DETECTING UNSAFE DRIVING PATTERNS USING DISCRIMINATIVE LEARNING

Yue Zhou†, Wei Xu‡, Huazhong Ning†, Yihong Gong‡, Thomas S. Huang†

†Beckman Institute, University of Illinois at Urbana-Champaign, Urbana, IL 61801
‡NEC Laboratories America, Inc., Cupertino, CA 95014

ABSTRACT

We propose a discriminative learning approach for fusing multi-channel sequential data with application to detect unsafe driving patterns from multi-channel driving recording data. The fusion is performed using a discriminatively trained graphical model - conditional random field (CRF). The proposed approach offers several advantage over existing information fusing approaches. First, it derives its classification power by directly modelling and maximizing the conditional probability. Second, it represents the variable dependency in an undirected graph, which is very efficient in inference. Third, it does not require to label all the training data and utilizes both labelled and unlabelled data efficiently by semi-supervised learning algorithms. The proposed approach is evaluated on driving recording data collected from driving simulator - STISIM. Experiments show it outperforms the simple discriminative classifier (SVM) and generative model (HMM).

1. INTRODUCTION

Unsafe driving patterns are often implicated in traffic accidents. Although it is difficult to determine the exact number of accidents due to unsafe driving patterns, it is widely believed that the two are highly correlated. In efforts to shed light on road safety problems, considerable research has been devoted to developing mathematical model and algorithms to evaluate and predict the operators' driving statuses and performances. [7] [3] use video image and physiological data to tracking eyes' and facial movements for driver fatigue monitoring and prediction. A computational state-based model of driver behavior is developed in [5] using Hidden Markov Model (HMM) and after training it is able to predict when the driver is about to brake or turn. Research is also performed to monitor the behavior of the driver from the recording of the hardware system under the control of the driver such as steering wheel, acceleration, braking, gear changing [1].

Most technologies detect unsafe driving behaviors from real-time recording of multi-sensory signals such as driving performance, physiological activities, videos, and etc. Although imperfectly and inaccurately, each type of sensory signal provides some characterizations of the system status. However, we believe that if different types of sensory data are combined systematically under a unified framework, it will
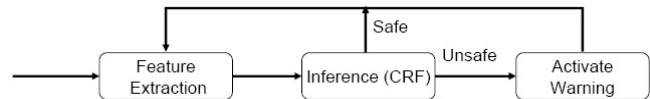

Fig. 1. Flowchart of the unsafe driving pattern detector.

reduce the uncertainty and ambiguity present in the single-source information inference algorithm.

In this paper, we focus on the problem of inferring hidden driving patterns from multiple sequential signals in driving recordings. The proposed algorithm is tested in the data collected from driving simulators. Fig 1 gives a brief overview of the proposed algorithm applied to driving pattern classification task. The driver's performance data is sampled and recorded in real-time. Features are extracted from the raw data. The feature vector is then fed into the inference model to estimate the current driving state - either safe or unsafe. The warning system will be activated if unsafe driving pattern is detected. The parameters of the inference model are learnt from training data and labels. In training, we effectively combine the labelled and unlabelled data using semi-supervised learning algorithms.

The rest of the paper is organized as follows: Section 2 discuss the CRF model and its learning/inference. We show how to use semi-supervised learning method to combine labelled and unlabelled data for training in Section 3. Section 4 describe the data collection and feature extraction. Encouraging experimental results are demonstrated in Section 5.

2. CONDITIONAL RANDOM FIELD (CRF)

The problem of fusing multi-channel sequential signal can be viewed as inferring the hidden patterns from observation vectors - a classification problem. However, it is not trivial to directly apply discriminative classifiers such as support vector machine (SVM) because the observations are temporally correlated. One alternative approach is generative model such as Hidden Markov Model (HMM). The drawback of HMM is that it models the observations independently given the hidden variables, and this assumption is too strong in many cases. For example, in HMM we can not use the feature generated
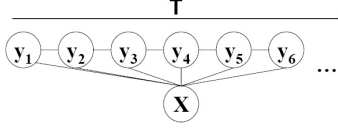
Fig. 2. A diagram for linear-chain CRF.

from the whole observation such as the total number of accidents/violations. Because of the problems addressed above, we choose Conditional Random Field (CRF) - a discriminatively trained graphical model as our inference model.

2.1. Definition of CRF

Conditional random field (CRF) [4] is a discriminatively-trained undirected graphical model of which hidden modes have Markov property conditioned on the observation. CRF brings together the merits of discriminative and generative models. The features are combined using exponential functions, which are trained discriminatively to maximize the conditional probability of the label given observation. But like generative models, the probabilities are propagated through neighboring nodes to obtain globally optimal labelling, which allow us to infer efficiently. Due to the sequential nature of driving record data, in this paper we focus on linear-chain CRF.

The feature function of node y_i in the CRF is defined as:

$$F_i(Y, X) = \sum_k \lambda_k \cdot f_k(y_i, y_{i-1}, X)$$

where f_k is the feature. The conditional probability of label sequence y given input sequence x is given by:

$$p(Y|X) = \frac{\exp(\sum_i F_i(Y, X))}{Z(X)} \qquad (1)$$

where i ranges over input nodes and $Z(X)$ is the partition function:

$$Z(X) = \sum_Y \exp(\sum_i F_i(Y, X))$$

In this paper we focus on linear-chain CRF although CRF was originally defined in general graphs. In linear-chain CRF the hidden variables are sequential connected in a linear-chain structure, where the marginal probabilities can be computed in linear time by dynamic programming. Fig 2 gives an example of linear-chain CRF.

2.2. CRF inferences

The inference problem of CRF is to find the best label Y such that the log likelihood is maximized: $\arg\max_Y P(Y|X, \theta)$, where θ is the parameters.

The marginal probabilities of CRF can be inferred using dynamic programming in linear time analog to the Viterbi algorithm for HMM.

2.3. Training CRF parameters

The training problem of CRF is defined as follows: given the training data set y_i, find the best model parameters such that the log likelihood $\log p(Y|X)$ is maximized. There is no close-form solution to this problem. However, because the log likelihood function is convex, it can be maximized numerically by many optimization methods, e.g. quasi-Newton method.

The partial derivative of the log likelihood function is computed as follows:

$$\frac{\partial \log p(Y|X)}{\partial \lambda_k} = \sum_{i=1}^N f_k(Y, X) - \sum_{i=1}^N \sum_{Y'} f_k(Y', X) \cdot P(Y'|X) - \frac{\lambda_k}{\sigma^2}$$

$$(2)$$

where N is the number of training data. The equation shows that the kth partial derivative of the log likelihood function is the difference between the empirical expected value of feature $f_k(Y, X)$ and the model's expected value of $f_k(Y', X)$ minus the derivative of the Gaussian prior. The Gaussian prior over the parameter space provides smoothing to help cope with sparsity in the training data [2].

3. TRAINING WITH LABELLED AND UNLABELLED DATA

3.1. The labelled and unlabelled data

Training the CRFs requires the labels of the whole data sequence. In our case, unfortunately, it is very time-consuming and expensive to label all the data sequences. The reason is that although we can always assign unsafe labels to the data with accidents or violations, the reverse argument is not true - driving patterns without accidents or violations may not be safe because unsafe driving patterns do not necessarily lead to accidents or violations. Therefore, we have to manually look at the long driving performance recording and assign label to each segment. It requires the efforts of human annotators who have fairly amount of driving experiences. Another problem with this labelling scheme is that the data labelling is subjective especially for a few non-trivial cases and the labels from different annotators may not agree with each other.

In our system we use the semi-supervised learning techniques, which effectively combine the labelled and unlabelled data.

3.2. Semi-supervised learning with Gaussian fields

To utilize the unlabelled data, we use the semi-supervised learning algorithm proposed in [6]. This approach is based on

Gaussian fields defined on a weighted graph over the labelled and unlabelled data. The weights of the graph are predefined in terms of the similarity function between observations. We give a brief overview of the algorithm in the following:

1. Construct a undirect graph whose nodes correspond to each data point - labelled or unlabelled.

2. Compute the weight of the graph edge from a similarity function, e.g. a Gaussian function.

3. Assign a real-valued function f to each node with the constraint that f equals the label for the labelled data.

4. Minimize the quadratic graph energy function

$$E(f) = \frac{1}{2} \sum_{i,j} w_{i,j}(f(i) - f(j))^2 \qquad (3)$$

5. Assign labels to the unlabelled data according to its f function value.

4. DATA PREPROCESSING AND FEATURE COMPUTING

4.1. Data collection and preprocessing

The driver performance data was collected by the STISIM car driving simulators (http://www.systemstech.com/). The simulator was instrumented to save data on the driver's performance during the driving course. The data is sampled at the frequency of 30Hz. Fig 3(a) shows a picture of the interface of STISIM driving simulator.

In our system, we use 9 channels of driving performance recording: throttle, brake, steering wheel, position, speed, acceleration, lane position, distance to same lane vehicle, distance to incoming vehicle. Because our CRF model is general, we may include more types of data in the future without modifying the whole framework. Fig 3(b) shows a sample plot of the driving performance data.

Because the sampling rate of driving performance data is very high, the raw data is first evenly segmented by sliding-windows of one-second size. The statistics of the data within each sliding window are calculated, e.g. minimum, maximum, mean, variance, first-order derivative and etc. A Gaussian Mixture Model (GMM) is then estimated using *Expectation Maximization (EM)* algorithm of each statistic for safe driving pattern and unsafe driving pattern respectively.

4.2. Features for CRF

The CRF relaxes the observation independence assumption in HMM and estimates the conditional probability given the whole observations. Therefore, we have much more flexibilities in choosing features for the model. We generally categorize the features into two types:
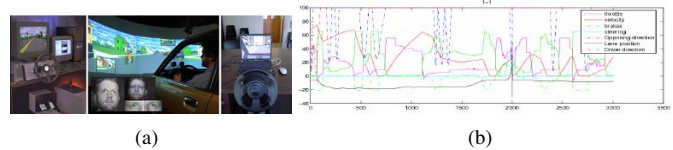

(a)   (b)

Fig. 3. STISIM driving simulator (a)The simulator interface; (b)A plot of the collected driving recording data.

1. **HMM-style features** The HMM-style features are pretty much the same as the features we use for HMM. The features are in the forms of (a) $f(y_i, x_i)$ which associates one hidden variable with its corresponding observation (b) $f(y_{i-1}, y_i)$ which represents the graph dynamics.

2. **CRF-style features** The CRF-style features are more generalized and associate one hidden segments or two neighboring hidden segments with the whole observation. The features are in the form of (a) $f(y_i, X)$ (b) $f(y_i, y_{i-1}, X)$.

The features in CRFs are much more expressive than in HMMs because they depend on the whole observations. In modelling the driving recording data, however, we restrict that all the features are causal, i.e., they only depend on the observation on or before current time.

4.3. Feature extraction from driving recording data

The statistical data of each sliding window x_i is compared with the GMM model to compute the feature $f(y_i, x_i)$. The system dynamic $f(y_{i-1}, y_i)$ is modelled using a Gaussian function.

The CRF-style features capture the dependency between the hidden state and the observation. For example, we aggregate the number of past accidents and violations as a bias for unsafe driving patterns, with features of form $f(y_i, X)$. Other features of $f(y_i, X)$ include the dependency between velocity variance and y_i, the dependency between average velocity and y_i and etc. The relationship between observation variance and system dynamic is modelled using features of form $f(y_i, y_{i-1}, X)$. We also penalize very short safe driving segment because it is very unlikely in real situation.

5. EXPERIMENTS

5.1. Experiment setting

Our experiments are carried out in the following way. The driving performances used in experiments are collected and recorded from the same driving courses taken by 15 different individuals. The performance is evaluated by cross-validation

	SVM	HMM	CRF
$P(A|U)$	0.059	0.063	0.081
$P(A|S)$	0.013	0.0103	0.0072
$P(U|A)$	0.539	0.580	0.700
$\frac{P(A|U)}{P(A|S)}$	4.5	6.1	11.3

Table 1. Experimental result of different classifiers

- each time we use 14 individuals for training and 1 individual for testing.

5.2. Performance evaluation

For comparison, we test other classifiers such as SVM, HMM with the same data. The SVM is directly applied to the sliding-window statistics with RBF kernel. For HMM, the likelihood is modelled using GMM. The CRF is trained with both labelled and unlabelled data.

One problem with evaluation is how to justify unsafe driving pattern. As we mention in the previous section, we believe accident and unsafe driving patterns are correlated and the probability of accident under unsafe driving patterns should be much high than under safe driving patterns. Let us denote A as accident, U as unsafe driving pattern, S as safe driving pattern. The conditional probabilities of accident under safe/unsafe driving patterns - $P(A|S)$ and $P(A|U)$ can be used to evaluate the accuracy of the classifiers indirectly. We report the the likelihood $P(A|U)$, $P(A|S)$, the posterior $P(U|A)$ and the likelihood ratio $\frac{P(A|U)}{P(A|S)}$ for SVM, HMM, and CRF respectively in Table 1. In all experiments 30% of the training data is labelled. The proposed algorithm increases the likelihood ratio $\frac{P(A|U)}{P(A|S)}$ from 4.5 for SVM up to 11.3. High $\frac{P(A|U)}{P(A|S)}$ indicates that the probability of accident is much higher in unsafe driving pattern than in safe driving pattern, which corresponds to our assumption that the unsafe driving pattern implicates accident.

We also compare the performances of the proposed algorithm given different sizes of labelled data in Fig 4. It shows that $P(A|U)$ and $P(A|S)$ are improved with more labelled data, which implicates better classification performances.

6. CONCLUSION

We have presented a discriminative classification framework for fusing multi-channel sequential data to label hidden patterns. It only requires partial data to be labelled. The learning and inference model is based on conditional random field (CRF). The proposed algorithm also uses Gaussian field and harmonic function to combine labelled and unlabelled data for training. Experiments demonstrate significant improvement in performance compared to individual classifier (SVM) and generative model (HMM).
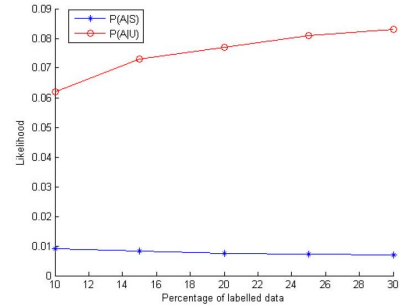

Fig. 4. The performance change with the increase of labelled data.

Future work includes exploring more expressive features for the model, and incorporating continuous hidden labels. We are exploring semi-Markov CRF to further improve the performance of the inference. It may also be possible to apply other semi-supervised learning algorithm to the proposed approach.

7. REFERENCES

[1] P. Artaud, S. Planque, C. Lavergne, H. Cara, P. de Lepine, C. Tarriere, and B. Gueguen. An on-board system for detecting lapses of alertness in car driving. In *14th Int. Conf. Enhanced Safety of Vehicles*, volume 1, Munich, Germany, 1994.

[2] S. F. Chen and R. Rosenfeld. A survey of smoothing techniques for me models. In *IEEE Trans, on Speech and Audio Processing*, volume 8, pages 37–50, 2000.

[3] A. Haro, M. Flickner, and I. Essa. Detecting and tracking eyes by using their physiological properties, dynamics, and appearance. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 01:1163, 2000.

[4] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proc. 18th International Conf. on Machine Learning*, pages 282–289, 2001.

[5] A. Pentland and A. Liu. Modeling and prediction of human behavior. In *Neural Computation*, volume 11, pages 229–242, 1999.

[6] Xiaojin Zhu and Zoubin Ghahramani and John Lafferty. Semi-Supervised Learning Using Gaussian Fields and Harmonic Functions. In *Proceedings of International Conference on Machine Learning*, 2003.

[7] Z. Zhu and Q. Ji. Real time and non-intrusive driver fatigue monitoring. In *7th IEEE International Conference on Intelligent Transportation Systems*, Washington DC, 2004.