Boosted Markov Networks for Activity Recognition

Tran The Truyen [1], Hung Hai Bui [2], Svetha Venkatesh [3]
Department of Computing, Curtin University of Technology
GPO Box U 1987, Perth, WA, Australia.
[1,3] {trantt2,svetha}@cs.curtin.edu.au, [2]bui@ai.sri.com

Abstract

We explore a framework called boosted Markov networks to combine the learning capacity of boosting and the rich modeling semantics of Markov networks and applying the framework for video-based activity recognition. Importantly, we extend the framework to incorporate hidden variables. We show how the framework can be applied for both model learning and feature selection. We demonstrate that boosted Markov networks with hidden variables perform comparably with the standard maximum likelihood estimation. However, our framework is able to learn sparse models, and therefore can provide computational savings when the learned models are used for classification.

1. Introduction

Recognising human activities using sensors is currently a major challenge in research. Typically, the information extracted directly from sensors is either not discriminative enough or is too noisy to infer activities occurring in the scene. Human activities are complex and evolve dynamically over time. Temporal probabilistic models such as hidden Markov models (HMMs) and dynamic Bayesian networks (DBNs) have been the dominant models used to solve the problem [1], [4], [19]. However, these models make a strong assumption in the generative process by which the data is generated in the model. This makes the representation of complex sensor data very difficult, and possibly results in large models.

Markov networks (MNs) (also known as Markov random fields) offer an alternative approach, especially in form of conditional random fields (CRFs) [10]. In CRFs, the observation is not modelled, and so we have the freedom to incorporate overlapping features, multiple sensor fusion, and long-range dependencies into the model. The discriminative nature and the underlying graphical structure of the MNs make it especially suitable to the problem of human activity recognition.

Boosting is a general framework to gradually improve the performance of the weak learner (which can be just slightly better than a random guess). A popular version called AdaBoost [6], [14], [15] forces the weak learner to focus more on hard-to-learn examples from the examples seen so far. The final strong learner is the weighted sum of all weak learners added to this ensemble in each iteration. Most of the work so far with boosting involves only unstructured output, except for a few occasions, such as the work in [2], [5], [18].

We are motivated to use boosting for parameter estimation of Markov networks (which we call boosted Markov networks (BMNs)), as recent results have shown the close relationship between boosting and the maximum likelihood estimation (MLE) [8], [11]. Furthermore, we use the inherent capacity of boosting for feature selection integrated with learning. We are motivated by studies that the typical log-linear models imposed on Markov networks can easily be overfitted with little data and many irrelevant features, but can be overcome by the use of explicit feature selection, either independently or integrated with learning (e.g. see [13]).

Previous work [2], [5], [18] have explored the BMNs showing promising results. The BMNs integrate the discriminative learning power of boosting and the rich semantics of the graphical model of Markov networks. In this paper, we further explore alternative methods for this approach. To handle hidden variables as in the standard MLE, we extend the work by Altun *et al.* [2]. This is a variant of the multiclass boosting algorithm called AdaBoost.MR [6], [15]. We also suggest an approximation procedure for the case of intractable output structures. The proposed framework is demonstrated through our experiments in which boosting can provide a comparable performance to MLE. However, since our framework uses sparse features it has the potential to provide computational savings during recognition.

The novelty of the paper is two-fold: a) We present the first work in applying boosting for activity recognition, and b) we derive a boosting procedure for structured models with missing variables and use a parameter update based on quadratic approximation instead of loose upper bounds as in [2] to speed up the learning.

The organisation of the paper is as follows. The next section reviews related work. Section 3 presents the concept of boosted Markov networks (BMNs) detailing how boosting is employed to learn parameters of Markov networks. Section 4 goes into the details of efficient computation for tree-structured and general networks. We then describe our experiments and results in applying the model to video-based human activity recognition in an indoor environment. The last section discusses any remaining issues.

2. Related work

Our work is closely related to that in [5], where boosting is applied to learn parameters of the CRFs using gradient trees [7]. The objective function is the log-likelihood in the standard MLE setting, but the training is based on fitting

0-7803-9399-6/05/$20.00 © 2005 IEEE

regression trees in a stage-wise fashion. The final decision function is in the form of a linear combination of regression trees. [5] employs functional gradients for the log-loss in a manner similar to LogitBoost [8], whilst we use the original gradients of the exponential loss of AdaBoost [6], [14], [15]. Thus [5] learns the model by maximising the likelihood of the data, while we are motivated by minimising the errors in the data directly. Moreover, [5] indirectly solves the structured model learning problem of MLE by converting the problem to a standard unstructured learning problem with regression trees. In contrast, we solve the original structured learning problem directly without the structured-unstructured conversion. In addition, the paper in [5] does not incorporate hidden variables as our work does.

Another work, [18], integrates the message passing algorithm of belief propagation (BP) with a variant of LogitBoost [8]. Instead of using the per-network loss as in [5], the authors of [18] employ the per-label loss (e.g. see [3] for details of the two losses), that is they use the marginal probabilities. Similar to [5], [18] also converts the structured learning problem into a more conventional unstructured learning problem. The algorithm thus alternates between a message passing round to update the local per-label log-losses, and a boosting round to update the parameters. However, as the BP is integrated in the algorithm, it is not made clear on how to apply different inference techniques when the BP fails to converge in general networks. It is also unclear on how to extend the method to deal with hidden variables.

There have been a number of attempts to exploit the learning power of boosting applied to structured models other than Markov networks, such as dynamic Bayesian networks (DBNs) [9], Bayesian network classifier [21], and HMMs [20].

3. Boosted Markov networks

A. Markov networks

We are interested in learning a structured model in which inference procedures can be performed. A typical inference task is decoding, e.g. to find the most probable label set (configuration) $y*$ for the network given an observation x $y* = \arg\max_y p(y|x)$. For a single node network, this is often called the classification problem. For a network of N nodes, the number of configurations is exponentially large. We assume the conditional exponential distribution (a.k.a. log-linear model)

$$p(y|x) = \frac{1}{Z(x)} \exp(F(x,y)) \tag{1}$$

where $Z(x) = \sum_y \exp(F(x,y))$ is the normalisation factor, and $F(x,y) = \sum_k \lambda_k f_k(x,y)$. $\{f_k\}$ are features or weak hypotheses in boosting literature. This conditional model is known as the conditional random fields (CRFs) [10]. The decoding reduces to $y* = \arg\max_y F(x,y)$.

Often, in Markov networks the following decomposition is used

$$f_k(x,y) = \sum_c f_k(x,y_c) \tag{2}$$

where c is the index of the clique defined by the structure of the network. This decomposition is essential to obtain a factorised distribution, which is vital for efficient inference in tree-like structures using dynamic programming.

Denote by $y = (v,h)$ the visible and hidden variables (which are represented by nodes in the network). Given n i.i.d observations $\{x_i\}_{i=1}^n$, maximum likelihood learning in MNs minimises the log-loss

$$L_{log} = -\sum_i \log p(v^i|x^i) = \sum_i \log \frac{1}{p(v^i|x^i)} \tag{3}$$

B. Exponential loss for incomplete data

We view the activity labeling and segmentation problems as classification where the number of distinct classes is exponentially large, i.e. $|Y|^N$, where $|Y|$ is the size of label set, and N is the number of nodes in the Markov network. Following the development by Altun et al. [2], we define a new *expected ranking loss* [15] to incorporate hidden variables as follows

$$L_{rank} = \sum_i \sum_h p(h|v^i,x^i) \sum_{v \neq v^i} \delta[\Delta F(x^i,v,h) > 0] \tag{4}$$

where $\Delta F(x^i,v,h) = F(x^i,v,h) - F(x^i,v^i,h)$, and $\delta[z]$ is the indicator function of whether the predicate z is true. This rank loss captures the expectation of moments in which the classification is wrong because if it is true, then $\max_v F(x^i,v,h) = F(x^i,v^i,h)$ implying $F(x^i,v,h) < F(x^i,v^i,h) \forall v \neq v_i$.

However, it is much more convenient to deal with a smooth convex loss and thus we formulate an upper bound of the rank loss, e.g. the exponential loss

$$L_{exp} = \sum_i \sum_h p(h|v^i,x^i) \sum_v \exp(\Delta F(x^i,v,h)) \tag{5}$$

It is straightforward to check that (5) is indeed an upper bound of (4). It can be seen that (4) includes the loss proposed in [2] as a special case when all variables are observed, i.e. $y = v$. It is essentially an adapted version of AdaBoost.M2 proposed in [6].

A difficulty with this formulation is that we do not know the true conditional distribution $p(h|v^i,x^i)$. First, we approximate it by the learned distribution at the previous iteration. Thus, the conditional distribution is updated along the way, starting from some guessed distribution, for example, a uniform distribution. Second, we assume the log-linear model as in (1), leading to

$$\sum_v \exp(\Delta F(x^i,v,h)) = \frac{\sum_v \exp(F(x^i,v,h))}{\exp(F(x^i,v^i,h))} = \frac{1}{p(v^i|h,x^i)}$$

which can be fed into (5) to obtain

$$L_{exp} = \sum_i \sum_h \frac{p(h|v^i,x^i)}{p(v^i|h,x^i)} = \sum_i \frac{1}{p(v^i|x^i)} \tag{6}$$

We can notice the similarity between the exponential loss in (6) and the log-loss in (3) as log(.) is a monotonically increasing function. The difference is the exponential scale used in (6) with respect to features $\{f_k\}$ compared to the linear scale in (3).

C. Boosting-based learning

The typical boosting process has many rounds, each of which selects the best weak hypothesis and finds the weight for this hypothesis to minimise the loss. Translated in our context, the boosting-based learning searches for the best feature f_j and its coefficient to add to the ensemble $F^{t+1} = F^t + \alpha^t f_j$ so that the loss in (5) is minimised.

$$(\alpha^t, j) = \arg\min_{\alpha,k} L_{exp}(t, \alpha, k), \text{ where} \quad (7)$$

$$L_{exp}(t, \alpha, k) = \sum_i E_{h|v^i,x^i,t}[\sum_v \exp(\Delta F^{i,t} + \alpha \Delta f_k^i)]$$

where $E_{h|v^i,x^i,t}[z(h)] = \sum_h p(h|v^i, x^i, t)z(h)$; $F^{i,t}$ and f_k^i are shorthands for $F^t(x^i, v, h)$ and $f_k(x^i, v, h)$, respectively. Note that this is just an approximation to (5) because we fix the conditional distribution $p(h|v^i, x^i, t)$ obtained from the previous iteration. However, this still makes sense since the learning is incremental, and thus the estimated distribution will get closer to the true distribution along the way. Indeed, this captures the essence of boosting: during each round, the weak learner selects the weak hypothesis that best minimises the following loss over the weighted data distribution (see [15])

$$(\alpha^t, j) = \arg\min_{\alpha,k} \sum_i \sum_{v,h} D(i, v, h, t) \exp(\alpha \Delta f_k^i) \quad (8)$$

where the weighted data distribution $D(i, v, h, t) = p(h|v^i, x^i, t) \exp(\Delta F^{i,t})/Z(i, t)$ with $Z(i, t)$ being the normalising constant. Since the data distribution does not contain α, (8) is identical to (7) up to a constant.

D. Beam search

It should be noted that boosting is a very generic framework to boost the performance of the weak learner. Thus we can build more complex and stronger weak learners by using some ensemble of features and then later fit them into the boosting framework. However, here we stick to simple weak learners, which are features, to make the algorithm compatible with the MLE.

We can select a number of top features and associated coefficients that minimise the loss in (8) instead of just one feature. This is essentially a beam search with specified beam size S.

E. Regularisation

We employ the l_2 regularisation term to make it consistent with the popular Gaussian prior used in conjunction with the MLE of Markov networks. It also maintains the convexity of the original loss. The regularised loss becomes

$$L_{reg} = L_{non-reg} + \sum_k \frac{\lambda_k^2}{2\sigma_k^2} \quad (9)$$

where $L_{non-reg}$ is either L_{log} for MLE in (3) or L_{exp} for boosting in (5). Note that the regularisation term for boosting does not have the Bayesian interpretation as in the MLE setting but is simply a constraint to prevent the parameters from growing too large, i.e. the model fits the training data too well, which is clearly sub-optimal for noisy and unrepresentative data. The effect of regularisation can be numerically very different for the two losses, so we cannot expect the same σ for both MLE and boosting.

4. Efficient computation

Straightforward implementation of the optimisation in (7) or (8) by sequentially and iteratively searching for the best features and parameters can be impractical if the number of features is large. This is partly because the objective function, although it can be tractable to compute using dynamic programming in tree-like structures, is still expensive. We propose an efficient approximation which requires only a few vectors and an one-step evaluation. The idea is to exploit the convexity of the loss function by approximating it with a convex quadratic function using second-order Taylor's expansion. The change due to the update is approximated as

$$\Delta J(\alpha, k) \approx \left.\frac{dJ(\alpha, k)}{d\alpha}\right|_{\alpha=0} \alpha + \frac{1}{2} \left.\frac{d^2 J(\alpha, k)}{d\alpha^2}\right|_{\alpha=0} \alpha^2 \quad (10)$$

where $J(\alpha, k)$ is a shorthand for $L(t, \alpha, k)$. The selection procedure becomes $(\alpha^t, j) = \arg\min_{\alpha,k} J(\alpha, k) = \arg\min_{\alpha,k} \Delta J(\alpha, k)$. The optimisation over α has an analytical solution $\alpha_t = -J'/J''$.

Once the feature has been selected, the algorithm can proceed by applying an additional line search step to find the best coefficient as $\alpha^t = \arg\min_\alpha L(t, \alpha, j)$. One way to do is to repeatedly apply the update based on (10) until convergence.

Up to now, we have made an implicit assumption that all computation can be carried out efficiently. However, this is not the case for general Markov networks because most quantities of interest involve summation over an exponentially large number of network configurations. Similar to [2], we show that dynamic programming exists for tree-structured networks. However, for general structures, approximate inference must be used.

There are three quantities we need to compute: the distribution $p(v_i|x_i)$ in (6), the first and second derivative of J in (10). For the distribution, we have $p(v^i|x^i) = \sum_h p(v^i, h|x^i) = Z(v^i, i)/Z(i)$ where $Z(v^i, i) = \sum_h \exp(\sum_c F(x^i, v_c^i, h_c))$ and $Z(i) = \sum_y \exp(\sum_c F(x^i, y_c))$. Both of these partition functions are in the form of sum-product, thus, they can be computed efficiently using a single pass through the tree-like structure. The first and second derivatives of J are then

$$J'|_{\alpha=0} = \sum_i E_{h|v^i,x^i,t}[\sum_v \exp(\Delta F^{i,t}) \Delta f_k^i] \quad (11)$$

$$J''|_{\alpha=0} = \sum_i E_{h|v^i,x^i,t}[\sum_v \exp(\Delta F^{i,t})(\Delta f_k^i)^2] \quad (12)$$

Expanding (11) yields

$$J'|_{\alpha=0} = \sum_i \frac{1}{p(v^i|x^i, t)} \sum_{v,h} p(v, h|x^i, t) \Delta f_k^i \quad (13)$$

Note that f_k^i has the additive form as in (2) so $\Delta f_k(x^i, y) = \sum_c \Delta f_k(x^i, y_c)$. Thus (13) reduces to

$$J'|_{\alpha=0} = \sum_i \frac{1}{p(v^i|x^i,t)} \sum_c \sum_{y_c} p(y_c|x^i,t)\Delta f_k(x^i, y_c) \quad (14)$$

which now contains clique marginals and can be estimated efficiently for tree-like structure using a downward and upward sweep. For general structures, loopy belief propagation can provide approximate estimates. Details of the procedure are omitted here due to space constraints.

However, the computation of (12) does not enjoy the same efficiency because the square function is not decomposable. To make it decomposable, we employ Cauchy's inequality to yield the upper bound of the change in (10)

$$(\Delta f_k(x^i, y))^2 = (\sum_c \Delta f_k(x^i, y_c))^2 \leq |C| \sum_c \Delta f_k(x^i, y_c)^2$$

where $|C|$ is the number of cliques in the network.

The update using $\alpha = -J'/\tilde{J}''$, where $\tilde{J}''$ is the upper bound of the second derivative J'', is rather conservative, so it is clear that a further line search is needed. Moreover, it should be noted that the change in (10), due to the Newton update, is $\Delta\tilde{J}(\alpha, k) = -0.5(J')^2/\tilde{J}''$, where $\Delta\tilde{J}$ is the upper bound of the change ΔJ due to Cauchy's inequality, so the weak learner selection using the optimal change does not depend on the scale of the second derivative bound of $\tilde{J}''$. Thus the term $|C|$ in Cauchy's inequality above can be replaced by any convenient constant.

The complexity of our boosting algorithm is the same as that in the MLE of the Markov networks. This can be verified easily by taking the derivative of the log-loss in (3) and comparing it with the quantities required in our algorithm.

5. Experimental results

We restrict our attention to the linear-chain structure for efficient computation because it is sufficient to learn from the video data we capture. For all the experiments reported here, we train the model using the MLE along with the limited memory quasi-Newton method (L-BFGS) and we use the proposed boosting scheme with the help of a line search, satisfying Amijo's conditions. For regularisation, the same σ is used for all features for simplicity and is empirically selected. In the training data, only 50% of labels are randomly given for each data slice in the sequence. For the performance measure, we report the per-label error and the average $F1$-score over all distinct labels.

A. Data and feature extraction

In this paper, we evaluate our boosting framework on video sensor data. However, the framework is applicable to different type of sensors and is able to fuse different sensor information. The observational environment is a kitchen and dining room with two cameras mounted to two opposite ceiling corners (Figure 1). The observations are sequences of noisy coordinates of a person walking in the scene extracted from video using a background subtraction algorithm. The

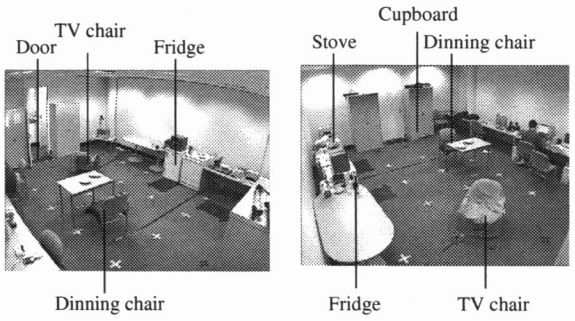

Fig. 1: The environment and scene viewed from the two cameras.

data was collected in our recent work [12], consisting of 90 video sequences spanning three scenarios: SHORT_MEAL, HAVE_SNACK and NORMAL_MEAL. We pick a slightly smaller number of sequences, of which we select 42 sequences for training and another 44 sequences for testing. Unlike [12], we only build flat models, thus separate models for each scenario are learned and tested. The labels are sequences of the 'activities' the person is performing, for example *'going-from-door-to-fridge'*. For training, only partial label sets are given, leaving a portion of the label missing. For testing, the labels are not provided, and the labels obtained from the decoding task are compared against the ground-truth.

It turns out that informative features are critical to the success of the models. At the same time, the features should be as simple and intuitive as possible to reduce manual labour. At each time slice τ, we extract a vector of 5 elements from the observation sequence $g(x, \tau) = (X, Y, u_X, u_Y, s = \sqrt{u_X^2 + u_Y^2})$, which correspond to the (X, Y) coordinates, the X & Y velocities, and the speed, respectively. These observation features are approximately normalised so that they are of comparable scale.

B. Effect of feature selection

Following previous boosting applications (e.g.[16]), we employ very simple decision rules: a weak hypothesis (feature function) returns a real value if certain conditions on the training data points are met, and 0 otherwise.

We design three feature sets. The first set, called the *activity-persistence*, captures the fact that activities are in general persistent. The set is divided into data-association features

$$f_{l,m}(x, y_\tau) = \delta[y_\tau = l] g_m(x, \tau) \quad (15)$$

where $m = 1, .., 5$, and label-label features

$$f_{l,m}(x, y_{\tau-1}, y_\tau) = \delta[y_{\tau-1} = y_\tau]\delta[y_\tau = l] \quad (16)$$

Thus the set has $K = 5|Y| + |Y|$ features, where $|Y|$ is the size of the label set.

The second feature set consists of *transition-features* that are intended to encode the activity transition nature as follows

$$f_{l1,l2,m}(x, y_{\tau-1}, y_\tau) = \delta[y_{\tau-1} = l1]\delta[y_\tau = l2] g_m(x, \tau) \quad (17)$$

Thus the size of the feature set is $K = 5|Y|^2$.

TABLE 1: PERFORMANCE ON THREE DATA SETS, ACTIVITY-PERSISTENCE FEATURES. SM = SHORT_MEAL, HS = HAVE_SNACK , NM = NORMAL_MEAL, AGTHM = ALGORITHM, ITRS = NUMBER OF ITERATIONS, FTRS = NUMBER OF SELECTED FEATURES, % FTRS = PORTION OF SELECTED FEATURES.

Data	SM	SM	HS	HS	NM	NM
Agthm	MLE	Boost	MLE	Boost	MLE	Boost
σ	∞	∞	∞	∞	∞	∞
error(%)	10.3	16.6	12.4	14.5	9.7	17.2
F1(%)	86.0	80.2	84.8	82.1	87.9	77.4
itrs	100	500	100	200	100	200
# ftrs	30	30	30	30	42	35
% ftrs	100	100	100	100	100	83.3

The third set, called the *context set*, is a generalisation of the second set. Observation-features now incorporate neighbouring observation points within a sliding window of width W

$$g_m(x, \tau, \epsilon) = g_m(x, \tau + \epsilon) \tag{18}$$

where $\epsilon = -W_l, ..0, ..W_u$ with $W_l + W_u + 1 = W$. This is intended to capture the correlation of the current activity with the past and the future, or the temporal *context* of the observations. The second feature set is a special case with $W = 1$. The number of features is a multiple of that of the second set, which is $K = 5W|Y|^2$.

The boosting studied in this section has the beam size $S = 1$, i.e. each round picks only one feature to update its weight. Tables 1, 2, 3 show the performance of the training algorithms on test data of all three scenarios (SHORT_MEAL, HAVE_SNACK and NORMAL_MEAL) for the three feature sets, respectively. Note that the infinite regularisation factor σ means that there is no regularisation. In general, sequential boosting appears to be slower than the MLE because it updates only one parameter at a time. For the activity persistence features (Table 1), the feature set is very compact but informative enough so that the MLE attains a reasonably high performance. Due to this compactness, the feature selection capacity is almost eliminated, leading to poorer results as compared with the MLE.

However, the situation changes radically for the activity transition feature set (Table 2) and for the context feature set (Table 3). When the observation context is small, i.e. $W = 1$, boosting consistently outperforms the MLE whilst maintaining only a partial subset of features ($< 50\%$ of the original feature set). The feature selection capacity is demonstrated more clearly with the context-based feature set ($W = 11$), where less than 9% of features are selected by boosting for the SHORT_MEAL scenario, and less than 3% for the NORMAL_MEAL scenario. The boosting performance is still reasonable despite the fact that a very compact feature set is used. There is therefore a clear computational advantage when the learned model is used for classification.

C. Learning the activity-transition model

We demonstrate in this section that the activity transition model can be learned by both the MLE and boosting. The

TABLE 2: PERFORMANCE ON ACTIVITY TRANSITION FEATURES

Data	SM	SM	HS	HS	NM	NM
Agthm	MLE	Boost	MLE	Boost	MLE	Boost
σ	∞	∞	∞	∞	∞	∞
error(%)	18.6	10.1	13.0	10.8	15.0	16.5
F1(%)	75.8	89.3	86.8	85.7	81.4	80.9
itrs	59	200	74	100	53	100
# ftrs	125	57	125	44	245	60
% ftrs	100	45.6	100	35.2	100	24.5

TABLE 3: PERFORMANCE ON CONTEXT FEATURES WITH WINDOW SIZE $W = 11$

Data	SM	SM	HS	HS	NM	NM
Agthm	MLE	Boost	MLE	Boost	MLE	Boost
σ	2	2	∞	∞	∞	∞
error(%)	15.3	9.6	9.4	11.2	9.3	16.6
F1(%)	81.6	87.7	89.3	86.6	87.7	78.1
itrs	51	200	22	100	21	100
# ftrs	1375	115	1375	84	2695	80
% ftrs	100	8.36	100	6.1	100	3.0

transition feature sets studied previously do not separate the transitions from data, so the transition model may not be correctly learned. We design another feature set, which is the bridge between the activity-persistence and the transition feature set. Similar to the activity persistence set, the new set is divided into data-association features, as in (15), and label-label features

$$f_{l1,l2}(y_{\tau-1}, y_\tau) = \delta[y_{\tau-1} = l1]\delta[y_\tau = l2] \tag{19}$$

Thus the set has $K = 5|Y| + |Y|^2$ features.

Given the SHORT_MEAL data set, and the activity transition matrix in Table 4, the parameters corresponding to the label-label features are given in Tables 5 and 6, as learned by boosting and MLE, respectively.

At first sight, it may be tempting to select non-zero parameters and their associated transition features, and hence the corresponding transition model. However, as transition features are non-negative (indicator functions), the model actually penalises the probabilities of any configurations that activate negative parameters exponentially, since $p(y|x) \propto$

TABLE 4: ACTIVITY TRANSITION MATRIX OF SHORT_MEAL DATA SET

Activity	1	2	3	4	11
1	1	1	0	0	0
2	0	1	1	0	1
3	0	0	1	1	0
4	0	0	0	1	0
11	0	0	0	0	1

TABLE 5: PARAMETER MATRIX OF SHORT_MEAL DATA SET LEARNED BY BOOSTING

Activity	1	2	3	4	11
1	1.8	0	-5904.9	-5904.9	0
2	-5904.9	3.6	0	-5904.9	0
3	-5904.9	-5904.9	2.425	0	-5904.9
4	-5904.9	-5904.9	-5904.9	2.4	-5904.9
11	-5904.9	-5904.9	-5904.9	-5904.9	2.175

TABLE 6: PARAMETER MATRIX OF SHORT_MEAL DATA SET LEARNED BY MLE

Activity	1	2	3	4	11
1	10.81	4.311	-5.7457	-5.3469	-1.8398
2	-2.2007	15.056	3.6388	-5.6644	0.41921
3	-5.3565	-2.3131	9.3656	1.6575	-2.3736
4	-5.4103	-4.556	-4.1142	7.1332	-5.2976
11	-3.17	-0.09001	-2.9518	-4.8741	8.9128

$\exp(\lambda_k f_k(y_{\tau-1}, y_\tau))$. Therefore, huge negative parameters practically correspond to improbable configurations. If we replace all non-negative parameters in Table 5 and 6 by 1, and the rest by 0, we actually obtain the transition matrix in Table 4. The difference between boosting and MLE is that boosting penalises the improbable transitions much more severely, thus leading to much sharper decisions with high confidence. Note that for this data set, boosting learns a much more correct model than the MLE, with an error rate of 3.8% ($F1 = 93.7\%$), in constrast to 15.6% ($F1 = 79.5\%$) by the MLE without regularisation, and 11.8% ($F1 = 85.0\%$) by the MLE with $\sigma = 5$.

D. Effect of beam size

Recall that the beam search described in Section 3-D allows the weak hypothesis to be an ensemble of S features. When $S = K$, all the parameters are updated in parallel, so it is essentially similar to the MLE, and thus no feature selection is performed. We run a few experiments with different beam sizes S, starting from 1, which is the main focus of this paper, to the full parameter set K. As S increases, the number of selected features also increases. However, we are quite inconclusive about the final performance. It seems that when S is large, the update is quite poor, leading to slow convergence. This is probably because the diagonal matrix resulting from the algorithm is not a good approximation to the true Hessian used in Newton's updates. It suggests that there exists a good, but rather moderate, beam size that performs best in terms of both the convergence rate and the final performance.

An alternative is just to minimise the exponential loss in (5) directly by using any generic optimisation method (e.g. see [2], [3]). However, this approach, although may be fast to converge, loses the main idea behind boosting, which is to re-weigh the data distribution on each round to focus more on hard-to-classify examples as in (8). These issues are left for future investigation.

6. CONCLUSIONS AND FURTHER WORK

We have presented a scheme to exploit the discriminative learning power of the boosting methodology and the semantically rich structured model of Markov networks and integrated them into a boosted Markov network framework which can handle missing variables. We have demonstrated the performance of the newly proposed algorithm over the standard maximum-likelihood framework on video-based activity recognition tasks. Our preliminary results on structure learning using boosting indicates promise. Moreover, the built-in capacity of feature selection by boosting suggests an interesting application area in small footprint devices with limited processing power and batteries. We plan to investigate how to select the optimal feature set online by hand-held devices given the processor, memory and battery status.

Although empirically shown to be successful in our experiments, the performance guarantee of the framework is yet to be proven, possibly following the large margin approach as in [14], [17], or the asymptotic consistency in the statistics literature as with the MLE. In the application to sensor networks, we intend to explore methods to incorporate richer sensory information into the weak learners, and to build more expressive structures to model multi-level and hierarchical activities.

ACKNOWLEDGMENTS

The Matlab code of the limited memory quasi-Newton method (L-BFGS) to optimise the log-likelihood of the CRFs is adapted from S. Ulbrich.

REFERENCES

[1] J. K. Aggarwal, Q. Cai, "Human motion analysis: A review", *Computer Vision and Image Understanding*, Vol. 73, No. 3, 1999, pp. 428-440.
[2] Y. Altun, T. Hofmann, M. Johnson, "Discriminative Learning for Label Sequences via Boosting", *Advances in Neural Information Processing Systems 15*, 2003, pp. 977-984.
[3] Y. Altun, M. Johnson, T. Hofmann, "Investigating Loss Functions and Optimization Methods for Discriminative Learning of Label Sequences", In: *Proceedings of the 2003 Conference on Empirical Methods in NLP*, 2003.
[4] H.H. Bui, S. Venkatesh, G. West, "Policy recognition in the abstract hidden Markov model", *Journal of Articial Intelligence Research*, Vol. 17, 2002, pp. 451-499.
[5] T.G. Dietterich, A, Ashenfelter, Y. Bulatov, "Training Conditional Random Fields via Gradient Tree Boosting", In: Proc. 21 st International Conference on Machine Learning, 2004, Banff, Canada.
[6] Y. Freund,R.E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting", *Journal of Computer and System Sciences*, Vol. 55, No. 1, 1997, pp. 119–139.
[7] J. H. Friedman, "Greedy function approximation: a gradient boosting machine", *Ann. Statist*, Vol. 29, No. 5, 2001, pp. 1189–1232.
[8] J. Friedman, T. Hastie, R. Tibshirani, "Additive logistic regression: a statistical view of boosting", *The Annals of Statistics*, Vol. 28, No. 2, 2000, pp. 337–374.
[9] A. Garg, V. Pavlovic, J. Rehg, "Boosted Learning in Dynamic Bayesian Networks for Multimodal Speaker Detection", In: *Proc. the IEEE*, 2003.
[10] J. Lafferty, A. McCallum, F. Pereira, "Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data", In: *Proc. 18th International Conf. on Machine Learning*, 2001, pp. 282–289.
[11] G. Lebanon, J. Lafferty, "Boosting and Maximum Likelihood for Exponential Models", *Advances in Neural Information Processing Systems 14*, 2002, pp. 447-454.
[12] N. Nguyen, D. Phung, S. Venkatesh, H. H. Bui "Learning and detecting activities from movement trajectories using the hierarchical hidden Markov models". In: *Proc. CVPR*, San Diego, CA, June 2005.
[13] S. D. Pietra, V. D. Pietra, J. Lafferty, "Inducing Features of Random Field", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 19, No. 4, 1997, pp. 380–393.
[14] R.E. Schapire, Y. Freund, P. Bartlett, W.S. Lee, "Boosting the Margin: A New Explanation for the Effectiveness of Voting Methods", *The Annals of Statistics*, Vol. 26, No. 5, 1998, pp. 1651-1686.
[15] R.E. Schapire, Y. Singer, "Improved Boosting Using Confidence-rated Predictions", *Machine Learning*, Vol. 37, No. 3, 1999, pp. 297–336.
[16] R.E. Schapire, Y. Singer, "BoosTexter: A Boosting-based System for Text Categorization", *Machine Learning*, Vol. 39, No. 2/3, 2000, pp. 135-168.
[17] B. Taskar, C. Guestrin, D. Koller, "Max-Margin Markov Networks", *Advances in Neural Information Processing Systems 16*, 2004.
[18] A. Torralba, K.P. Murphy, W.T. Freeman, "Contextual Models for Object Detection Using Boosted Random Fields", *Advances in Neural Information Processing Systems 17*, 2005, pp. 1401-1408.
[19] J. Yamato, J. Ohya, K. Ishii, "Recognizing human action in time-sequential images using hidden Markov models", In: *CVPR*, 1992, pp. 379-385.
[20] P. Yin, I. Essa, J.M. Rehg, "Asymmetrically Boosted HMM for Speech Reading", In: *Proc. CVPR*, Vol. 2, No. 2, 2004, pp. 755-761.
[21] Y. Jing, V. Pavlovic, J.M. Rehg, "Efficient discriminative learning of Bayesian network classifiers via boosted augmented naive Bayes", In: *Proc. ICML*, 2005.