patch size, but kept the same for two types of test. Large patch size have fewer offset values to be estimated, therefore we use smaller depth size, while smaller patch size have many offset values which creates more training samples and requires larger depth for better estimation. The variance threshold is set to 0.1 and minimum number of samples to 10 for all our experiments. A smaller variance will force the




Fig. 4. Cumulative distribution of estimation error λ for patch sizes of 10x10, 14x4, and 17x17. The figure shows only the cumulative distribution for first few λ 's.

training samples to split if the offset values are too different. The total number of possible binary tests for t_f is $\frac{(w_p \times h_p)(w_p \times h_p - 1)}{2}$, which is large. Hence a fixed number of tests are evaluated (200 in our case) at every node, and for each test the pixel pairs are picked randomly. The total number of possible binary tests in the case of $t_{\mu f}$ is $w_p \times h_p$, as it compares a pixel value to the average value of the patch. Since the number of test evaluation is small, each node evaluates all the test and selects the best one based on (7).

Training of a decision tree proceeds by giving all the samples at the root node and recursively splitting the training samples using (7) until it reaches the maximum depth or the variance of the samples at a node reaches below a specified threshold value. At test time, a patch is passed through the tree, and the leaf node gives an offset estimate (\hat{x}, \hat{y}) . Since we want to measure how close the estimated offset is to the true offset we use squared L_2 norm to evaluate the estimation error:


$$\lambda = (\hat{x} - x)^2 + (\hat{y} - y)^2 \quad (9)$$

where (x, y) is the true offset value of the patch.

To inspect the distribution of error λ , we define $g(\lambda)$ as the number of test patches that have estimation error of λ , and the cumulative distribution of estimation error as $c(\lambda) = \sum_{j=0}^{\lambda} g(j)$. Figure 4 shows the cumulative distribution of estimation error for square patch sizes of 10, 14 and 17. We see that there is only a slight performance difference between the two types of test.

4.2 Evaluation on benchmark face database

We now compare the detection rate and speed of scanning with and without bbx estimation on CMU+MIT [13] and Fleuret [14] frontal face databases. These



(a) estimated locations (our approach)

(b) regular grid spacing

Fig. 5. A dot shows the top left corner of the subwindow given to the classifier. (a) estimated locations with our approach, (b) for regular grid spacing.

databases have a total of 375 images and 1085 faces of various size. A pyramid based scanning approach [3] is used to detect faces at different scales. Multiple detections are merged by averaging the detection within a certain radius which is a function of scale. The estimated eye coordinate of merged detection are compared with ground truth eye coordinates using Jesorsky measure [15], which is set to 0.6 for all our experiments. We describe first the different face classifiers used for our experiments, then show how the location on the grid gets modified with bounding box prediction. We then show how our approach is better than standard scanning technique irrespective of the classifier used.

Main classifier. Many different classifiers and features are available for face detection task. We choose Modified Census Transform (MCT) features as it has been shown to be robust to lighting variations and does not require any preprocessing [16]. Boosting is used to train the classifiers and the cascade architecture is used for speeding up detection process. Two different classifiers are trained with different sets of face training data. One classifier is trained without perturbing the face training data (C_{mct-np}) and the other with perturbed face training data (C_{mct-p}). The perturbation consists of shifting the face by one pixel in x and y directions. We also use the OpenCV face classifier (C_{opencv}) [17] to compare the performance (detection rate and false alarms) with our approach for different grid spacing. Figure 5 shows the locations at which the main classifier is placed with and without bbx prediction.

Detection rate vs grid spacing. Figure 6 shows the performance of detection rate with respect to grid spacing. We can see that classifier C_{mct-p} performs similar to C_{opencv} . Training without perturbation has lower detection rate as the grid spacing is increased which is obvious. In the figure `mct-p-bbx` and `mct-np-bbx`, represents the performance with bbx prediction. It can be seen clearly that our approach has higher detection rate for both the MCT based classifiers. It can also be seen that there is an equal amount of drift when different classifiers are used

with and without bbx prediction. This might be due to that our approach is independent of the classifier used. Figure 7 shows the detection rate for different patch sizes when the grid spacing is varied. It can be seen that smaller patches are able to perform better compared to larger patches when the grid spacing is larger which goes along with our hypothesis (Section 3.2).




Fig. 6. Detection rate vs grid spacing with and without bbx prediction.




Fig. 7. Detection rate vs grid spacing with bbx prediction for different patch sizes.

Detection rate vs speed We analyze the computational complexity empirically. For an image we consider the total time for the scan to finish with and without bounding box prediction. For the given grid spacing obviously our approach takes more time, but we need to see for a given detection rate whether we gain any improvement in speed. In Figure 8, we can see that we can achieve better detection rate for a fixed time, or better time (increase in speed) for a fixed

detection rate. We also see that C_{mct-p} has more detection rate than C_{mct-np} , but our approach performs better when using the same classifier.

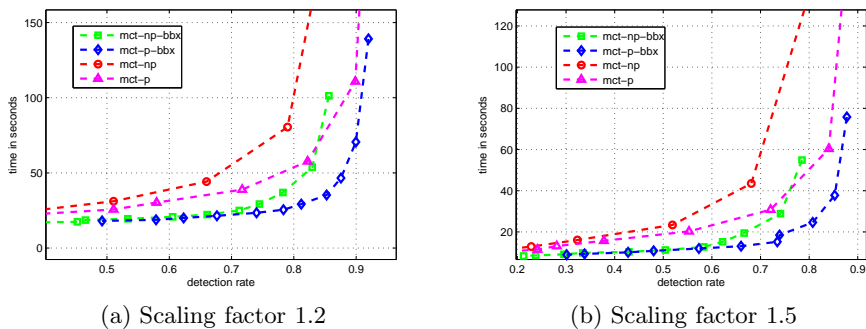


Fig. 8. Time in seconds to scan 375 images vs detection rate with and without bbx prediction.

Detection rate vs false alarm rate. In Figure 9 it can be seen that the false alarm rate for C_{opencv} and C_{mct-p} are similar, while C_{mct-np} has lower false alarm rate. This indicates that training without perturbation has lower false alarm at the price of detection rate. We can also see that the false alarm is lower when using bbx prediction for a fixed detection rate.

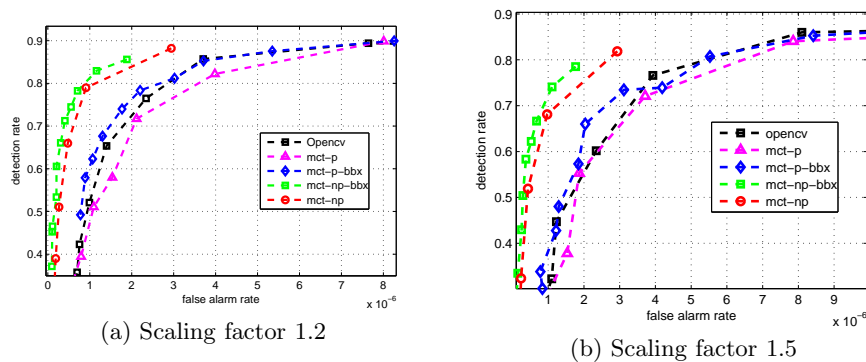


Fig. 9. Detection rate vs false alarm rate with and without bbx prediction.

4.3 Evaluation on Fddb face dataset

We also evaluate our approach on the Fddb dataset for EXP-2 [1]. It contains 2845 images with a total of 5171 faces. For this task we use a detector from

[18]. We use the “ROC.txt” and the evaluation code files from Fddb website (<http://vis-www.cs.umass.edu/fddb/results.html>) to generate the Figure 10(a) for comparison with [4] and [19]. We see that our detector is comparable to the state-of-art detectors. Here we show only discrete ROC curves. Figure 10(b) shows the ROC curves with our detector using the standard scanning technique and with our bbx prediction. The time taken to process 2845 images using standard scanning technique with grid spacing 2, 4, 6, and 8 are 480, 165, 105, and 86 seconds respectively, while for bbx prediction they are 562, 187, 119, and 86 seconds respectively. It is clear that as the grid spacing becomes smaller the bounding box prediction does not give any advantage, but when we look at Figure 10(b) for a given false positives (say 1000), bbx prediction with grid spacing 4 performs comparable to standard scanning technique with grid spacing 2, and also it is 3 times faster. Similarly, bbx with grid spacing 8 performs comparable to standard scanning technique with grid spacing 4, and its 2 times faster.

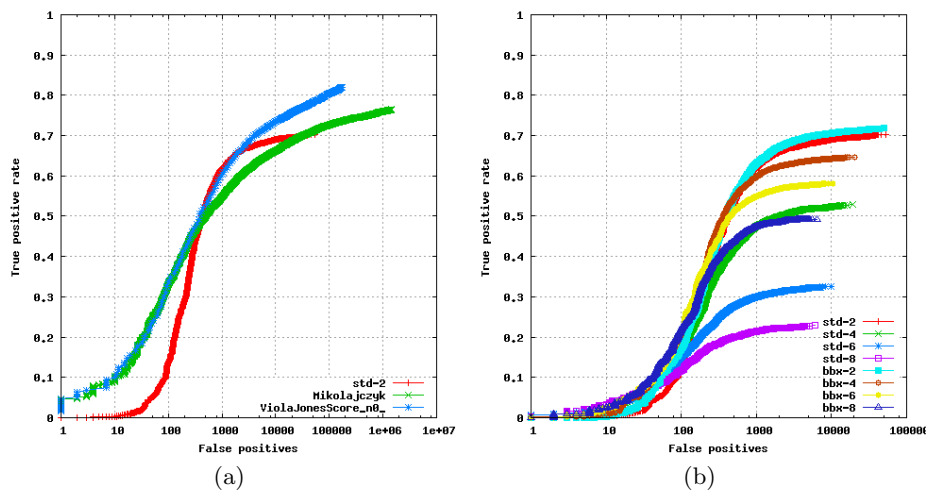


Fig. 10. Discrete ROC curves. (best viewed in color, x axis is plotted in log scale) (a) Comparison of our detector with [4] and [19]. (b) ROC curves for our detector with standard scanning (std-) and with bbx prediction (bbx-). The numbers 2,4,6, and 8 represent grid spacing used while scanning.

5 Conclusion and future work

We have presented a method to improve the detection rate and speed while increasing the grid spacing in a sliding window based scanning technique. The key idea is to predict the bounding box of the object with high performance (both in speed and accuracy) which is achieved by using a decision tree with simple binary test at each node. We have used benchmark face databases to

validate our approach and demonstrated that there can be an improvement in speed for a fixed detection rate, or better detection rate for a fixed time. The grid spacing could be chosen to reflect specific application need. We also showed the performance of our approach on the new larger FDDB dataset, with the common evaluation protocol, and we see that our hypothesis works on wide variety of images. What we observe is that, the pose and occlusion remains a challenge in face detection and it needs to be addressed. In our future work we plan to investigate if our approach can be easily extended to estimate the face pose. We also plan to use different set of features or a combination of features and context information to improve face detection in unconstrained environment.

Acknowledgments. The authors would like to thank the Swiss National Science Foundation, projects MultiModal Interaction and MultiMedia Data Mining (MULTI, 200020-122062) and Interactive Multimodal Information Management (IM2, 51NF40-111401) and the FP7 European MOBIO project (IST-214324) for their financial support.

References

1. Jain, V., Learned-Miller, E.: Fddb: A benchmark for face detection in unconstrained settings. Technical Report UM-CS-2010-009, University of Massachusetts, Amherst (2010)
2. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: IEEE Conference on Computer Vision and Pattern Recognition. Volume 1. (2005) 886–893
3. Rowley, H.A., Baluja, S., Kanade, T.: Neural network-based face detection. IEEE Transactions on Pattern Analysis and Machine Intelligence **20** (1998) 23–38
4. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: IEEE Conference on Computer Vision and Pattern Recognition. Volume 1. (2001) 511–518
5. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60** (2004) 91–110
6. Özuysal, M., Fua, P., Lepetit, V.: Fast keypoint recognition in ten lines of code. In: IEEE Conference on Computer Vision and Pattern Recognition, Minneapolis, USA (2007) 1–8
7. Beis, J., Lowe, D.: Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In: IEEE Conference on Computer Vision and Pattern Recognition. (1997) 1000–1006
8. Özuysal, M., Lepetit, V., Fua, P.: Pose estimation for category specific multi-view object localization. In: IEEE Conference on Computer Vision and Pattern Recognition, Miami, USA (2009) 778–785
9. Bileschi, S.M., Heisele, B.: Advances in component based face detection. In: IEEE International Workshop on Analysis and Modeling of Faces and Gestures. (2003)
10. Leibe, B., Leonardis, A., Schiele, B.: Combined object categorization and segmentation with an implicit shape model. In: Proceedings of the Workshop on Statistical Learning in Computer Vision, Prague, Czech Republic (2004)
11. Gall, J., Lempitsky, V.: Class-specific hough forests for object detection. In: IEEE Conference on Computer Vision and Pattern Recognition, Miami, USA (2009) 1–8

12. Zhou, S.K., Comaniciu, D.: Shape regression machine (2007)
13. Schneiderman, H., Kanade, T.: A statistical model for 3d object detection applied to faces and cars. In: IEEE Conference on Computer Vision and Pattern Recognition. (2000) 746–751
14. Fleuret, F.: Fast binary feature selection with conditional mutual information. *Journal of Machine Learning Research* **5** (2004) 1531–1555
15. Jesorsky, O., Kirchberg, K., Frischholz, R.: Robust face detection using the Hausdorff distance. In: International Conference on Audio- and Video-Based Biometric Person Authentication, Halmstad, Sweden (2001) 90–95
16. Froba, B., Ernst, A.: Face detection with the modified census transform. In: IEEE International Conference on Automatic Face and Gesture Recognition. (2004) 91–96
17. Lienhart, A.K., Pisarevsky, V.: Empirical analysis of detection cascades of boosted classifiers for rapid object detection. In: 25th Pattern Recognition Symposium. (2003) 297–304
18. Sauquet, T., Rodriguez, Y., Marcel, S.: Multiview Face Detection. *Idiap-RR Idiap-RR-49-2005, IDIAP* (2005)
19. Mikolajczyk, K., Schmid, C., Zisserman, A.: Human detection based on a probabilistic assembly of robust part detectors. In: European Conference on Computer Vision. (2004) 69–82