Incremental Training of a Detector Using Online Sparse Eigendecomposition

Sakrapee Paisitkriangkrai, Chunhua Shen, and Jian Zhang, *Senior Member, IEEE*

Abstract—The ability to efficiently and accurately detect objects plays a very crucial role for many computer vision tasks. Recently, offline object detectors have shown a tremendous success. However, one major drawback of offline techniques is that a complete set of training data has to be collected beforehand. In addition, once learned, an offline detector cannot make use of newly arriving data. To alleviate these drawbacks, online learning has been adopted with the following objectives: 1) the technique should be computationally and storage efficient; 2) the updated classifier must maintain its high classification accuracy. In this paper, we propose an effective and efficient framework for learning an adaptive online greedy sparse linear discriminant analysis model. Unlike many existing online boosting detectors, which usually apply exponential or logistic loss, our online algorithm makes use of linear discriminant analysis' learning criterion that not only aims to maximize the class-separation criterion but also incorporates the asymmetrical property of training data distributions. We provide a better alternative for online boosting algorithms in the context of training a visual object detector. We demonstrate the robustness and efficiency of our methods on handwritten digit and face data sets. Our results confirm that object detection tasks benefit significantly when trained in an online manner.

Index Terms—Asymmetry, feature selection, greedy sparse linear discriminant analysis (GSLDA), object detection, online linear discriminant analysis.

I. INTRODUCTION

REAL-TIME object detection plays an important role in many real-world vision applications. It is used as a preceding step in applications such as intelligent video surveillance, content based image retrieval, face and activity recognition. Object detection is a challenging problem due to large variations in visual appearances, object poses, illumination, camera

motion, etc. All these issues have made the problem very challenging from a machine vision perspective.

The literature on object detection is abundant. A thorough discussion on this topic can be found in several surveys: faces [1], [2], human and pedestrians [3], [4], eyes [5], vehicles [6], etc. In this paper, we review only the most relevant visual detection work, focusing on algorithms that operate directly on classification based visual object detection and incremental learning.

Object detection problems are often formulated as classification tasks where a sliding window technique is used to scan the entire image and locate objects of interest [7]–[10]. Viola and Jones [8] proposed an efficient detection algorithm based upon boosting algorithms [11]–[13] and a cascade classifier. Their detector is the first highly-accurate real-time face detector. They trained classifiers on data sets with a few thousand faces and a large number of nonfaces. During the training procedure, negative samples are gradually bootstrapped and added to the training set of boosting classifiers in the next stage. This method yields an extremely low false positive rate. However, a large number of faces and nonfaces are required in advance to cover different face appearances and poses. As a result, the computation cost and memory requirements for training an AdaBoost detector are extremely high. Viola and Jones spent weeks on training a detector with 6060 features (weak learners) on a face training set of 4916.

To speed up the training time bottleneck, a few approaches have been proposed. Pham and Cham [14] reduced the training time of weak learners by approximating decision stumps with class-conditional Gaussian distributions. Wu *et al.* [15] introduced a fast implementation of AdaBoost methods and proposed forward feature selection (FFS) for fast training. FFS ignores the reweighting step in boosting such that weak classifiers only need to be trained for once. Xiao *et al.* [16] applied distributed learning to learn their proposed dynamic cascade framework. They use over 30 desktop computers for parallel training. They managed to train a face detector on a training set with 500,000 positive samples and 10 billion negative samples in under 7 h. However, these techniques are not applicable to some real-world applications where a complete set of training samples is often not given in advance. Retraining the model each time new data arrive would increase the time complexity by the factor of N , where N is the number of newly arrived samples. Hence, developing an efficient adaptive object detector has become an urgent issue for many applications of object detection in diverse and changing environments. To alleviate this problem, online incremental learning algorithms have been proposed for this purpose.

Manuscript received June 30, 2009; revised January 22, 2010 and May 28, 2010; accepted June 02, 2010. Date of publication June 21, 2010; date of current version December 17, 2010. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Center of Excellence program. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Patrick Flynn.

S. Paisitkriangkrai is with the University of Adelaide, Adelaide, SA 5005, Australia (e-mail: paul.paisitkriangkrai@adelaide.edu.au).

J. Zhang is with the National ICT Australia (NICTA), Neville Roach Laboratory, Kensington, NSW 2052, Australia. They are also with the University of New South Wales, Sydney, NSW 2052, Australia (e-mail: jian.zhang@nicta.com.au).

C. Shen is with National ICT Australia (NICTA), Canberra Research Laboratory, Canberra, ACT 2601, Australia. He is also with the Australian National University, Canberra, ACT 0200, Australia (e-mail: chunhua.shen@nicta.com.au).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2010.2053548

Online learning was firstly introduced in the computational learning community. Since, boosting is one of the classifiers that have been successfully applied to many machine learning tasks, there has been considerable interest in applying boosting techniques on problems that require online learning. The first online version of boosting algorithms was proposed by Oza and Russel [17]. The algorithm works by minimizing the classification error while updating weak classifiers online. Grabner and Bischof [18] later applied online boosting to object detection and visual tracking. They proposed an online feature selection method, where a group of selectors is initialized randomly, each with its own feature pool. By interchanging weak learners based upon lowest classification error, the algorithm is able to capture the change in patterns induced by new samples. Huang *et al.* [19] proposed an incremental learning algorithm that adjusts a boosted classifier with domain-partitioning weak hypotheses to online samples. They showed that by incremental learning with few difficult unseen faces (e.g., faces with sun glasses or extreme illumination), the performance of online detectors can be significantly improved. Parag *et al.* [20] advocated an online boosting algorithm where parameters of weak classifiers are updated using weighted linear regressors to minimize the weighted least square error (LSE). In the context of pedestrian detection, Liu and Yu [21] introduced a gradient-based feature selection approach where parameters of weak classifiers are updated using gradient descent to minimize the weighted LSE. Nonetheless, most of these proposed techniques concentrated on the application of visual tracking or object classification with small training sets and few online data sets.¹ Hence, to date, it remains unclear whether there is any improvement in object detection by continuously updating the existing models with a sufficiently large training sample set. We will reveal this mystery in Section III-B.2.

Recently, Moghaddam *et al.* [22] presented a technique that combines the greedy approach with the efficient block matrix inverse formula. The proposed technique, termed greedy sparse linear discriminant analysis (GSLDA), speeds up the calculation time by 1000-fold compared with globally optimal solutions found by branch-and-bound search in the case of binary-classification problems. Paisitkriangkrai *et al.* [23] later applied the GSLDA algorithm to face detection and showed very convincing results. Their GSLDA face detector has shown to outperform AdaBoost based face detector due to the nature of the training data (the distribution of face and nonface samples is highly imbalanced). The objective of this work is to design an efficient incremental GSLDA algorithm that can accommodate new data efficiently while preserving a promising classification performance.

Unlike classical linear discriminant analysis (LDA) where a lot of online learning techniques have been designed and proposed [24]–[26], there are very few works on incremental learning for *sparse* LDA. One of the difficulties might be due to the fact that the sparse LDA problem is nonconvex and NP-hard. It is not straightforward to design an incremental solution for sparse LDA. In this work, we design an algorithm

that efficiently learns and updates the sparse LDA classifier. Our online sparse LDA classifier not only incorporates new data efficiently but also improves the classification accuracy as new data become available. In brief, we extend the work of [23] with an efficient online update scheme. Our method modifies the weight of linear discriminant functions to adapt to new data sets. This update process generalizes the weight of linear discriminant functions and results in accuracy improvements on real data and real applications.

The key contributions of this work can be summarized as follows.

- 1) We propose an efficient incremental greedy sparse LDA classifier for training an object detector in an incremental fashion. The online algorithm integrates the GSLDA based feature selection with our adaptation schemes for updating weights of linear discriminant functions and the linear classifier threshold. Our updating algorithm is very efficient. We neither replace weak learners nor throw away any weak learners during updating phase.
- 2) Our online GSLDA serves as a better (in terms of performance) alternative to the standard online boosting [17] for training detectors. To our knowledge, it is the first time to apply the online sparse linear discriminant analysis algorithm to object detection.
- 3) Finally, we have conducted extensive experiments on several data sets that have been used in the literature. Experimental results confirm that incremental learning with online samples is beneficial to the initial classifier. Our algorithm can efficiently update the classifier when new instances are inserted while achieving comparable classification accuracy to the batch algorithm.² Our findings indicate that online learning plays a crucial role in object detection, especially when the initial number of training samples is small. Note that when trained with few positive samples, the detector often underperforms since it fails to capture appearance variations of target objects. By applying our online technique, the classification performance can be further improved at the cost of minor increases in training time.

The rest of the paper is organized as follows. Section II begins by introducing the concept of LDA and GSLDA. We then propose our online GSLDA object detector. The results of numerous experiments are presented in Section III. We conclude the paper in Section IV.

II. ALGORITHMS

For ease of exposition, symbols and their denotations used in this paper are summarized in Table I. In this section, we begin by introducing the basic concept of classical LDA and GSLDA. We then propose our online greedy sparse LDA (OGSLDA).

A. Classical LDA

LDA deals with the problem of finding weights of linear discriminant functions. Let us assume that we have a set of training

¹For example, in [21], the authors trained the initial classifier with 366 positive samples and 2540 negative patches, and incrementally updated with 366 online positive samples and 2450 online negative patches.

²We use the terms “batch learning” and “offline learning” interchangeably in this paper.

patterns $\mathbf{x} = [x_1, x_2, \dots, x_M]^\top$ where each of which is assigned to one of two classes, C_1 and C_2 . We can find a weight vector $\mathbf{w} = [w_1, w_2, \dots, w_M]^\top$ and a threshold w_0 such that

$$\begin{aligned} \mathbf{w}^\top \mathbf{x} + w_0 &\geq 0 \quad (\mathbf{x} \in C_1) \\ \mathbf{w}^\top \mathbf{x} + w_0 &< 0 \quad (\mathbf{x} \in C_2). \end{aligned} \quad (1)$$

In general, we seek the vector $[w_0, w_1, w_2, \dots, w_M]$ that best satisfies (1). The data are said to be linearly separable if for all $\mathbf{x}$, (1) is satisfied.

An intuitive objective that one can take is to find a linear combination of variables that can separate two classes as much as possible. The computed linear combination reduces dimensions of samples to one dimension. The classical criterion proposed by Fisher is the ratio of between-class to within-class variances, which can be written as

$$J = \frac{\sum_{c \in C} \sum_{\mathbf{x} \in c} (\mathbf{w}^\top \mathbf{m}_c - \mathbf{w}^\top \bar{\mathbf{m}})^2}{\sum_{c \in C} \sum_{\mathbf{x} \in c} (\mathbf{w}^\top \mathbf{x} - \mathbf{w}^\top \mathbf{m}_c)^2} = \frac{\mathbf{w}^\top S_b \mathbf{w}}{\mathbf{w}^\top S_w \mathbf{w}} \quad (2)$$

$$S_b = \sum_{c \in C} N_c (\mathbf{m}_c - \bar{\mathbf{m}})(\mathbf{m}_c - \bar{\mathbf{m}})^\top \quad (3)$$

$$S_w = \sum_{c \in C} \sum_{\mathbf{x} \in c} (\mathbf{x} - \mathbf{m}_c)(\mathbf{x} - \mathbf{m}_c)^\top. \quad (4)$$

Here, $\mathbf{m}_c$ is the mean of class c , $\bar{\mathbf{m}}$ is the global mean, N_c is the number of instances in class c , S_b and S_w are the so-called between-class and within-class scatter matrices. The numerator of (2) denotes the distance between projected means and the denominator denotes the variance of the pooled data. We want to find linear projections, $\mathbf{w}$, that maximize J , the distance between means of two classes while minimizing the variance within each class. The solution can be obtained by generalized eigendecomposition (GEVD). The optimal solution $\mathbf{w}$ is the eigenvector corresponding to the maximal eigenvalue and can be expressed as [27]

$$\mathbf{w} \propto S_w^{-1}(\mathbf{m}_1 - \mathbf{m}_2). \quad (5)$$

If we further assume that the data are normally distributed and that distributions in the original space have identical covariance matrices, an optimal threshold, w_0 , can be calculated from

$$w_0 = -\frac{1}{2}(\mathbf{m}_1 + \mathbf{m}_2)^\top S_w^{-1}(\mathbf{m}_1 - \mathbf{m}_2) - \log \left(\frac{\Pr(C_2)}{\Pr(C_1)} \right). \quad (6)$$

Here, $\Pr(C_1)$ and $\Pr(C_2)$ are priori probabilities of class C_1 and C_2 , respectively. This threshold can be interpreted as the mid-point between two projected means, shifted by the log of the ratio between priori probabilities of two classes.

B. GSLDA

In this section, we briefly present the offline implementation of the greedy sparse LDA algorithm [22], [28]. The sparse version of classical LDA is to solve

$$\begin{aligned} &\underset{\mathbf{w}}{\text{maximize}} && \frac{\mathbf{w}^\top S_b \mathbf{w}}{\mathbf{w}^\top S_w \mathbf{w}} \\ &\text{subject to} && \text{Card}(\mathbf{w}) = k \end{aligned} \quad (7)$$

where $\text{card}(\mathbf{w}) = k$ is an additional sparsity constraint, $\text{card}(\cdot)$ denotes the ℓ_0 norm, which counts the nonzero components of the vector. k is an integer preset by a user. Due to this additional sparsity constraint, the problem is nonconvex and NP-hard. In [28], Moghaddam *et al.* presented a technique to compute optimal sparse linear discriminants using branch-and-bound approach. Nevertheless, finding the exact global optimal solutions for high-dimensional data is infeasible. The algorithm was extended in [22] with new sparsity bounds and efficient block matrix inverse techniques to speed up the computation time by 1000-fold. The technique works by sequentially adding the new variable which yields the maximum eigenvalue (forward selection) until the number of nonzero components, $\text{card}(\mathbf{w})$, is equal to the integer set by the user.

In [23], Paisitkriangkrai *et al.* learned an object detector using the GSLDA algorithm. The training procedure is described in Algorithm 1. First, the set of selected features is initialized to an empty set. The algorithm then trains all weak learners and store their results into a lookup table (lines 1, 2). At every round, the output of each weak learner is examined and the weak learner that most separates the two classes is sequentially added to the list (lines 3, 4). Mathematically, Algorithm 1 sequentially selects the weak learner whose output yields the maximal eigenvalue. Weak learners are added until the target learning goal is met. The authors of [23] use an asymmetric node learning goal to build a cascade object detector using GSLDA.

Algorithm 1 The training procedure for building an offline GSLDA object detector

Input:

- A positive training set and a negative training set;
- A set of features $\{f_i; i \in [1, M]\}$;

Output:

- A set of selected weak learners $\{h_i; i \in [1, T]\}$ that best separates the training set;

1 foreach feature do

2 Train a weak learner (e.g., a decision stump parameterized by a threshold θ that results in the smallest classification error) on the training set;

3 while the target goal is not met do

4 Add the best weak learner (e.g., decision stump) that yields the maximum class separation to the set of selected weak learners;

C. Incremental Learning of GSLDA Classifiers

The major challenge of GSLDA object detectors in real-world applications is that a complete set of training samples is often not given in advance. As new data arrive, the between-class and within-class scatter matrices, S_b and S_w , will change accordingly. In offline GSLDA, the value of both matrices would have to be recomputed from scratch. However, this approach is unacceptable due to its heavy computation and storage requirements. First, the cost of computing both matrices grows with

TABLE I
NOTATION

Notation	Description
C_1, C_2	Class 1 (positive class), class 2 (negative class)
N	The number of training samples in each classifier
N_1, N_2	The number of training samples in the first and second class, respectively
M	The size of feature sets (for decision stumps, this is also equal to the number of weak learners)
T	The number of features to be selected
$\mathbf{X}$	Data matrix
$\mathbf{x}$	The new instance being inserted
$\bar{\mathbf{m}}$	The global mean of training samples
$\mathbf{m}_1, \mathbf{m}_2$	The mean (centroid) of the first and second class, respectively
Σ_1, Σ_2	The covariance of the first and second class
μ_1, μ_2	The projected mean of the first and second class
σ_1, σ_2	The projected covariance of the first and second class
$S_b, \tilde{S}_b$	Between-class scatter matrix and its updated value after the new instance $\mathbf{x}$ has been inserted
$S_w, \tilde{S}_w$	Within-class scatter matrix and its updated value
$\mathbf{w}$	Weights of linear discriminant functions (also referred to as weak learners' coefficients in the context)
w_0	The linear classifier threshold

the number of training samples. As a result, the algorithm will run slower and slower as time progresses. Second, the batch approach uses the entire set of training data for each update. In other words, the previous training data need to be stored for the retraining purpose.

In order to overcome these drawbacks, we propose an on-line learning algorithm, termed online greedy sparse LDA (OGSLDA). The OGSLDA algorithm consists of two phases: the initial offline learning phase and the incremental learning phase. The training procedure in the initial phase is similar to the algorithm outlined in Algorithm 1. Here, we assume that the number of training samples available initially is adequate and well represents the true density. In the second phase, the learned covariance matrices are updated in an incremental manner.

It is important to point out that a number of incremental LDA-like approximation algorithms have been proposed in the literature [25], [29]. Ye *et al.* [25] proposed an efficient LDA-based incremental dimension reduction algorithm which applied QR decomposition and QR-updating techniques for memory and computation efficiency. Kim *et al.* [29] proposed an incremental LDA by applying the concept of the sufficient spanning set approximation in each update step. However, we did not find any of existing LDA-like algorithms appropriate to our problems. Based upon our preliminary experiments, the projection matrix determined in subspace often gives worse discriminant power than that from full space. This might be due to their dimension reduction algorithms which reduced the between-class and within-class scatter matrices to a much smaller size. Our on-line GSLDA is guaranteed to build the same between-class and within-class scatter matrices as batch GSLDA given the same training data. The reason why we need not worry about large dimensions in our algorithms is because applying sparse LDA in our initial phase already reduces the number of dimensions that we have to deal with. Hence, given the same set of features, the accuracy of our online GSLDA is better than the existing incremental LDA-like approximation algorithms. The only expensive computation in our algorithms is eigenanalysis. In order to avoid

the high computation complexity of continuously solving generalized eigendecomposition, we apply the efficient matrix inversion updating techniques based upon the inverse Sherman–Morrison formula. As a result, our incremental algorithm is very robust and efficient.

In this section, we first introduce an efficient method that incrementally updates both the within-class and between-class scatter matrices as new observations arrive. Then, we describe an approach to update the classifier threshold. Finally, we analyze the storage and training time complexity of our proposed approaches.

1) *Incremental Update of Between-Class and Within-Class Matrices:* Since, GSLDA assumes a Gaussian distribution, the incremental update of class mean and class covariance can be computed very quickly. The techniques used to update both matrices can be easily derived. The procedure proceeds in three steps:

- Step 1) updating the between-class scatter matrix S_b ;
- Step 2) updating the within-class scatter matrix S_w ;
- Step 3) updating inverse of the within-class scatter matrix S_w^{-1} .

Updating the between-class scatter matrix: The definition of the between-class scatter matrix is given in (3). For two classes (C_1 and C_2), S_b can be simplified to

$$S_b = \frac{N_1 N_2}{N} (\mathbf{m}_1 - \mathbf{m}_2)(\mathbf{m}_1 - \mathbf{m}_2)^\top. \quad (8)$$

The expression can be interpreted as the scatter of class 1 with respect to the scatter of class 2. Let $\mathbf{x}$ be a new instance being inserted. The updated $\tilde{\mathbf{m}}_1$ and $\tilde{\mathbf{m}}_2$ can be calculated from

$$\tilde{\mathbf{m}}_1 = \begin{cases} \mathbf{m}_1 + \frac{\mathbf{x} - \mathbf{m}_1}{N_1 + 1}, & \text{if } \mathbf{x} \in C_1 \\ \mathbf{m}_1, & \text{otherwise} \end{cases} \quad \tilde{\mathbf{m}}_2 = \begin{cases} \mathbf{m}_2, & \text{if } \mathbf{x} \in C_1 \\ \mathbf{m}_2 + \frac{\mathbf{x} - \mathbf{m}_2}{N_2 + 1}, & \text{otherwise.} \end{cases} \quad (9)$$

Updating the within-class scatter matrix: The covariance of random vectors $\mathbf{X}$ is a square matrix Σ where $\Sigma = \mathbb{E}[(\mathbf{X} - \mathbb{E}[\mathbf{X}])(\mathbf{X} - \mathbb{E}[\mathbf{X}])^\top]$. Given the new instance $\mathbf{x}$, the updated covariance matrix is given by

$$\tilde{\Sigma} = ([\mathbf{X}, \mathbf{x}] - \tilde{\mathbf{m}}\mathbf{1}^\top)([\mathbf{X}, \mathbf{x}] - \tilde{\mathbf{m}}\mathbf{1}^\top)^\top. \quad (10)$$

Here, $\tilde{\mathbf{m}}$ is an updated mean after new instance has been inserted and $\mathbf{1}$ is a column vector with each entry being 1. Its dimensionality should be clear from the context. Note that in (10), we leave out the constant term since it makes no difference to the final solution

$$\begin{aligned} [\mathbf{X}, \mathbf{x}] - \tilde{\mathbf{m}}\mathbf{1}^\top &= [\mathbf{X}, \mathbf{x}] - \mathbf{m}_1\mathbf{1}^\top + \mathbf{m}_1\mathbf{1}^\top - \tilde{\mathbf{m}}\mathbf{1}^\top \\ &= [\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{x} - \mathbf{m}_1] - (\tilde{\mathbf{m}} - \mathbf{m}_1)\mathbf{1}^\top. \end{aligned}$$

Substitute the previous expression into (10) and let $\mathbf{u} = \mathbf{x} - \mathbf{m}_1$ and $\mathbf{v} = \tilde{\mathbf{m}} - \mathbf{m}_1$

$$\begin{aligned} \tilde{\Sigma} &= ([\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}] - \mathbf{v}\mathbf{1}^\top)([\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}] - \mathbf{v}\mathbf{1}^\top)^\top \\ &= ([\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}][\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}]^\top - [\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}](\mathbf{v}\mathbf{1}^\top)^\top \\ &\quad - (\mathbf{v}\mathbf{1}^\top)[\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}]^\top + (\mathbf{v}\mathbf{1}^\top)(\mathbf{v}\mathbf{1}^\top)^\top) \\ &= \Sigma + \mathbf{u}\mathbf{u}^\top - [\mathbf{X} - \mathbf{m}_1\mathbf{1}^\top, \mathbf{u}]\mathbf{v}^\top \end{aligned}$$

$$\begin{aligned}
& -\mathbf{v}([\mathbf{X}-\mathbf{m}\mathbf{1}^\top, \mathbf{u}\mathbf{1}^\top]^\top + (N+1)\mathbf{v}\mathbf{v}^\top \\
& = \Sigma + \mathbf{u}\mathbf{u}^\top - (N\mathbf{m} - N\mathbf{m} + \mathbf{u})\mathbf{v}^\top \\
& \quad - \mathbf{v}(N\mathbf{m} - N\mathbf{m} + \mathbf{u})^\top + (N+1)\mathbf{v}\mathbf{v}^\top \\
& = \Sigma + \mathbf{u}\mathbf{u}^\top - \mathbf{u}\mathbf{v}^\top - \mathbf{v}\mathbf{u}^\top + (N+1)\mathbf{v}\mathbf{v}^\top \\
& = \Sigma + (\mathbf{u}-\mathbf{v})(\mathbf{u}-\mathbf{v})^\top + N\mathbf{v}\mathbf{v}^\top \\
& = \Sigma + (\mathbf{x}-\tilde{\mathbf{m}})(\mathbf{x}-\tilde{\mathbf{m}})^\top + N(\tilde{\mathbf{m}}-\mathbf{m})(\tilde{\mathbf{m}}-\mathbf{m})^\top. \tag{11}
\end{aligned}$$

Note that $\mathbf{X}\mathbf{1} = \mathbf{m}\mathbf{1}^\top\mathbf{1} = N\mathbf{m}$. Next, we consider updating within-class scatter matrix. Let $\mathbf{x}$ be a new instance being inserted. The updated matrix, $\tilde{S}_w$, can be calculated from

$$\tilde{S}_w = \begin{cases} \tilde{S}_1 + \Sigma_2, & \text{if } \mathbf{x} \in C_1 \\ \Sigma_1 + \tilde{S}_2, & \text{otherwise.} \end{cases} \tag{12}$$

Updating inverse of the within-class scatter matrix: As mentioned in [22] that the computational complexity of 2-class GSLDA relies heavily on calculating the within-class scatter matrix inversion. In order to efficiently update the matrix inversion, we make use of the technique called balanced incomplete factorization which is based upon the inverse Sherman–Morrison formula proposed by Bru *et al.* [30]. Let Σ be the square matrix of size $M \times M$ which can be written as

$$\Sigma = \Sigma_0 + \mathbf{p}_1\mathbf{q}_1^\top + \mathbf{p}_2\mathbf{q}_2^\top. \tag{13}$$

Here, we assume that Σ_0 is nonsingular and $\mathbf{p}_1, \mathbf{p}_2, \mathbf{q}_1, \mathbf{q}_2 \in \mathbb{R}^M$. The inverse of Σ is given by

$$\Sigma^{-1} = \Sigma_0^{-1} - \Sigma_0^{-1}UD^{-1}V^\top\Sigma_0^{-1} \tag{14}$$

where $D^{-1} = \begin{bmatrix} r_1^{-1} & 0 \\ 0 & r_2^{-1} \end{bmatrix}$, $U = [\mathbf{p}_1 \ \mathbf{p}_2 - (\mathbf{q}_1^\top\Sigma_0^{-1}\mathbf{p}_2/r_1)\mathbf{p}_1]$
 $V = [\mathbf{q}_1 \ \mathbf{q}_2 - (\mathbf{q}_2^\top\Sigma_0^{-1}\mathbf{p}_1/r_1)\mathbf{q}_1]$, $r_1 = 1 + \mathbf{q}_1^\top\Sigma_0^{-1}\mathbf{p}_1$

$$r_2 = 1 + \left(\mathbf{q}_2 - \frac{\mathbf{q}_2^\top\Sigma_0^{-1}\mathbf{p}_1}{r_1}\mathbf{q}_1 \right)^\top \Sigma_0^{-1}\mathbf{p}_2.$$

The updated inverse of the within-class scatter matrix can be written as

$$\tilde{S}_w^{-1} = S_w^{-1} - S_w^{-1}UD^{-1}VS_w^{-1} \tag{15}$$

where $\mathbf{p}_1 = \mathbf{q}_1 = \mathbf{x} - \tilde{\mathbf{m}}$, $\mathbf{p}_2 = N_c(\tilde{\mathbf{m}} - \mathbf{m})$ and $\mathbf{q}_2 = \tilde{\mathbf{m}} - \mathbf{m}$ (from (11) and (13)).

2) *Updating Weak Learners' Coefficients and Thresholds:* Given the updated within-class matrix, $\tilde{S}_w^{-1}$, and between-class matrix, $\tilde{S}_b$, updated weights of linear discriminant functions can now be calculated from matrix-vector multiplication using (5). To complete the linear classifier, the threshold w_0 has to be obtained. Three criteria can be adopted. The first criterion is to apply the optimal Bayesian classifier in the projected space. In other words, the selected threshold should be the value in which the 1-D distribution function of two classes in the projected line

is equal. The mean and variance in the transformed space can be calculated as

$$\mu_c = \mathbf{w}^\top \mathbf{m}_c, \quad \sigma_c = \mathbf{w}^\top \Sigma_c \mathbf{w}. \tag{16}$$

If we let $X_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $X_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$. The optimal threshold is calculated as the point in which the 1-D density function of two classes is equal. Let $\log \Pr(x_1) = \log \Pr(x_2)$. After some algebraic expansions and simplifications, we can write the expression in the second-order polynomial

$$ax^2 + bx + c = 0$$

where $a = -(1/2\sigma_1^2) + (1/2\sigma_2^2)$, $b = (\mu_1/\sigma_1^2) - (\mu_2/\sigma_2^2)$ and $c = (\mu_2^2/2\sigma_2^2) + \log(\sigma_2) - (\mu_1^2/2\sigma_1^2) + \log(\sigma_1)$. The quadratics have two roots

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

In our implementations, we choose the threshold, w_0 , to be the value between the two class means

$$w_0 = x \quad \text{where} \quad \mu_1 < x < \mu_2. \tag{17}$$

The second criterion is to choose the threshold which yields a high detection rate with a moderate false alarm rate. This asymmetric criterion is often adopted in the cascade classifier framework [8]. Let $\phi(Z) = (1/\sqrt{2\pi}) \int_{-\infty}^Z \exp(-(1/2)u^2)du$ be the cumulative distribution function (CDF) of the standard normal random variable Z . If $X \sim \mathcal{N}(\mu_1, \sigma_1^2)$, the CDF of X is $\phi(Z)$ where $Z = (X - \mu_1)/\sigma_1$. Let the miss rate by p , the threshold which yields $1 - p$ detection rate can be calculated as

$$w_0 = \mu_1 + Z\sigma_1 = \mu_1 + \phi^{-1}(p)\sigma_1. \tag{18}$$

The last criterion is to set the threshold to be the projected mean of the negative class. This threshold helps us ensure the target asymmetric learning goal (moderate (50%) false positive rate with high detection rate). The threshold for the last criterion is

$$w_0 = \mu_2. \tag{19}$$

The previous three threshold updating rules might look oversimple. However, in [31], a few numerical simulations were performed on multidimensional normally distributed classes and real-life data taken from UCI machine learning repository. It is reported that selecting threshold using the simple approach as (17) often leads to smaller classification error than the traditional Fisher's approach (6).

Most online boosting algorithms modify parameters of weak learners to adapt to new data. For example, Grabner and Bischof updated parameters of weak learners using Kalman filtering [18]; Parag *et al.* updated the parameter using linear regression [20]; Liu and Yu updated the parameter using gradient descent [21], etc. We have found that extreme care has to be taken when we consider updating weak learners' parameters for applications of object detection. To demonstrate this, we generate an artificial asymmetric data set similar to the one used in [32].

Fig. 1. Toy data sets. $\times$ and $\circ$ represent positive and negative samples, respectively. Top row: no update. The parameters of weak learners do not get updated. Middle row: Gaussian model. The linear classifier threshold is calculated from updated mean and variance (using (17)). Bottom row: LSE model. The linear classifier threshold is updated using linear regression. The leftmost column shows the classifier thresholds on the initial training set (50 positive and 50 negative training points). The middle column shows the classifier thresholds with 25 new positive and 25 new negative points inserted. The rightmost column shows the thresholds with 50 new positive and 50 new negative inserted. Due to the asymmetry of data distributions, updating the parameters of weak learners could result in performance deterioration.

We then learn two different incremental linear weak classifiers with different parameter updating schemes:

- 1) incrementally update the model based upon a Gaussian distribution similar to [18];
- 2) incrementally update linear coefficients and intercept to minimize the LSE using linear regression similar to [20] (here, we assume uniform sample weights).

In this experiment, each weak learner represents a linear function with different coefficients (slopes). Each weak learner has one updatable parameter, i.e., linear classifier threshold (intercept). We apply the GSLDA algorithm and select the weak learner with minimal classification error. Based upon the selected weak learner, we continuously insert new samples and update the linear classifier threshold. Fig. 1 plots nine different linear classifier thresholds. The top row shows the linear classifier with no parameter updating. Middle row shows the linear classifier with the Gaussian updating rule. Bottom row shows the linear classifier using the linear regression algorithm. The first column shows different classifier thresholds on the initial training set. The middle and last columns show updated classifier thresholds with new data being inserted. We found that the top two classifier thresholds (no update and Gaussian) perform very similarly. LSE seems to perform worse when more new data are inserted. The reason may be attributed to the asymmetry of the data. When data are linearly separable, we can see that the regressor works very well. Based upon our results, we believe that parameter updating algorithms could significantly weaken the performance of weak learners if not applied properly. Hence, in this work, we decide not to update parameters of weak learners. By not updating weak learners' parameters, we can further speed up the computation.

The online GSLDA framework is summarized in Algorithm 2. Note that here we only use forward search of the original GSLDA algorithm of [22] and [28]. In [23], we have shown

that forward selection plus backward elimination only *slightly* improves the detection performance but with extra computation.

Algorithm 2 The Online GSLDA Algorithm

Given:

- The initial set of weak learners $\{h_i; i \in [1, T]\}$ trained using offline GSLDA on small initial data;

Input:

- New training datum I and its corresponding class label $y \in \{1, 2\}$;
- The current between-class covariance matrix, S_b ;
- The inverse of within-class covariance matrix, S_w^{-1} ;

Output:

- The updated between-class covariance matrix, $\tilde{S}_b$;
- The updated inverse of within-class covariance matrix, $\tilde{S}_w^{-1}$;
- The updated weak learners' coefficients, $\tilde{\mathbf{w}}$;
- The classifier threshold, $\tilde{w}_0$

1 Classify the new datum I using the given weak learners, $\mathbf{x} = [h_1(I), h_2(I), \dots, h_T(I)]$;

2 Update S_b with $\mathbf{x}$ using (8) and (9);

3 Update S_w^{-1} using (15);

4 Recalculate weak learners' coefficients, $\mathbf{w}$, using (5);

5 Update classifier threshold, w_0 , based upon node learning goal [(17) for minimal classification error, $\min((18), (19))$] for asymmetric node learning goal (see Section II-C.2);

Fig. 2. Comparison of (a) training and (b) test error rates between GSLDA and boosting approaches on USPS digits data sets [33].

3) *Computational Complexity of Incremental Learning:* Because the initial training of online GSLDA is the same as of-fine GSLDA, we briefly explain the time complexity of GSLDA [23]. Let us assume we choose decision stumps as our weak learners. Let the number of training samples be N . Finding an optimal threshold of each feature needs $O(N \log N)$.³ Assume that the size of feature sets is M . The time complexity for training weak learners is $O(MN \log N)$. During GSLDA learning, we need to find mean $O(N)$, variance $O(N)$ and correlation $O(T^2)$ for each feature. Since, we have M features and the number of weak learners to be selected is T , the total time complexity for offline GSLDA is $O(MN \log N + MNT + MT^3)$.

Given the selected set of weak learners, the time complexity of online GSLDA when a new training image, I , is inserted can be calculated as follows. Because the number of weak learners is T , the total time complexity to calculate a vector of weak classifiers' outputs, $\mathbf{x}$, in Step 1) is $O(T)$. It also takes $O(T)$ to update the class mean in Step 2). At Step 3), calculating U , V , r_1 , r_2 take $O(T^2)$. In this step, the order in which we calculate the matrix-matrix multiplication affects the overall efficiency. Since we are dealing with a small matrix chain multiplication, it is possible to go through each possible order and pick the most efficient one. For (14), we perform matrix-matrix multiplication in the following order $(((\Sigma_0^{-1}U)D^{-1})(V^T \Sigma_0^{-1}))$. The number of operations required to compute $(\Sigma_0^{-1}U)$ is $O(T \times T \times 2)$, $((\Sigma_0^{-1}U)D^{-1})$ is $O(T \times 2 \times 2)$, $(V^T \Sigma_0^{-1})$ is $O(2 \times T \times T)$ and $(((\Sigma_0^{-1}U)D^{-1})(V^T \Sigma_0^{-1}))$ is $O(T \times 2 \times T)$. Hence, the complexity of updating matrix inversion is still in the order of $O(T^2)$. Because the size of within-class matrices is $T \times T$, the matrix-vector multiplication in Step 4) takes $O(T^2)$. Updating the classifier threshold in Step 5) takes $O(T^2)$ for the first criterion (First, we find the projected mean and covariance, $O(T)$ and $O(T^2 + T)$, respectively. Then, we solve a closed-form second-degree polynomial). The second criterion in Step 5) takes $O(T^2)$ (Again, the time complexity of projected mean and covariance is $O(T)$ and $O(T^2 + T)$). The third criterion in Step 5) takes $O(T)$ (Here, we only have to calculate

³One usually sorts the 1-D feature using Quicksort, which has a complexity $O(N \log N)$.

the dot product of two vectors). Hence, the time complexity of Step 5) is at most $O(T^2)$. Therefore, the total time complexity for online GSLDA with the insertion of a new training image is at most $O(\underbrace{N_0 M \log N_0 + N_0 MT}_{\text{Offline}} + \underbrace{MT^3 + T^2}_{\text{Online}})$. Here,

N_0 is the number of initial training samples which assumed to be small. Note that the speed-up of online GSLDA over batch GSLDA is noticeable, i.e., $\underbrace{O(NT^2)}_{\text{Online}} \ll \underbrace{O(N^2 \log N)}_{\text{Batch}}$, when more instances are inserted into the training set ($N \gg N_0$).

In terms of memory usage, between-class scatter matrix takes up $O(2T)$. The inverse of within-class scatter matrix occupies $O(T^2)$. For the first and second criteria in Step 5), we also need to keep the covariance matrices of Σ_1 and Σ_2 which takes up $O(2T^2)$. Hence, the extra memory requirements for online GSLDA are at most $O(3T^2 + 2T)$. Given that the selected number of weak classifiers in each cascade layer is often small ($T < 200$), the time and memory complexity of online GSLDA is almost negligible.

III. EXPERIMENTS

A. USPS Digits Classification

In our first experiment, we use 16×16 pixels USPS digits "3" and "5" data sets. Data sets consist of 406 training instances and 418 test instances for the digit "3," 361 training instances and 355 test instances for the digit "5." We use the raw intensity value as our features. Hence, the total number of features is 256. We first compare a strong classifier learned using of-fine GSLDA [22] with three offline boosting algorithms: AdaBoost [11], Asymmetric Boosting [32] and FloatBoost [34]. For [32], we set the asymmetric parameter to be 1.1 based upon cross-validation. The performance is measured by classification error rates on both training and test sets. Results are shown in Fig. 2 and following observations can be made. GSLDA outperforms boosting approaches when the selected number of features is small. However, when the number of features increases, we observe the performance of all classifiers to be very similar.

In our next experiment, we compare offline GSLDA against online GSLDA. For online learning, we randomly select

Fig. 3. Top: classification error rates by offline GSLDA and online GSLDA on 16×16 pixels USPS digits data sets [33]. The number of initial training data for online GSLDA is (a) 30%, (b) 50%, and (c) 70% of the available training data. All experiments, except batch GSLDA (trained with full training sets), are run 10 times. The mean of errors are plotted. Bottom: classification error rates by online GSLDA and online boosting [17]. The number of initial training data is (d) 30%, (e) 50%, and (f) 70% of the available training data. All experiments are run ten times.

30/50/70 percent training samples as the training set. Incremental updating is performed with the remaining training instances being inserted one at a time. We use decision stumps as weak learners for both classifiers. All experiments, except batch GSLDA (trained with full training sets), are run 10 times. The mean of classification errors are plotted. Fig. 3(a)-(c) shows the achieved classification error rates by batch GSLDA and online GSLDA. The horizontal axis shows the cardinality of feature coefficients, i.e., the number of weak classifiers, and the vertical axis indicates the classification error rate on test data. We observe a trend that the error rate decreases when we train with more training instances. It is important to point out that in this experiment the error rate of online GSLDA is very close to that by batch GSLDA. We also train offline GSLDA classifiers with 30%, 50%, and 70% training data. We observe an increase in error rates of GSLDA (30% training data) when the number of dimensions increases. This is expected since it is common for a classifier to overfit with large dimensions and a small sample size.

Next, we compare the performance of online GSLDA with online boosting proposed in [17]. For each weak classifier, we build a model by estimating the univariate normal distribution with weighted mean and variance for digits “3” and “5.” We update the weak classifier by incrementally updating the mean and variance using weighted version of (9) and (10). Results are shown in Fig. 3(d)–(f). The test error of online boosting decreases as the initial number of training samples increases. We

observe the performance of online boosting to be remarkably worse than the performance of online GSLDA.

Fig. 4(a) and (c) shows the achieved classification error rates by batch GSLDA and online GSLDA with 25 and 100 dimensions (features). The horizontal axis shows the portion of training data instances and the vertical axis indicates the classification error rate. We observe a trend that the error rate decreases when more and more training data instances are involved. Online GSLDA not only performs well on given data sets but it is also very efficient. We compare the computation cost between batch GSLDA and incremental GSLDA in Fig. 4(b) and (d). As shown, the execution time of online GSLDA is significantly smaller than that of batch GSLDA as the number of training samples grows.

B. Frontal Face Detection

Due to its efficiency, Haar-like rectangle features [8] have become a popular choice as image features in the context of face detection. Similar to the work in [8], the weak learning algorithm known as decision stumps and Haar-like rectangle features are used here due to their simplicity and efficiency. Following experiments compare the offline GSLDA and online GSLDA learning algorithm.

1) *Performances on Single-Node Classifiers:* We conduct two experiments in this section. The first experiment compares a single strong classifier learned using AdaBoost [8], AsymBoost [32], offline GSLDA [23] and our proposed online

Fig. 4. Comparison of classification error rates and computation cost between online GSLDA and batch GSLDA on 16×16 pixels USPS digits data sets [33]. We set the number of features to (a) and (b) 25, (c) and (d) 100.

GSLDA. Data sets consist of 1000 mirrored face examples (Fig. 7) and 10,000 bootstrapped nonface examples. Faces were cropped and rescaled to 24×24 pixels. For nonface examples, we initially select 1000 random nonface patches from nonface images. The other 9000 nonface patches are added to the initial pool of training data by bootstrapping.⁴

We train three offline face detectors using AdaBoost, AsymBoost and GSLDA. Each classifier consists of 200 weak classifiers. Three classifiers are tested on challenged face videos, David Ross indoor data sets and trellis data sets,⁵ which are publicly available on the internet. Both videos contain large lighting variation, cast shadows, unknown camera motion, and tilted faces with in-plane and out-of-plane rotation. The first video contains 761 frames of a person moving from a dark to a bright area. Since, the first few video frames has a very low contrast (almost impossible to see faces), we ignore the first 100 frames. The second video contains 501 frames of a person moving underneath a trellis with large illumination changes and cast shadows.

In this experiment, we use the scanning window technique to locate faces. We set the scaling factor to 1.2 and window

⁴We incrementally construct new nonface samples using a trained classifier of [8].

⁵<http://www.cs.toronto.edu/~dross/ivt/>.

shifting step to 1 pixel. The patch with highest classification score is classified as targets. In other words, there is only one selected face in each frame. The criterion similar to the one used in PASCAL VOC Challenge [35] is adopted here. Detections are considered true or false positives based upon the area of overlap with ground truth bounding boxes. To be considered a correct detection, the area of overlap between the predicted bounding box, B_p , and ground truth bounding box, B_{gt} , must exceed 50% by the formula

$$\frac{\text{area}(B_p \cap B_{gt})}{\text{area}(B_p \cup B_{gt})} > 50\%.$$

For online GSLDA, predicted faces in previous frames are used to update the GSLDA model. Note that the updated patches could contain both true positives (faces) and false positives (misclassified nonfaces). After the update process, the classifier predicts a single patch with highest classification score in the next frame as the face patch. This learning technique is similar to semisupervised learning where the classifier makes use of the unlabeled data in conjunction with a small amount of labeled data. Note that unlike the work in [18] where both positive and negative patches are used to incrementally update their model, we only make use of positive patches.

Fig. 5. Comparison of offline AdaBoost based frontal face detector [8] (top row), AsymBoost based face detector [32] (second row), GSLDA based face detector [23] (third row) and our proposed OGSLDA face detector (last row). All detectors are trained initially with 1000 faces and 10,000 nonfaces. Online GSLDA is incrementally updated with patches classified as faces from previous video frames. The first video (*david indoor*) contains 761 frames of a person moving from a dark to a bright area undergoing large lighting and pose changes (frames 150, 250, 350, 409, 450, 494, and 592). The second video (*trellis*) contains 501 frames of a person moving underneath a trellis with large illumination change (frames 50, 85, 182, 231, 287, 386, and 457).

TABLE II
PERFORMANCE ON FOUR DIFFERENT FRONTAL FACE DETECTORS
ON DAVID INDOOR AND TRELLIS VIDEO

	detection rate	
	indoor sequence	trellis sequence
AdaBoost [8]	57.8%	35.3%
AsymBoost [32]	68.7%	37.5%
GSLDA [23]	70.3%	48.5%
Our proposed OGSLDA	83.1%	62.1%

Table II compares four face detectors in terms of their performance. We observe that the performance of AdaBoost face detectors is the worst. This is not surprising since the distributions of training data are highly skewed (1000 faces and 10,000 nonfaces). Viola and Jones also pointed out this limitation in [32]. Face detectors trained using AsymBoost and GSLDA perform quite similar on the first video. Our results are consistent with those reported in [23]. Table II shows that online GSLDA performs best. Based upon our observations, incrementally updating the GSLDA model improves detection results significantly at small increase in computation time. Fig. 5 compares

empirical results between offline GSLDA and our proposed online GSLDA.

Finally, we compare receiver operating characteristic (ROC) curves between the offline GSLDA model (1000 faces and 10,000 nonfaces) and the online GSLDA model (initially trained with 1000 faces and 10,000 nonfaces plus updated with 661 patches classified as faces). In this experiment, we set the scaling factor to 1.2 and window stepping size to 1 pixel. The technique used for merging overlapping windows is similar to [8]. Detections are considered true or false positives based upon the area overlap with ground truth bounding boxes. We shift the classifier threshold and plot ROC curves (Fig. 6). Clearly, updating the trained model with relevant training data increases the overall performance of classifiers.

In the next experiment, we compare the performance of single strong classifiers learned using offline GSLDA and online GSLDA on frontal faces database. The database consists of 10,000 mirrored faces. Faces were cropped and rescaled to images of size 24×24 pixels. For nonface examples, we randomly selected 10,000 random nonface patches from nonface

Fig. 6. Comparison of ROC curves between offline and online GSLDA on David Ross indoor data sets (top) and trellis data sets (bottom).

Fig. 7. Random sample of face images for training.

TABLE III
SIZE OF TRAINING AND TEST SETS USED ON THE SINGLE NODE CLASSIFIER

#	data splits	faces/split	non-faces/split
Train	3	2000	2000
Test	2	2000	2000

images obtained from the internet. The collected patches are split into three training sets and two test sets. Each set contains 2000 face examples and 2000 nonface examples (Table III). For each experiment, three different classifiers are generated, each by selecting two out of three training sets and the remaining training set for validation.

In this experiment, we train 30, 50, and 100 weak learners of Haar-like features. The performance is measured by the test

error rate. Results are shown in Fig. 8 and following observations can be made. The error of both classifiers drops as the number of training samples increases. The error rate of batch GSLDA drops at a slightly faster rate than online GSLDA. This is not surprising. For batch learning, the previous set of training samples along with a new sample are used to update the decision stumps every time a new sample is inserted. For each update, the batch algorithm throws away previously selected weak classifiers and reselects the new 30, 50, and 100 weak classifiers. As a result, the training process is time consuming and requires a large amount of storage. In contrast, online GSLDA relies on initial trained decision stumps. The new instance does not update previously trained decision stumps but between-class and within-class scatter matrices. The process is suboptimal compared to batch GSLDA. However, the slight increase in performance of batch GSLDA over online GSLDA (0.7% drop in test error rate for 100 weak classifiers) comes at a much higher storage cost and significantly higher computation cost.

2) *Performances on Cascades of Strong Classifiers*: In this experiment, we use mirrored faces from previous experiment to compare batch learning and online learning. The number of initial positive samples used in each experiment is varied. We use 500, 1000, and 5000 faces to initially train a face detector. In each experiment, we trained four different cascaded detectors. The first cascaded detector is the same as in Viola and Jones [8], i.e., face data sets used in each cascade stage are the same while the nonface samples used in each cascade layer are collected from false positives of previous stages of cascades (bootstrapping). The cascade training algorithm terminates when there are not enough negative samples to bootstrap.

The second, third, and forth face detectors are trained initially with the technique similar to the first cascaded detector. However, the second cascaded face detector is incrementally updated with new negative examples collected from false positives of previous stages of cascades. The third cascaded face detector is incrementally updated with 5000 unseen faces. The final face detector is incrementally updated with both false positives from previous stages and unseen faces. For each face detector, weak classifiers are added to the cascade until the predefined objective is met. In this experiment, we set the minimum detection rate in each cascade stage to be 99% and the maximum false positive rate to be 50%.

We tested our face detectors on MIT+CMU frontal face test sets. Fig. 9 shows a comparison between ROC curves produced by online GSLDA. ROC curves in Fig. 9(a) show that the online GSLDA classifier outperforms offline GSLDA at all false positive rates when initially trained with 500 faces. Incrementally updating the GSLDA model with unseen faces (+5000 faces) yields a better result than updating the model with new false positives from previous stages of cascades (+ 10^6 negative patches). The online classifier performs best when updated with both new positive and negative patches. Fig. 9(b) shows a comparison when the number of initial training samples has been increased to 1000 faces. The performance gap between GSLDA and online GSLDA is now smaller. We observe the performance of both GSLDA and online GSLDA (+ 10^6 negative patches) to be very similar. This indicates that the cascade learning framework proposed by Viola and Jones might have al-

Fig. 8. Comparison of classification error rates between batch GSLDA and online GSLDA. The number of weak learners (decision stumps on Haar-like features) in each experiment is (a) 30, (b) 50, and (c) 100. The error of both classifiers drops as the number of training samples increases.

Fig. 9. Comparison of ROC curves on MIT+CMU face test set. The four detectors are trained using (a) 500 faces, (b) 1000 faces, and (c) 5000 and 10,000 mirrored faces.

ready incorporated the benefit of massive negative patches. Incremental learning with new negative instances does not seem to improve the performance of cascaded detectors any further. Another way to explain results of our findings is to use the concept of linear asymmetric classifier (LAC) proposed in [15]. In [15], the asymmetric node learning goal is expressed as

$$\begin{aligned} & \underset{\mathbf{w}, w_0}{\text{maximize}} \quad \Pr_{\mathbf{x} \sim (\mathbf{m}_1, \Sigma_1)} \{ \mathbf{w}^\top \mathbf{x} \geq w_0 \} \\ & \text{subject to} \quad \Pr_{\mathbf{y} \sim (\mathbf{m}_2, \Sigma_2)} \{ \mathbf{w}^\top \mathbf{y} < w_0 \} = \beta. \end{aligned} \quad (20)$$

Since, the problem has no closed-form solution, the authors developed an approximate solution when $\beta = 0.5$. To find a closed-form solution, the authors assumed that $\mathbf{w}^\top \mathbf{x}$ is Gaussian for any $\mathbf{w}$, class C_2 distribution is symmetric and the median value of the class C_2 distribution is close to its mean. The direction $\mathbf{w}$ can then be approximated by

$$\underset{\mathbf{w} \neq 0}{\text{maximize}} \quad \frac{\mathbf{w}^\top (\mathbf{m}_1 - \mathbf{m}_2)}{\sqrt{\mathbf{w}^\top \Sigma_1 \mathbf{w}}}. \quad (21)$$

From their objective functions, the only difference between FDA (5) and LAC (21) is that the pooled covariance matrix of FDA, $\Sigma_1 + \Sigma_2$, is replaced by the covariance matrix of class C_1 , Σ_1 . In other words, when train the classifier with the asymmetric node learning goal for the cascade learning framework, the variance of negative classes becomes less

relevant. In contrast, new instances of positive classes affect both the numerator and denominator in (21). Hence, it is easier to notice the performance improvement when new positive instances are inserted. Our results seem to be consistent with their derivations.

We further increase the number of initial training faces to 5000. All face detectors now seem to perform very similar to each other. We conjecture that this is the best performance that our cascaded detectors with provided training sets can achieve on MIT+CMU data sets. Results of GSLDA face detectors trained with 10,000 faces and 10,000 nonfaces seem to support our assumptions [Fig. 9(c)]. To further improve the performance, different cascade algorithms, e.g., soft-cascade [36], WaldBoost [37], multiexit classifiers [38], etc. and a combination with other types of features, e.g., edge orientation histograms (EOH) [39], covariance features [40], etc., can also be experimented. Fig. 10 shows a comparison of the computation cost between batch GSLDA and online GSLDA. The horizontal axis shows the number of weak learners (decision stumps) and the vertical axis indicates the training time in minutes. From the figure, online learning is much faster than training the batch GSLDA classifier as the number of weak learners grows. On average, our online classifier takes less than 1.5 ms to update a strong classifier of 200 weak learners on a standard PC with the use of GNU scientific library (GSL).⁶

⁶<http://www.gnu.org/software/gsl/>.

Fig. 10. Comparison of the training time between GSLDA and the online GSLDA algorithm. The first and second GSLDA detectors are trained with 5000 faces and 5000 nonfaces, and 10,000 faces and 10,000 nonfaces, respectively. Online GSLDA is initially trained with 5000 faces and 5000 nonfaces and updated with one million new patches. Notice that there is a slight increase in training time even though we incrementally update with $200\times$ more training samples.

C. Discussion

It can be argued that given a large set of training images, the AdaBoost based object detector can be trained to achieve a high detection performance. However, we think that this statement is not always true. In object detection, there are usually large variations due to visual appearances, object poses, illuminations and camera motions. Based upon our experiments, AdaBoost can overfit and gives poor performance when there exists a large overlapping between object classes, e.g., a mix of frontal, tilted and profile faces, or a mix of faces with and without occlusions, etc. To detect a turning object, a number of researchers have proposed techniques which employ object detection with object tracking [41], [42]. Nonetheless, in this paper, we show that it is possible to detect turning objects using object detection technique *alone*. In our experiments, we have shown that our online version of GSLDA can perform very similarly to the offline counterpart in terms of detection performance but computationally cheaper. We have also shown that the online GSLDA performs better than the standard online boosting, which is considered as the state-of-the-art.

IV. CONCLUSION

In this work, we have proposed an efficient online object detection algorithm. Unlike many existing algorithms which are based upon online boosting, our framework makes use of GSLDA based feature selection which aims to maximize the class-separation criterion. Our experimental results show that our incremental algorithm does not only perform comparable to the batch GSLDA algorithm but is also much more efficient. On USPS digits data sets, our online algorithm with decision stumps as weak learners outperforms online boosting with class-conditional Gaussian distributions. Our extensive experiments on face detections reveal that it is almost always

beneficial to incrementally train the detector with online samples. Ongoing works include the search for more accurate and efficient online weak learners.

REFERENCES

- [1] M.-H. Yang, D. J. Kriegman, and N. Ahuja, "Detecting faces in images: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 1, pp. 34–58, Jan. 2002.
- [2] E. Hjelm and B. K. Low, "Face detection: A survey," *Comput. Vis. Image Understand.*, vol. 83, no. 3, pp. 236–274, 2001.
- [3] S. Munder and D. M. Gavrila, "An experimental study on pedestrian classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 11, pp. 1863–1868, Nov. 2006.
- [4] D. Gerónimo, A. M. López, A. D. Sappa, and T. Graf, "Survey on pedestrian detection for advanced driver assistance systems," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 7, pp. 1239–1258, Jul. 2010.
- [5] P. Campadelli, R. Lanzarotti, and G. Lipori, "Eye localization: A survey," *NATO Sci. Ser.*, vol. 18, pp. 234–245, 2007.
- [6] Z. Sun, G. Bebis, and R. Miller, "On-road vehicle detection: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 5, pp. 694–711, May 2006.
- [7] C. Papageorgiou and T. Poggio, "A trainable system for object detection," *Int. J. Comput. Vis.*, vol. 38, no. 1, pp. 15–33, 2000.
- [8] P. Viola and M. J. Jones, "Robust real-time face detection," *Int. J. Comput. Vis.*, vol. 57, no. 2, pp. 137–154, 2004.
- [9] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, San Diego, CA, 2005, vol. 1, pp. 886–893.
- [10] P. Wang, C. Shen, H. Zheng, and Z. Ren, "Training a multi-exit cascade with linear asymmetric classification for efficient object detection," in *Proc. IEEE Int. Conf. Image Process.*, Hong Kong, 2010, pp. 61–64.
- [11] R. E. Schapire, "Theoretical views of boosting and applications," in *Proc. Int. Conf. Algorithmic Learn. Theory*, Tokyo, Japan, 1999, pp. 13–25.
- [12] C. Shen and H. Li, "Boosting through optimization of margin distributions," *IEEE Trans. Neural Netw.* vol. 21, no. 4, pp. 659–666, Apr. 2010.
- [13] C. Shen and H. Li, "On the dual formulation of boosting algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.* vol. 32, no. 12, pp. 2216–2231, Dec. 2010.
- [14] M. T. Pham and T. J. Cham, "Fast training and selection of Haar features using statistics in boosting-based face detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, Rio de Janeiro, Brazil, 2007, pp. 1–7.
- [15] J. Wu, S. C. Brubaker, M. D. Mullin, and J. M. Rehg, "Fast asymmetric learning for cascade face detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 3, pp. 369–382, Mar. 2008.
- [16] R. Xiao, H. Zhu, H. Sun, and X. Tang, "Dynamic cascades for face detection," in *Proc. IEEE Int. Conf. Comp. Vis.*, Rio de Janeiro, Brazil, 2007, pp. 1–8.
- [17] N. C. Oza and S. Russell, "Online bagging and boosting," in *Proc. Artif. Intell. Statist.*, Key West, FL, 2001, pp. 105–112.
- [18] H. Grabner and H. Bischof, "On-line boosting and vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, New York, 2006, vol. 1, pp. 260–267.
- [19] C. Huang, H. Ai, T. Yamashita, S. Lao, and M. Kawade, "Incremental learning of boosted face detector," in *Proc. IEEE Int. Conf. Comput. Vis.*, Rio de Janeiro, Brazil, 2007.
- [20] T. Parag, F. Porikli, and A. Elgammal, "Boosting adaptive linear weak classifiers for online learning and tracking," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, Anchorage, AK, 2008, pp. 1–8.
- [21] X. Liu and T. Yu, "Gradient feature selection for online boosting," in *Proc. IEEE Int. Conf. Comp. Vis.*, Rio de Janeiro, Brazil, 2007, pp. 1–8.
- [22] B. Moghaddam, Y. Weiss, and S. Avidan, "Fast pixel/part selection with sparse eigenvectors," in *Proc. IEEE Int. Conf. Comput. Vis.*, Rio de Janeiro, Brazil, 2007, pp. 1–8.
- [23] S. Paisitkriangkrai, C. Shen, and J. Zhang, "Efficiently training a better visual detector with sparse eigenvectors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Miami, FL, 2009, pp. 1129–1136.
- [24] S. Pang, S. Ozawa, and N. Kasabov, "Incremental linear discriminant analysis for classification of data streams," *IEEE Trans. Syst., Man, Cybern. B*, vol. 35, no. 5, pp. 905–914, Oct. 2005.
- [25] J. Ye, Q. Li, H. Xiong, H. Park, R. Janardan, and V. Kumar, "IDR/QR: An incremental dimension reduction algorithm via QR decomposition," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 9, pp. 1208–1222, Sep. 2005.

- [26] H. Zhao and P. C. Yuen, "Incremental linear discriminant analysis for face recognition," *IEEE Trans. Syst., Man, Cybern. B*, vol. 38, no. 1, pp. 210–221, Feb. 2008.
- [27] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, 2nd ed. Hoboken, NJ: Wiley, 2001.
- [28] B. Moghaddam, Y. Weiss, and S. Avidan, "Generalized spectral bounds for sparse LDA," in *Proc. Int. Conf. Mach. Learn.*, Pittsburgh, PA, 2006, pp. 641–648.
- [29] T.-K. Kim, S.-F. Wong, B. Stenger, J. Kittler, and R. Cipolla, "Incremental linear discriminant analysis using sufficient spanning set approximations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Minneapolis, MN, 2007, pp. 1–8.
- [30] R. Bru, J. M. J. Mas, and M. Tuma, "Balanced incomplete factorization," *SIAM J. Sci. Comput.*, vol. 30, pp. 2302–2318, 2008.
- [31] L. G. Rueda, "An efficient approach to compute the threshold for multi-dimensional linear classifiers," *Pattern Recognit.*, vol. 37, no. 4, pp. 811–826, 2004.
- [32] P. Viola and M. J. Jones, "Fast and robust classification using asymmetric adaboost and a detector cascade," in *Proc. Adv. Neural Inf. Process. Syst.*, Vancouver, Canada, 2002, pp. 1311–1318.
- [33] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.
- [34] S. Z. Li and Z. Zhang, "FloatBoost learning and statistical face detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 9, pp. 1112–1123, Sep. 2004.
- [35] The PASCAL Visual Object Classes Challenge (VOC 2007) [Online]. Available: <http://www.pascal-network.org/challenges/VOC/voc2007/index.html>
- [36] L. Bourdev and J. Brandt, "Robust object detection via soft cascade," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, San Diego, CA, 2005, pp. 236–243.
- [37] J. Sochman and J. Matas, "Waldboost—Learning for time constrained sequential detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, San Diego, CA, 2005, vol. 2, pp. 150–156.
- [38] M. T. Pham, V. D. D. Hoang, and T. J. Cham, "Detection with multi-exit asymmetric boosting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Anchorage, AK, 2008.
- [39] K. Levi and Y. Weiss, "Learning object detection from a small number of examples: The importance of good features," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Washington, DC, 2004, vol. 2, pp. 53–60.
- [40] C. Shen, S. Paisitkriangkrai, and J. Zhang, "Face detection from few training examples," in *Proc. IEEE Int. Conf. Image Process.*, San Diego, CA, 2008, pp. 2764–2767.
- [41] K. Okuma, A. Taleghani, N. de Freitas, J. Little, and D. G. Lowe, "A boosted particle filter: Multitarget detection and tracking," in *Proc. Eur. Conf. Comput. Vis.*, Prague, Czech Republic, 2004, vol. 3021/2004, pp. 28–39.
- [42] T. Chateau, V. Gay-Belille, F. Chausse, and J.-T. Lapreste, "Real-time tracking with classifiers," *Dyn. Vis.*, vol. 4358, pp. 218–231, 2007.

Sakrapee Paisitkriangkrai received the B.E. degree in computer engineering, the M.E. degree in biomedical engineering, and the Ph.D. degree from the University of New South Wales, Sydney, Australia, in 2003 and 2010, respectively.

He is currently with the University of Adelaide, Australia. His research interests include pattern recognition, image processing, and machine learning.

Chunhua Shen received the Ph.D. degree from the School of Computer Science, University of Adelaide, Adelaide, Australia in 2005; and the M.Phil. degree from Mathematical Sciences Institute, Australian National University, Canberra, Australia, in 2009.

Since Oct. 2005, he has been working with the Computer Vision Program, National ICT Australia (NICTA), Canberra Research Laboratory, Canberra, Australia, where he is a Senior Research Fellow and holds a continuing research position. His main research interests include statistical machine learning

and its applications in computer vision and image processing.

Jian Zhang (M'98–SM'04) received the Ph.D. degree in electrical engineering from the University College, University of New South Wales, Australian Defence Force Academy, Australia, in 1997.

He is a Principal Researcher at National ICT Australia (NICTA), Sydney, Australia. He is also a conjoint Associate Professor at the University of New South Wales. His research interests include image/video processing, video surveillance and multimedia content management.

Dr. Zhang is currently an associate editor of the *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY* and the *Journal on Image and Video Processing*.