Multiple-Instance Hidden Markov Model for GPR-Based Landmine Detection

Achut Manandhar, Peter A. Torrione, Leslie M. Collins, *Senior Member, IEEE*, and Kenneth D. Morton, *Member, IEEE*

***Abstract*—Hidden Markov models (HMMs) have previously been successfully applied to subsurface threat detection using ground penetrating radar (GPR) data. However, parameter estimation in most HMM-based landmine detection approaches is difficult since object locations are typically well known for the 2-D coordinates on the Earth's surface but are not well known for object depths underneath the ground/time of arrival in a GPR A-scan. As a result, in a standard expectation maximization HMM (EM-HMM), all depths corresponding to a particular alarm location may be labeled as target sequences although the characteristics of data from different depths are substantially different. In this paper, an alternate HMM approach is developed using a multiple-instance learning (MIL) framework that considers an unordered set of HMM sequences at a particular alarm location, where the set of sequences is defined as positive if at least one of the sequences is a target sequence; otherwise, the set is defined as negative. Using the MIL framework, a collection of these sets (bags), along with their labels is used to train the target and nontarget HMMs simultaneously. The model parameters are inferred using variational Bayes, making the model tractable and computationally efficient. Experimental results on two synthetic and two landmine data sets show that the proposed approach performs better than a standard EM-HMM.**

***Index Terms*—Ground penetrating radar (GPR), hidden Markov model (HMM), landmine detection, multiple-instance learning (MIL), variational Bayes (VB).**

I. Introduction

LANDMINES pose a serious threat to soldiers and civilians worldwide and also provide major challenges to agriculture, infrastructure and road development in post-conflict regions [1]. In order to demine the affected areas, several techniques have been developed to detect subsurface threats. Since electromagnetic induction (EMI) sensors can only detect metal mines, a variety of other techniques are being explored, most of which exploit the electromagnetic characteristics of the mines or the mine casing; for example, ground penetrating radar (GPR), infrared/hyperspectral methods, acoustic/seismic methods, etc. [2]. GPR can detect nonmetal or low-metal content subsurface targets but naive processing of GPR also flags anomalies such as roots, rocks, potholes, and other clutter objects as potential threats. To reduce false alarms, current GPR-based landmine detection systems first identify potential threat locations using a computationally efficient technique known as a prescreening algorithm [3]. These prescreener alarms are then further discriminated into target and nontarget classes using more sophisticated methods [4]–[10].

One of the more successful techniques for GPR anomaly discrimination uses hidden Markov models (HMMs) to characterize 2-D GPR data over multiple depth bins at a particular alarm location as sequences of time samples [11]–[14]. This technique is motivated by the typical hyperbolic shape of landmines in 2-D GPR data, which can be represented by an HMM sequence of time samples as shown in Fig. 1. The parameter estimation procedures in most HMM approaches are affected by the unknown alarm depth locations in the training set of 3-D GPR data. The prescreener alarm locations are known in 2-D coordinates on the Earth's surface, but not for individual target depths. Due to the absence of depth labels, in a standard expectation maximization HMM (EM-HMM), sequences at depth bins both below and above a target are all labeled as targets despite being substantially different in their characteristics. The uncertainty in the sequence labels with depth complicates learning the HMM parameters, potentially degrading the classifier's performance. Obtaining alarm labels at individual depths requires hand labeling, which is extremely time consuming and potentially error prone. While new research efforts are in progress to automate an efficient labeling process [15], this work proposes an alternate approach based on a multiple-instance learning (MIL) framework to incorporate the label uncertainty into the classifier training process. In the MIL-based approach, a set of sequences at a particular alarm location is considered a bag. A bag is labeled positive if at least one of its sequences is a positive sequence; otherwise, the bag is labeled negative. This work integrates HMMs into a MIL framework and develops a new approach called the multiple-instance hidden Markov model (MiHMM).

A variety of HMM approaches have been developed for threat detection using GPR [11]–[14] and other technologies [16]–[18]. In [11], gradient-based features characterizing the edge strengths in diagonal and antidiagonal directions in 2-D GPR data were used to learn HMMs with discrete and continuous observation models. The performance of [11] was improved by discriminately training an HMM using different techniques [12], [13]. Since the features in [11] are constrained to characterize only the diagonal and antidiagonal edges, more generalizable features were constructed in [14] for time-domain GPR data by implementing Gabor filters on 2-D GPR data. Despite

Manuscript received October 1, 2013; revised March 13, 2014 and June 19, 2014; accepted June 28, 2014. This work was supported by the U.S. Army RDECOM CERDEC Night Vision and Electronic Sensors Directorate through the Army Research Office (ARO) under Grant W911NF-09-1-0487 and Grant W911NF-06-1-0357.

The authors are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: kdm@ee.duke.edu).

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/TGRS.2014.2346954

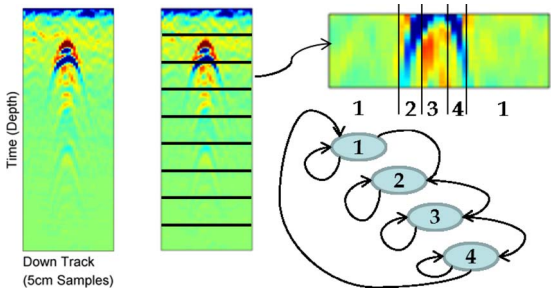

Fig. 1. From left to right, (a) a typical hyperbolic response of a buried landmine in 2-D GPR data, (b) the data divided into multiple depth bins, where each depth bin is a sequence of down track samples, and (c) a sequence represented as an HMM sequence whose states correspond to background and three edges—rising, flat, and falling.

advances made to the set of features applied to HMM-based landmine detection, none of the above approaches have considered the label uncertainty of the sequences due to the multiple-instance nature of landmine data, although all approaches have noted that this label uncertainty is a concerning issue.

MIL has been widely applied in different domains, including landmine detection. A vast majority of these efforts have been related to content-based image retrieval, image classification, and object recognition problems. Many MIL approaches have been inspired by the Diverse Density approach [19]–[22], maximum-margin classifiers [23], [24], instance selection methods [25], [26], and boosting [27]. More recent applications of MIL are in object tracking [28], [29]. MIL methods have also been recently applied in the subsurface target detection and have been shown to perform better than their non-MIL counterparts. A random set framework for MIL (RSFMIL) [30] and multiple-instance Relevance Vector Machines (MIRVM) [31] were applied to detect landmines using remotely sensed hyperspectral images in [32].

Most of the existing MIL approaches are unsuitable for time-series data and do not generalize well across different applications. Recently, a MIL framework has been used in HMM-based landmine detection [33], which implements Diverse Density (DD) [19] on time series data. Although their preliminary work shows improvement over a standard EM-HMM, the DD itself suffers from limitations that have been addressed in subsequent efforts [24] and [34]. In this paper, our previous nonparametric Bayesian multiple-instance learning (NPBMIL) approach [35] is extended to incorporate time-varying GPR data by learning HMMs in a MIL framework (MiHMM). The proposed approach is generalizable to different applications by using an appropriate observation model, and in this work, we consider both discrete (multinomial) and continuous (Gaussian) features. The proposed method uses variational Bayes (VB) to estimate the model parameters, making parameter inference fast and computationally efficient. The performance of the proposed MiHMM method is compared against NPBMIL and a standard EM-HMM using feature developed in [11] as well as a new set of features generated using a very successful computer vision technique called the histogram of oriented gradients (HOG) [36].

The remainder of this work is organized as follows. Section II describes the two different MiHMM approaches and the parameter inference for both. Section III describes the simulated data sets and two landmine data sets used in this work and compares the performance of the proposed approach with a standard EM-HMM and NPBMIL. Finally, Section IV concludes this paper and discusses directions for potential future work.

II. Multiple-Instance Hidden Markov Models

In the MIL framework, data is represented by a collection of bag/label pairs, $\{\mathbf{X}_n, Y_n\}_{N=1}^{N}$. A bag, $\mathbf{X}_n$, consists of a set of instance/label pairs, $\mathbf{X}_n = \{\mathbf{x}_{ni}, y_{ni}\}_{i=1}^{I_n}$, $i \in \{1, \ldots, I_{ni}\}$, $n \in \{1, \ldots, N\}$, where N denotes the total number of bags and I_n denotes the total number of instances in the nth bag. For any bag, $\mathbf{X}_n$, only the bag label, Y_n, is observed, whereas its instance labels, $\{y_{ni}\}_{i=1}^{I_n}$, are unobserved. In the context of sequential data, each instance, $\mathbf{x}_{ni}$, can be represented by a sequence of time samples, $\{\mathbf{x}_{ni,t}\}_{t=1}^{T_{ni}}$, $t \in \{1, \ldots, T_{ni}\}$, where T_{ni} is the total number of time samples in the ith instance in the nth bag. For a binary MIL problem, the latent instance labels influence the bag labels as follows: If the time samples in an instance, $\mathbf{x}_{ni}$, $i \in \{1, \ldots, I_n\}$, are generated from probability densities $p(\mathbf{x}_{ni}|y_{ni} = 1)$, then the bag label, Y_n, is 1; otherwise Y_n is 0.

As described earlier, the HMM approaches developed for GPR-based landmine detection [11], [13], [14] do not incorporate the inherent uncertainty in the instance labels. In the absence of instance labels, a standard EM-HMM represents all instances in multiple depths at a particular spatial location with their corresponding bag label. Consequently, instances below and above a target instance are labeled as targets although their characteristics might be substantially different from targets. In an attempt to localize the regions corresponding to the target, Gaussian Markov random fields (GMRFs) have been used to segment the information-bearing region of the GPR image from the background of the image [15]. However, image segmentation can be a very time consuming procedure. Moreover, the image segmentation procedure does not incorporate the time-varying characteristics of landmine response, which should be closely coupled in the HMM learning process.

In order to overcome these limitations, this work proposes an HMM approach based on a MIL framework (MiHMM), in which bags of instances are drawn from a hierarchical mixture of two HMMs. Time samples in at least one of the instances in a $Y = 1$ bag are drawn from probability densities $p(\mathbf{x}_{ni,t}|y_{ni} = 1)$, and time samples in all of the instances in a $Y = 0$ bag are generated from probability densities $p(\mathbf{x}_{ni}|y_{ni} = 0)$. In contrast to the previous HMM approaches, MiHMM learns the mixture of HMMs simultaneously.

A. Multiple-Instance Hidden Markov Models

The HMM described in [11] forms a model that can be used to represent sequences from two different hypotheses $p(\mathbf{x}_{ni}|y_{ni} = 0)$ and $p(\mathbf{x}_{ni}|y_{ni} = 1)$. Using the HMM as a building block, we construct a generative model for time-series data in a MIL scenario. For a binary MIL problem, it is assumed that the time samples in any instance in a negative bag $(Y_n = 0)$ are sequentially generated from $p(\mathbf{x}_{ni,t}|\mathbf{z}_{ni,t}, y_{ni} = 0)$,

whereas the time samples in any instance in a positive bag ($Y_n = 1$) are generated from a hierarchical mixture of two HMMs as follows: First the hidden instance label, y_{ni}, is drawn from a multinomial distribution with mixing proportions, $\boldsymbol{\eta}$, and based on y_{ni}, the time samples in the instance are sequentially generated from either $p(\mathbf{x}_{ni}|y_{ni} = 0)$ or $p(\mathbf{x}_{ni}|y_{ni} = 1)$.

A latent variable $\boldsymbol{\zeta}_{ni}$ assigns one of the two HMMs to the instance $\mathbf{x}_{ni}$ as follows:

$$\boldsymbol{\zeta}_{ni} \begin{cases} \sim Mult(\boldsymbol{\eta}), \boldsymbol{\eta} \sim Dir(\boldsymbol{\alpha}_0) & \text{if } Y_n = 1 \\ = 0 & \text{if } Y_n = 0 \end{cases}$$

where $Y_n = 1$ denotes the bag is positive, $Y_n = 0$ denotes the bag is negative, and $\boldsymbol{\eta} = \{\eta_0, \eta_1\}$ denote the mixing proportions of the two HMMs for an instance in a $Y = 1$ bag.

Let $\mathcal{H}_m = \{\boldsymbol{\Theta}_k^m, \Pi_k^m, \mathbf{A}^m\}$ denote the parameters of the probability densities, initial state probabilities, and state transition probabilities of the mth HMM, $m \in \{0, 1\}$. Let $\mathbf{z}_{ni,t}^m$ be the discrete multinomial latent variable that describes which component density in the mth HMM is responsible for generating the corresponding observation, $\mathbf{x}_{ni,t}$. The data generation process can be explained as follows:

$$\mathbf{z}_{ni,1}^m|\boldsymbol{\Pi}^m \sim Mult(\boldsymbol{\Pi}^m) \tag{1}$$

$$\mathbf{z}_{i,t}^m|\left\{z_{i,t-1}^m = k, \mathbf{A}_{k\bullet}^m\right\} \sim Mult\left(\mathbf{A}_{k\bullet}^m\right), k \in \{1, \ldots, K^m\}, \quad t \in \{1, \ldots, T_i\} \tag{2}$$

$$\mathbf{x}_{ni,t}|\mathbf{z}_{ni,t}^m \sim \prod_{k=1}^{K^m} p\left(\mathbf{x}_{ni,t}|\boldsymbol{\Theta}_k^m\right)^{z_{ni,t}^m} \tag{3}$$

$$\mathbf{X}_n|\{\mathcal{H}_0, \mathcal{H}_1\} \sim \prod_{i=1}^{I_n} \prod_{m=\{0,1\}} p(\mathbf{x}_{ni}|\mathcal{H}_m)^{\zeta_{ni}}. \tag{4}$$

The uncertainty in the parameters can be expressed using conjugate priors

$$\boldsymbol{\Pi}^m|\boldsymbol{\Lambda}_0^m \sim Dir\left(\boldsymbol{\Lambda}_0^m\right) \tag{5}$$

$$\mathbf{A}_{k\bullet}^m|\mathbf{a}_{0\bullet}^m \sim Dir\left(\mathbf{a}_{0\bullet}^m\right), \mathbf{A}_{k\bullet}^m = [A_{k1}^m, \ldots, A_{kK}^m], \quad \mathbf{a}_{0\bullet}^m = [a_{01}^m, \ldots, a_{0K}^m] \tag{6}$$

$$\boldsymbol{\Theta}_k^m|\mathbf{h}_0^m \sim p\left(\boldsymbol{\Theta}_k^m|\mathbf{h}_0^m\right) \tag{7}$$

$$\boldsymbol{\eta}|\boldsymbol{\alpha}_0 \sim Dir(\boldsymbol{\alpha}_0). \tag{8}$$

Fig. 2 shows a graphical model for the MiHMM model and differs from the HMM model in [11] in two important aspects: (1) there are two HMMs in Fig. 2, and (2) a latent variable, $\boldsymbol{\zeta}_{ni}$ is introduced that assigns an instance in a $Y = 1$ bag to one of the two HMMs. The dashed plates, denoting the exponential family component densities, can be replaced by one of the two specific densities, i.e., multinomial (MN) and multivariate normal (MVN).

1) MN: In the discrete landmine features developed in [11], each time sample is represented by one of the 16 discrete observations, which can be modeled by a multinomial probability density. Consider a MN probability density with dimensionality D parameterized by a probability vector $\boldsymbol{\Theta}_k^m = \mathbf{p}_k^m$ such that $\mathbf{x}_{ni,t}|\mathbf{p}_k^m \sim Mult(\mathbf{p}_k^m)$. For a MN probability density, $\mathbf{x}_{ni,t} = [x_{ni,t1}, \ldots, x_{ni,tD}]'$, where $x_{ni,td}$ denotes the count of the dth outcome and p_{kd}^m denotes the probability of that outcome. A special case of MN probability density, where each sample has exactly one outcome, can be used to model the time samples in the discrete landmine data set. The conjugate prior for $\mathbf{p}_k^m$ is the Dirichlet density

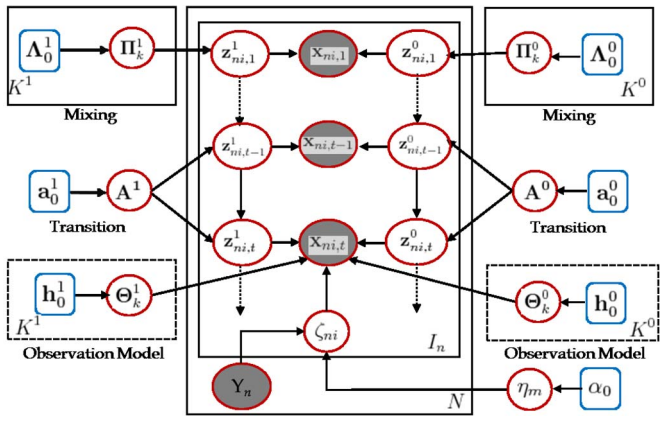

Fig. 2. MiHMM1 graphical model. The latent variable $\boldsymbol{\zeta}_{ni}$ assigns one of the HMMs to the instance $\mathbf{x}_{ni}$. Given a $Y_n = 1$ bag, the HMMs mix with proportion $\boldsymbol{\eta}$. Given an HMM, the latent variable $\mathbf{z}_{ni,t}^m$ assigns one of the mixture components, parameterized by $\{\boldsymbol{\Theta}_k^m\}_{k=1}^K$, to the sample $\mathbf{x}_{ni,t}$. Excluding the first state, all other successive states in a sequence are determined by the state transition probabilities $\mathbf{A}^m$. The first state is determined by the initial state mixing proportion $\boldsymbol{\Pi}^m$. This model is modified to have a shared observation model in MiHMM2, whose graphical model is shown in Fig. 3.

$$\mathbf{p}_k^m|\boldsymbol{\Omega}_0^m \sim Dir\left(\boldsymbol{\Omega}_0^m\right). \tag{9}$$

2) MVN: In the HOG landmine features, described in Section III-B2, each time sample is described by a 9-D feature vector. Intuitively, this feature vector gives a weighted distribution of edge orientations within a particular region in a 2-D GPR image, which can be modeled by a multivariate normal (MVN) probability density. A MVN density with dimensionality D is defined by a $[D \times 1]$ mean vector and a $[D \times D]$ covariance matrix. Let $\boldsymbol{\Theta}_k^m = \{\boldsymbol{\mu}_k^m, \boldsymbol{\Gamma}_k^m\}$ denote the parameters of the kth component in the mth HMM mixture model, where $\boldsymbol{\mu}_k^m$ denotes the mean and $\boldsymbol{\Gamma}_k^m$ denotes the precision such that $\mathbf{x}_{ni,t}|\{\boldsymbol{\mu}_k^m, \boldsymbol{\Gamma}_k^m\} \sim \mathcal{N}(\boldsymbol{\mu}_k^m, \{\boldsymbol{\Gamma}_k^m\}^{-1})$. The conjugate prior for the joint density of $\boldsymbol{\mu}_k^m$ and $\boldsymbol{\Gamma}_k^m$ is the Normal-Wishart density as follows:

$$\boldsymbol{\Gamma}_k^m|\left\{\boldsymbol{\Phi}_0^m, \nu_0^m\right\} \sim Wish\left(\nu_0^m, \{\boldsymbol{\Phi}_0^m\}^{-1}\right) \tag{10}$$

$$\boldsymbol{\mu}_k^m|\left\{\boldsymbol{\Gamma}_k^m, \boldsymbol{\rho}_0^m, \beta_0^m\right\} \sim \mathcal{N}\left(\boldsymbol{\rho}_0^m, \{\beta_0^m \boldsymbol{\Gamma}_i^m\}^{-1}\right). \tag{11}$$

B. Parameter Inference

Based on the data generation process described in Section II-A, a joint probability density of the time samples and the model parameters can be obtained. Due to the model complexity, the computation of the posterior densities involves an intractable integration. This work utilizes variational Bayes (VB) to approximate the intractable integration by maximizing

a lower bound [37]. To achieve this goal, VB assumes parameterized posterior densities of the parameters can be factorized into conditionally independent groups of parameters. For MiHMM, the assumed conditional independence is as follows:

$$q\left(\{\boldsymbol{\zeta}_n, \mathbf{z}_n, \boldsymbol{\Pi}, \mathbf{A}, \boldsymbol{\Theta}\}|\mathbf{X}\right) = \prod_{i=1}^{I_n} \prod_{m=\{0,1\}} q(\zeta_{nm}|\mathbf{X}) \prod_{k=1}^{K^m} \left[q\left(\mathbf{z}_{ni}^m, \boldsymbol{\Pi}, \mathbf{A}|\mathbf{X}\right) q\left(\boldsymbol{\Theta}_k^m|\mathbf{X}\right)\right]. \tag{12}$$

The parameters of the assumed posterior densities can be optimized to approximate the true posterior densities by iterating a set of closed-form coupled update equations. The parameters of these coupled equations can be solved iteratively until the change in the negative free energy (NFE) is less than a small value, such as 10^{-10}. The set of coupled equations for the MiHMM model, which assumes standard conjugate posterior densities, are detailed in the supplementary materials.[1] These equations require several VB moments, denoted by $\langle\bullet\rangle$, and the variational average likelihood, $\langle \log p(\mathbf{x}_{ni,t}|\boldsymbol{\Theta}_k^m)\rangle_{q(\boldsymbol{\Theta}_k^m)}$, for each observation model analyzed in this work, both of which are given in the supplementary materials.

Let $\phi_{mni}^M = q(\zeta_{mni})$ denote the responsibility of the mth HMM for the instance $\mathbf{x}_{ni}$; and given the mth HMM, let $\phi_{ni,t}^m$ denote the responsibility of the kth state for the time sample $\mathbf{x}_{ni,t}$, and let $\xi(\mathbf{z}_{ni,t}|\mathbf{z}_{ni,t-1})$ correspond to the conditional posterior probability $q(\mathbf{z}_{ni,t-1}, \mathbf{z}_{ni,t}|\mathbf{X})$ for each of the $K \times K$ transition possibilities for $(\mathbf{z}_{ni,t-1}, \mathbf{z}_{ni,t})$ for the time sample $\mathbf{x}_{ni,t}$. The quantities $\phi_{ni,t}^m$ and $\xi(\mathbf{z}_{ni,t}|\mathbf{z}_{ni,t-1})$ are estimated using the forward-backward algorithm [38]. The responsibility ϕ_{mni}^M is defined as follows:

$$\phi_{mni}^M = q(\zeta_{ni} = m) = \frac{\widehat{\phi_{mni}^M}}{\widehat{\phi_{1ni}^M} + \widehat{\phi_{2ni}^M}} \tag{13}$$

$$\widehat{\phi_{mni}^M} = exp\left\{\sum_{t=1}^{T_{ni}} \langle \log p(\mathbf{x}_{ni,t}|\boldsymbol{\Theta}^m)\rangle_{q(\boldsymbol{\Theta}^m)} + \langle \log \eta_m \rangle_{q(\eta_m)}\right\}. \tag{14}$$

After estimating ϕ_{mni}^M, $\phi_{ni,t}^m$, and $\xi(\mathbf{z}_{ni,t}|\mathbf{z}_{ni,t-1})$, the posterior densities for the parameters of component k for the HMM m, $q(\boldsymbol{\Theta}_k^m)$, can be estimated, which are detailed for the two observation models considered in this work in the supplementary materials.

C. MiHMM With a Shared Observation Model

The MiHMM model developed in the previous section assumes that targets and false alarms in GPR data are generated from two separate HMMs with different parameters. Since the background portions of both target and nontarget data are similar, we can allow the background state to be shared across the two HMMs by modeling the data with two separate HMMs that share the same observation model but have unique initial state mixing and state transition probabilities. Motivated by this insight, the MiHMM model is modified by allowing the two HMMs to share a single observation model. Henceforth, the two types of MiHMMs will be referred to as MiHMM1 and MiHMM2, denoting different and shared observation models, respectively.

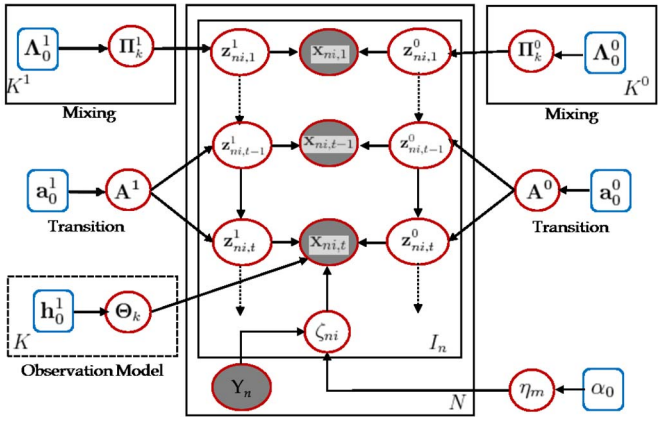

Fig. 3. MiHMM2 graphical model. The latent variable ζ_{ni} assigns one of the HMMs to the instance $\mathbf{x}_{ni}$. Given a $Y_n = 1$ bag, the HMMs mix with proportion $\boldsymbol{\eta}$. Given an HMM, the latent variable $\mathbf{z}_{ni,t}^m$ assigns one of the mixture components, parameterized by $\{\boldsymbol{\Theta}_k\}_{k=1}^K$, to the sample $\mathbf{x}_{ni,t}$. Excluding the first state, all other successive states in a sequence are determined by the state transition probabilities $\mathbf{A}^m$. The first state is determined by the initial state mixing proportion $\boldsymbol{\Pi}^m$. This difference between the models in Figs. 2 and 3 is that there is a single shared observation model in Fig. 3.

Assuming a shared observation model requires only few minor modifications in the method described in Section II-A and B. The parameters of the mth HMM are denoted by $\mathcal{H}_m = \{\boldsymbol{\Theta}, \Pi^m, \mathbf{A}^m\}$, where the shared probability densities are parameterized in terms of $\boldsymbol{\Theta}$ and hyper parameter $\mathbf{h}$, replacing the HMM specific parameter $\boldsymbol{\Theta}^m$ and hyper parameter $\mathbf{h}^m$ in Section II-A and B. As a consequence of the shared component densities, the two separate responsibilities, $\phi_{ni,t}^m, m \in \{0,1\}$, obtained from the forward-backward algorithm in two separate HMMs are combined into a single responsibility, $\phi_{ni,t}$, as follows:

$$\phi_{ni,t} \sum_{m\in\{0,1\}} \phi_{mni}^M \phi_{ni,t}^m. \tag{15}$$

The posterior density estimations are detailed for the two observation models in the supplementary materials. The graphical model of MiHMM2 in Fig. 3 is similar to that of MiHMM1 in Fig. 2 except that there is only one observation model in Fig. 3, which is shared by both HMMs.

Having estimated the parameters of the posterior densities of the MiHMM model, a new bag $\mathbf{X}_n$, can be classified using a likelihood ratio test calculated using the posterior predictive densities of the model parameters, which is detailed in the supplementary materials. Section III evaluates the efficacy of the proposed MiHMM approaches with the standard EM-HMM on various data sets.

[1]This paper has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the authors. This includes a pdf document detailing the parameter inference and the prediction of the proposed model. The material is 318 kB in size.

III. Experimental Analysis on MiHMM Data sets

The proposed MiHMM approaches are evaluated against the standard EM-HMM and the NPBMIL approach on two synthetic data sets and two landmine data sets. Motivated by the two landmine data sets, the two synthetic data sets are generated based on the two specific component densities, i.e., MN and MVN. In addition to comparing the receiver operating characteristics (ROC) performance of these algorithms on different data sets using 10-fold cross-validation, the following sections also analyze the parameters learned by these algorithms and the H_1 instances in a positive bag inferred by the proposed approach. For training the standard EM-HMM model, all instances in a positive bag were considered H_1 instances and used for training. This corresponds to a lack of information regarding the actual target location in depth. Of course, if the true target location is known, then that information can be incorporated into the HMM training. Several approaches exist to attempt to estimate the object depth; however, a comparison of these is beyond the scope of this paper.

A. Synthetic Data sets

For each synthetic data set, there are 200 negative bags, 200 positive bags, 15 time samples per instance, 10 instances per bag, and one H_1 instance per positive bag.

1) Multinomial Synthetic Data: In the MN data set, time samples in both H_0 and H_1 HMMs are generated from one of the four multinomial components. The HMMs have different transition probabilities, $\{\mathbf{A}^0, \mathbf{A}^1\}$, and initial state mixing probabilities, $\{\mathbf{\Pi}^0, \mathbf{\Pi}^1\}$. The four multinomial components are motivated by the four-state representation of a typical landmine signature in time-domain GPR data. Each multinomial component is a $[16 \times 1]$ probability vector, $\mathbf{p}_k, k \in \{1, \ldots, 4\}$, drawn from a uniform Dirichlet distribution. The transition probabilities are defined as follows: For $k \in \{1, \ldots, 4\}$, $\{\mathbf{A}^0_{k\bullet}\} = \{[.7, 0, 0, .3], [.3, .7, 0, 0], [0, .3, .7, 0], [0, 0, .3, .7]\}$ and $\{\mathbf{A}^1_{k\bullet}\} = \{[.7, .3, 0, 0], [0, .7, .3, 0], [0, 0, .7, .3], [.3, 0, 0, .7]\}$. The initial state mixing probabilities, $\{\mathbf{\Pi}^0, \mathbf{\Pi}^1\}$, are uniform. The choice of the number of time samples per instance and the dimension of the multinomial probability vector are motivated by the discrete landmine data described in Section III-B1. Fig. 4(a) shows that the MiHMMs outperform the standard EM-HMM by accurately modeling the H_1 HMM. Between the two MiHMMs, MiHMM2 does slightly better than MiHMM1 because although the data was generated assuming a single observation model, only MiHMM2 assumes a single shared observation model. MiHMM1 is required to estimate additional parameters with the same amount of data, resulting in a slight performance loss. It is also not surprising to see NPBMIL perform poorly because although the data was generated using an underlying Markov model, NPBMIL cannot adequately model the sequential nature of the data.

2) Multivariate Normal Synthetic Data: The MVN synthetic data set is motivated by the HOG landmine data described in Section III-B2. Similar to Section III-A1, in the MVN synthetic data set, time samples in both H_0 and H_1 HMMs are generated from a GMM composed of four components, represented by the means and covariances, $\{\boldsymbol{\mu}_k, \mathbf{\Sigma}_k\}, k \in \{1, \ldots, 4\}$. The means are located at [1,1], $[-1, 1]$, $[1, -1]$, and $[-1, -1]$ and all Gaussians have unit covariances. The transition probabilities and the initial state mixing probabilities are similar to Section III-A1. Similar to the results in Section III-A1, the ROC performance in Fig. 4(b) shows that HMM-based approaches outperform NPBMIL and MiHMMs outperform the standard EM-HMM.

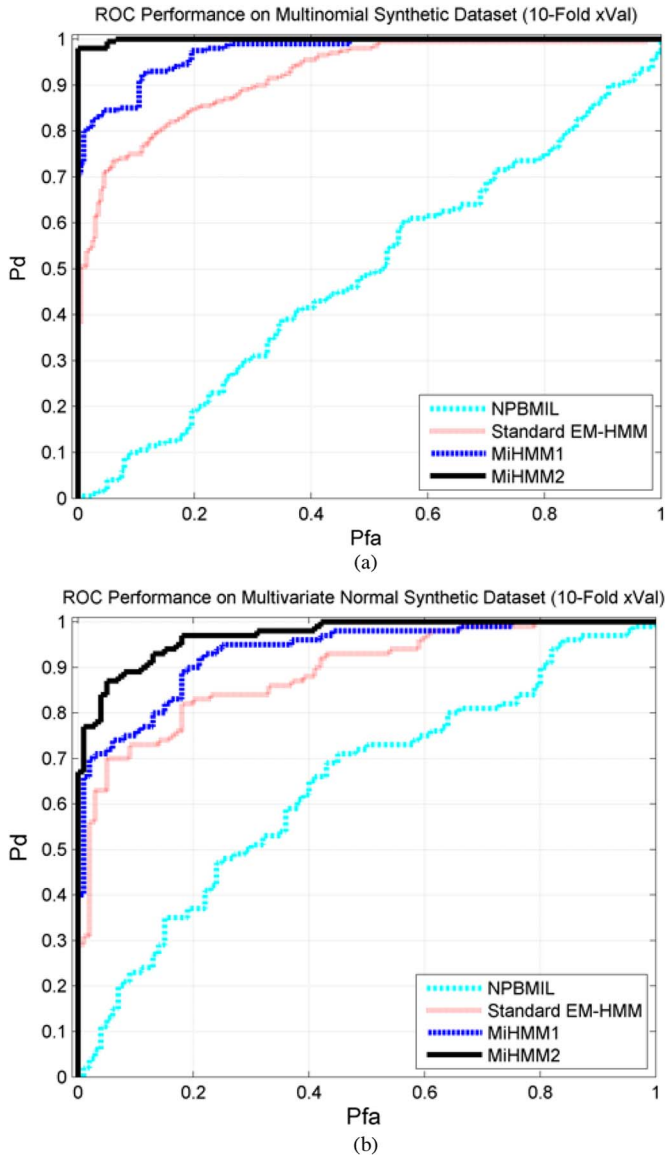

Fig. 4. ROC performance comparison of NPBMIL, standard EM-HMM, MiHMM1 (different observation models), and MiHMM2 (shared observation model) on (a) Multinomial and (b) Multivariate Normal synthetic data sets.

B. Landmine Data sets

This work implements the proposed MiHMM approaches for GPR-based landmine detection using two landmine data sets and evaluates their efficacy against the standard EM-HMM.

1) Discrete Features: The discrete features were obtained by extracting gradient features from the time-domain GPR data using the software provided by the authors of [11]. These

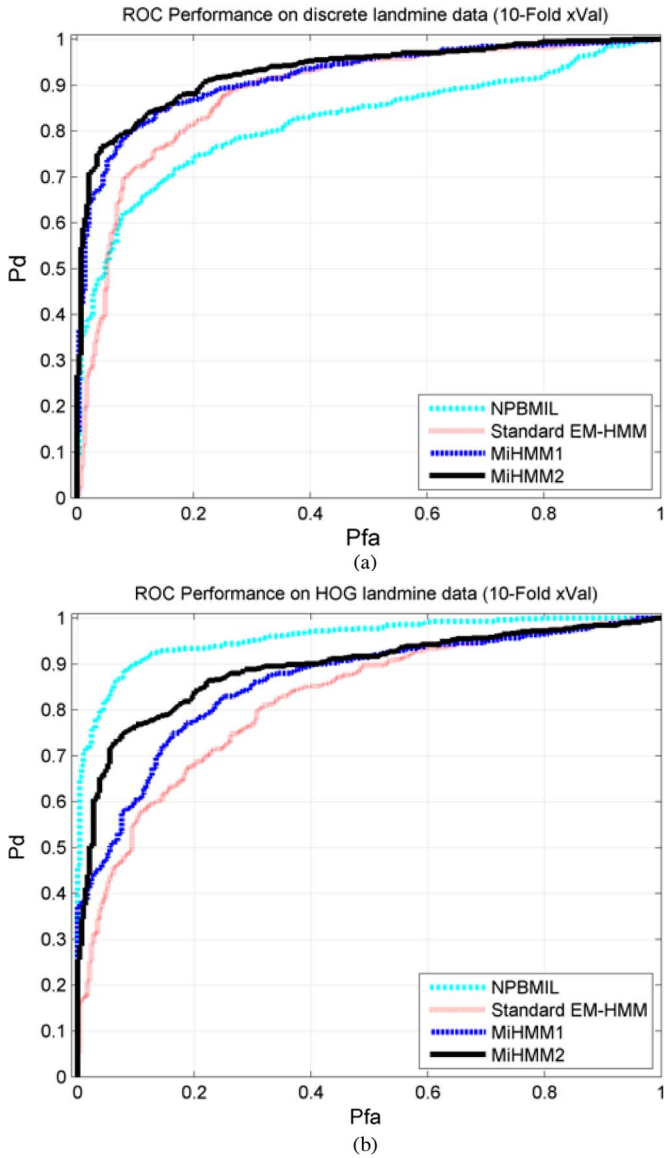

Fig. 5. ROC performance comparison of NPBMIL, standard EM-HMM, MiHMM1 (different observation models), and MiHMM2 (shared observation model) on (a) landmine discrete data [11] (b) landmine HOG data.

features describe the gradient information in the diagonal and antidiagonal edges. The HMM features extracted from multiple depth bins of a [downtrack $\times$ depth] GPR data at an alarm location can be considered a bag of instances. Each bag is composed of 22 instances and each instance is composed of 15 samples, where each sample is represented by one of the 16 possible discrete observations. A collection of bags along with the ground truth of buried threats in training data provide a data set suitable for multiple-instance learning, mainly because the depth locations of the buried targets cannot be easily obtained.

In this paper, the discrete HMM features are extracted from a total of 864 alarms, consisting of 574 targets. The ROC comparison in Fig. 5(a) demonstrates that MiHMMs outperform the standard EM-HMM. Since the discrete features were specifically generated to model the GPR data sequentially, it is not surprising that HMM-based approaches outperform NPBMIL on discrete features.

2) HOG Features: The histogram of oriented (HOG) features [36] are extracted from time-domain GPR data. Originally developed as a feature extraction technique for pedestrian detection [36], HOG features have been extensively used as a promising feature descriptor in many other computer vision applications [39], and recently for GPR-based landmine detection [40], [41]. HOG features represent distribution of pixel-wise gradients that are aggregated over a region of pixels, known as cells, and smoothed over groups of cells, known as blocks [36]. Intuitively, these features describe relative changes in image regions. Similar to Section III-B1, the HOG features extracted from multiple depth bins of a [downtrack $\times$ depth] slice of GPR data can be considered a bag of instances. For the HOG features, each bag is composed of 12 instances, and each instance, corresponding to an HMM sequence, is composed of 20 observation samples, where each sample is represented by a HOG-block [36], as shown in Fig. 6.

Similar to the discrete features in Section III-B1, the HOG features are also extracted from a total of 864 alarms, consisting of 574 targets. Once again, the ROC comparison in Fig. 5(b) exhibits the improvement of the MiHMM approaches over the standard EM-HMM, however for the HOG features, the NPBMIL approach outperforms all HMM techniques. Since the HOG features were generated to represent 2-D patches of GPR data without incorporating any sequential information, it is not unreasonable for NPBMIL to outperform the HMM-based approaches on HOG features.

The MiHMM approaches also provide the mixing proportion of H_1 and H_0 instances in a positive bag and can also infer the H_1 instances in a positive bag. Fig. 6(a) and (b) shows examples of 2-D GPR data with buried target and false alarm, respectively along with their corresponding HOG features per instance, and the probability of each instance to be a H_1 instance. In Fig. 6(a), the HOG features in depth bins 5–8 represent typical rising and falling edges of a buried landmine, and clearly, these depth bins have very high probability of being a H_1 instance. In contrast, all depth bins in a false alarm example in Fig. 6(b) have very low probability of being a H_1 instance.

C. Discussion

For the results presented in Section III-A and B, the number of mixture components was optimized for all approaches. While the choice of the number of components affects the performance of all approaches, the performance is fairly robust for a wide range of values, showing that a coarse optimization is sufficient in most cases. In terms of computational complexity, the MiHMMs take approximately twice as long to train compared with the standard EM-HMM. While the standard EM-HMM divides the training data into two different sets to train two HMMs separately, the MiHMMs train a hierarchical mixture of two HMMs simultaneously so that the MiHMMs require more memory during training than the standard EM-HMM. Among the two MiHMMs, MiHMM2 is faster than MiHMM1 because MiHMM2 has fewer parameters to be estimated. The experiments with synthetic data sets exhibit that the MiHMMs are capable of learning a more accurate model for H_1 instances, and as a result, the MiHMMs are more robust

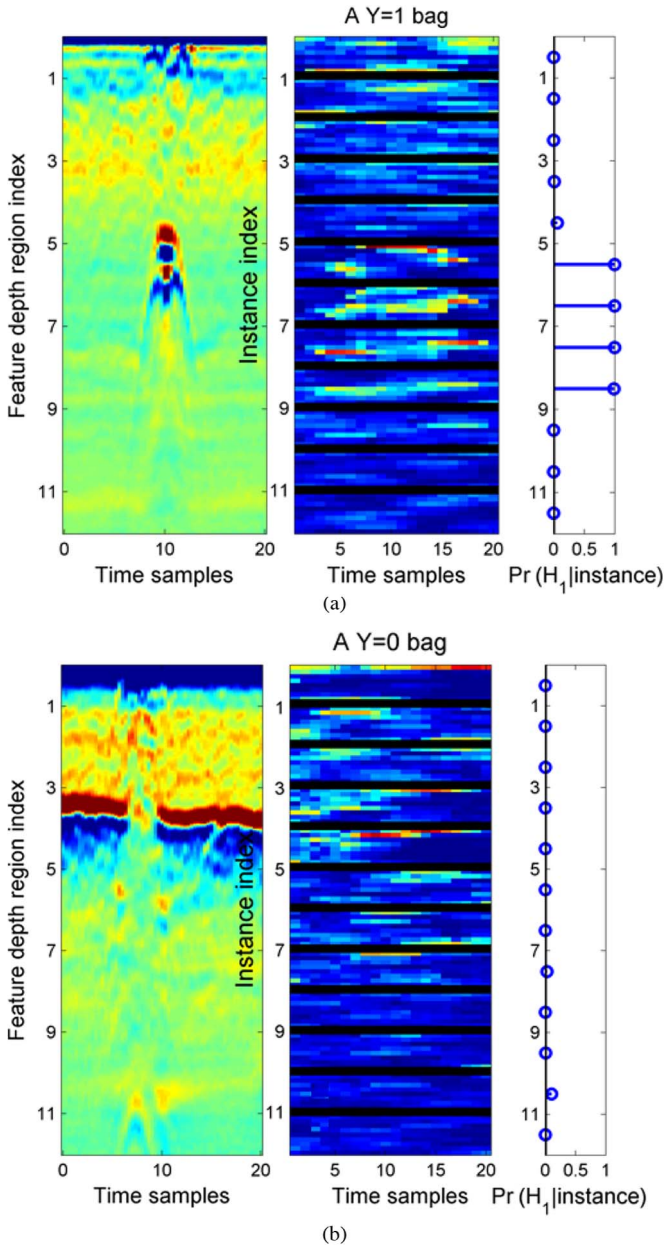

Fig. 6. Examples of 2-D GPR data with (a) buried target and (b) false alarm along with their corresponding HOG descriptors per instance and the probability of each instance to be a H_1 instance. The output $\Pr(H_1|\text{instance})$ illustrate that the proposed model is capable of identifying the positive instances in the bag, since the $\Pr(H_1|\text{instance})$ is large at the same depth regions where the target is present.

when there are fewer H_1 instances in a positive bag. The standard EM-HMM approach is not able to learn an accurate H_1 model and instead acts mostly as an anomaly detector. However, in data sets with a large number of H_1 instances per positive bag, the standard EM-HMM may perform comparably to the MiHMMs. The MiHMMs can also infer the probability of a H_1 instance in a positive bag. As mentioned in [42], a MIL approach is most beneficial when it is necessary to learn the H_1 instances in addition to classifying the positive bags from negative bags. For example, in the drug activity problem [43], it may be beneficial to draw examples of positive molecules and for subsurface threat detection, the depth information of the target can provide valuable insight to the threat removal system. In addition, the MiHMM approaches are easily generalizable to different data sets by using an appropriate observation model.

1) MiHMM1 Versus MiHMM2: Between the two MiHMMs, the choice of the model should be motivated by the data. MiHMM2 is more suitable in the landmine detection data sets implemented in this work because we expect the background portions of the data in both H1 and H0 instances to be similar. The MiHMM2 model enables these background states to be shared across the H1 and H0 data sets, which reduces the total number of parameters, and enables more robust parameter estimation. Since there are fewer parameters to estimate in the MiHMM2 model, we expect it to outperform the MiHMM1 model when the amount of training data is limited. As expected, as shown in Fig. 4(a) and (b), in the synthetic data sets, with very few data points, the performance improvement of MiHMM2 over MiHMM1 is more substantial than in the actual landmine data. However, for applications where it is unreasonable to learn a set of shared states between the two hypotheses, MiHMM1 should be preferred.

IV. Conclusion and Future Work

This work has outlined a generative model for binary classification of multiple-instance time-series data. The proposed MiHMM approaches extend a standard HMM approach to an HMM in a MIL framework that inherently incorporates the ambiguity in the individual sample labels. The proposed MiHMM approaches are evaluated against a standard EM-HMM approach on synthetic and landmine detection data sets. Results indicate that the proposed approach performs better than the standard EM-HMM technique by modeling the label uncertainty of the MIL nature of the data. The proposed parameter learning method for the MiHMM is based on variational Bayesian methods, which are fast and computationally efficient. For future work, it may be beneficial to explore the possibility of performing Markov Chain Monte Carlo (MCMC) inference to obtain better parameter estimates at the expense of additional computation and online VB to further limit the memory requirements of the algorithm. Since the target and the nontarget HMMs in the standard EM-HMM approach were trained independently, whereas a hierarchical mixture of the two HMMs was trained simultaneously in the proposed approach, it will be interesting to compare the proposed approach with a discriminately trained HMM. Furthermore, it will also be interesting to compare the performance of the proposed approach on hyperspectral data of the same target site.

Acknowledgment

The authors would like to thank the U.S. Army RDECOM CERDEC Night Vision and Electronic Sensors Directorate for providing supporting Grants (W911NF-09-1-0487 and W911NF-06-1-0357) through the Army Research Office and J. Bolton, J. Wilson, and P. Gader at the University of Florida.

REFERENCES

[1] Landmine and Cluster Munition Monitor, *Landmine Monitor Report 2012*, Landmine and Cluster Munition Monitor, Geneva, Switzerland, 2012. [Online]. Available: http://www.the-monitor.org/lm/2012/resources/Landmine_Monitor_2012.pdf

[2] J. MacDonald and J. R. Lockwood, "Alternatives for Landmine Detection," RAND Corp., Santa Monica, CA, USA, 2003.

[3] P. Torrione, C. Throckmorton, and L. Collins, "Performance of an adaptive feature-based processor for a wideband ground penetrating radar system," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 2, pp. 644–658, Apr. 2006.

[4] P. Torrione and L. Collins, "Application of texture feature classification methods to landmine/clutter discrimination in off-lane GPR data," in *Proc. IEEE IGARSS*, 2004, vol. 3, pp. 1621–1624.

[5] Q. Zhu and L. Collins, "Application of feature extraction methods for landmine detection using the Wichmann/Niitek ground-penetrating radar," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 1, pp. 81–85, Jun. 2005.

[6] P. Torrione and L. Collins, "Texture features for antitank landmine detection using ground penetrating radar," *IEEE Trans. Geoscie. Remote Sens.*, vol. 45, no. 7, pp. 2374–2382, Jul. 2007.

[7] T. Savelyev, L. Van Kempen, H. Sahli, J. Sachs, and M. Sato, "Investigation of time ndash;frequency features for GPR landmine discrimination," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 1, pp. 118–129, 2007.

[8] J. Wilson, P. Gader, W. H. Lee, H. Frigui, and K. Ho, "A large-scale systematic evaluation of algorithms using ground-penetrating radar for landmine detection and discrimination," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 8, pp. 2560–2572, Aug. 2007.

[9] K. Ho, L. Carin, P. D. Gader, and J. N. Wilson, "An investigation of using the spectral characteristics from ground penetrating radar for landmine/clutter discrimination," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 4, pp. 1177–1191, Apr. 2008.

[10] H. Frigui and P. Gader, "Detection and discrimination of land mines in ground-penetrating radar based on edge histogram descriptors and a possibilistic k-nearest neighbor classifier," *IEEE Trans. Fuzzy Syst.*, vol. 17, no. 1, pp. 185–199, Feb. 2009.

[11] P. Gader, M. Mystkowski, and Y. Zhao, "Landmine detection with ground penetrating radar using hidden Markov models," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 6, pp. 1231–1244, Jun. 2001.

[12] Y. Zhao, P. Gader, P. Chen, and Y. Zhang, "Training DHMMS of mine and clutter to minimize landmine detection errors," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 5, pp. 1016–1024, May 2003.

[13] K. C. H. Frigui and H. P. Gader, "Real-time landmine detection with ground-penetrating radar using discriminative and adaptive hidden Markov models," *EURASIP J. Appl. Signal Process.*, vol. 12, pp. 1867–1885, 2005.

[14] H. Frigui, O. Missaoui, and P. Gader, Landmine Detection Using Discrete Hidden Markov Models With Gabor Features, pp. 65 532A-1–65 532A-10, 2007. [Online]. Available: +http://dx.doi.org/10.1117/12.722241

[15] P. A. Torrione and L. Collins, Image Segmentation Techniques for Improved Processing of Landmine Responses in Ground-Penetrating Radar Data, pp. 655 329–655 329-12, 2007. [Online]. Available: http://dx.doi.org/10.1117/12.718128

[16] Y. Dong *et al.*, "Multi-aspect detection of surface and shallow-buried unexploded ordnance via ultra-wideband synthetic aperture radar," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 6, pp. 1259–1270, Jun. 2001.

[17] X. Liao, P. Runkle, and L. Carin, "Identification of ground targets from sequential high-range-resolution radar signatures," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 4, pp. 1230–1242, Oct. 2002.

[18] S. Ji, X. Liao, and L. Carin, "Adaptive multiaspect target classification and detection with hidden Markov models," *IEEE Sens. J.*, vol. 5, no. 5, pp. 1035–1042, Oct. 2005.

[19] O. Maron and T. Lozano-Perez, "A framework for multiple-instance learning," in *Proc. NIPS*, 1998, vol. 10, pp. 570–576.

[20] O. Maron and A. L. Ratan, "Multiple-instance learning for natural scene classification," in *Proc. 55th ICML*, San Francisco, CA, USA, 1998, pp. 341–349, Morgan Kaufmann Publishers Inc..

[21] Q. Zhang and S. A. Goldman, "EM-DD: An improved multiple-instance learning tecnique," in *Proc. NIPS*, 1998, vol. 13, pp. 1073–1080.

[22] W. Y. J. F. Qi Zhang and S. A. Goldman, "Content-based image retrieval using multiple-instance learning," in *Proc. 19th Int. Conf. Mach. Learn.*, 2002, pp. 682–689.

[23] T. H. S. Andrews and I. Tsochantaridis, "Support vector machines for multiple-instance learning," in *Proc. NIPS*, 2002, pp. 561–568.

[24] Y. Chen and J. Z. Wang, "Image categorization by learning and reasoning with regions," *J. Mach. Learn. Res.*, vol. 5, pp. 913–939, Dec. 2004.

[25] Y. Chen, J. Bi, and J. Wang, "Miles: Multiple-instance learning via embedded instance selection," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 28, no. 12, pp. 1931–1947, Dec. 2006.

[26] Z. Fu, A. Robles-Kelly, and J. Zhou, "Milis: Multiple instance learning with instance selection," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 33, no. 5, pp. 958–977, May 2011.

[27] P. Viola, J. C. Platt, and C. Zhang, "Multiple instance boosting for object detection," in *Proc. NIPS*, 2006, pp. 1419–1426, MIT Press.

[28] S. Ali and M. Shah, "Human action recognition in videos using kinematic features and multiple instance learning," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 32, no. 2, pp. 288–303, Feb. 2010.

[29] B. Babenko, M.-H. Yang, and S. Belongie, "Robust object tracking with online multiple instance learning," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 33, no. 8, pp. 1619–1632, Aug. 2011.

[30] J. Bolton, P. Gader, H. Frigui, and P. Torrione, "Random set framework for multiple instance learning," *Inf. Sci.*, vol. 181, no. 11, pp. 2061–2070, Jun. 2011.

[31] V. C. Raykar, B. Krishnapuram, J. Bi, M. Dundar, and R. B. Rao, "Bayesian multiple instance learning: Automatic feature selection and inductive transfer," in *Proc. 25th ICML*, New York, NY, USA, 2008, pp. 808–815.

[32] J. Bolton and P. Gader, "Application of multiple-instance learning for hyperspectral image analysis," *IEEE Geosci. Remote Sens. Lett.*, vol. 8, no. 5, pp. 889–893, Sep. 2011.

[33] S. Yuksel, J. Bolton, and P. Gader, "Landmine detection with multiple instance hidden Markov models," in *Proc. IEEE Int. Workshop MLSP*, 2012, pp. 1–6.

[34] R. Rahmani, S. A. Goldman, H. Zhang, S. R. Cholleti, and J. E. Fritts, "Localized content-based image retrieval," *IEEE Trans. Pattern Anal. Mach. Intel.*, vol. 30, no. 11, pp. 1902–1912, Nov. 2008.

[35] A. Manandhar, K. D. Morton, L. M. Collins, and P. A. Torrione, "Multiple Instance Learning for Landmine Detection Using Ground Penetrating Radar," in *Proc. SPIE*, J. H. Holloway, Jr. and J. T. Broach, Eds., 2012, vol. 8357, p. 835 721.

[36] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Comput. Soc. Conf. CVPR*, 2005, vol. 1, pp. 886–893.

[37] A. Q. Vaclav Smidl, *The Variational Bayes Method in Signal Processing*. Berlin, Germany: Springer-Verlag, 2006.

[38] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proc. IEEE*, vol. 77, no. 2, pp. 257–286, Feb. 1989.

[39] D. A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*. Upper Saddle River, NJ, USA: Prentice Hall, 2012, Professional Technical Reference.

[40] T. C. Havens *et al.*, Multiple Kernel Learning for Explosive Hazard Detection in Forward-Looking Ground-Penetrating Radar, pp. 83 571D-1–83 571D-15, 2012. [Online]. Available: http://dx.doi.org/10.1117/12.920482

[41] P. Torrione, K. Morton, R. Sakaguchi, and L. Collins, Histograms of Oriented Gradients for Landmine Detection in Ground-Penetrating Radar Data, vol. 52, no. 3, pp. 1539–1550, Mar. 2014.

[42] S. Ray and M. Craven, "Supervised versus multiple instance learning: An empirical comparison," in *Proc. 22nd ICML*, New York, NY, USA, 2005, pp. 697–704.

[43] R. E. L. Thomas G. Dietterich and T. Lozano-Perez, "Solving the multiple instance problem with axis-parallel rectangles," *Artif. Intel.*, vol. 89, no. 1/2, pp. 31–71, Jan. 1997.

Achut Manandhar was born in Kathmandu, Nepal. He received the B.S.E.E.(Hons.) degree from Fairleigh Dickinson University, Teaneck, NJ, USA, in 2008 and the M.S.E.E. degree from Duke University, Durham, NC, USA, in 2012. He is currently working toward the Ph.D. degree in electrical engineering at Duke University with research focused on statistical signal processing methods for landmine detection and video analysis.

Peter A. Torrione received the B.S.E.E. degree from Tufts University, Medford, MA, USA, in 1999, and the M.S.E.E. and Ph.D. degrees in electrical engineering from Duke University, Durham, NC, USA, in 2002 and 2008.

He is currently an Assistant Research Professor with Duke University, where his interests include physics-based statistical signal processing, pattern recognition, and machine learning.

Leslie M. Collins (M'96–SM'01) was born in Raleigh, NC,USA. She received the B.S.E.E. degree from the University of Kentucky, Lexington, KY, USA, in 1985, and the M.S.E.E. and Ph.D. degrees in electrical engineering, both from the University of Michigan, Ann Arbor, MI, USA, in 1986 and 1995, respectively.

From 1986 to 1990, she was a Senior Engineer with the Westinghouse Research and Development Center, Pittsburgh, PA, USA. She joined Duke University, Durham, NC, USA, in 1995 as an Assistant Professor and was promoted to Associate Professor in 2002. Her current research interests include incorporating physics-based models into statistical signal processing algorithms, and she is working toward applications in subsurface sensing as well was enhancing speech understanding by hearing impaired individuals.

Dr. Collins is a member of Tau Beta Pi, Eta Kappa Nu, and Sigma Xi.

Kenneth D. Morton (M'01) was born in York, PA, USA, in 1982. He received the B.S. degree in electrical and computer engineering from The University of Pittsburgh, Pittsburgh, PA, in 2004 and the M.S. and Ph.D. degrees in electrical and computer engineering from Duke University, Durham, NC, USA, in 2006 and 2010, respectively.

He is currently an Assistant Research Professor with Duke University with research focused on the development of Bayesian statistical models for a variety of signal processing and machine learning applications, including acoustic signal classification and landmine detection.

Dr. Morton is a member of Tau Beta Pi and Eta Kappa Nu.