Driver Behavior Recognition and Prediction in a SmartCar

Nuria Oliver and Alex P. Pentland

Media Laboratory
Massachusetts Institute of Technology (MIT)
20 Ames St E15-385
Cambridge MA 02139

ABSTRACT

This paper presents our SmartCar testbed platform: a real-time data acquisition and playback system and a machine learning –dynamical graphical models– framework for modeling and recognizing driver maneuvers at a tactical level, with particular focus on how contextual information affects the driver's performance. The SmartCar's perceptual input is multi-modal: four video signals capture the surrounding traffic, the driver's head position and the driver's viewpoint; and a real-time data acquisition system records the car's brake, gear, steering wheel angle, speed and acceleration throttle signals. We have carried out driving experiments with the instrumented car over a period of 2 months. Over 70 drivers have driven the SmartCar for 1.25 hours in the greater Boston area. Dynamical Graphical models, HMMs and potentially extensions (CHMMs), have been trained using the experimental driving data to create models of seven different driver maneuvers: passing, changing lanes right and left, turning right and left, starting and stopping. These models are essential to build more realistic automated cars in car simulators, to improve the human-machine interface in driver assistance systems, to prevent potential dangerous situations and to create more realistic automated cars in car simulators.

Keywords Driver behavior modeling, driver assistance, system architectures, situation awareness

1. INTRODUCTION

This paper presents our SmartCar testbed platform: a real-time data acquisition and playback system and a machine learning –graphical models– framework for modeling and recognizing driver maneuvers at a tactical level, with particular focus on how contextual information affects the driver's performance. The SmartCar's perceptual input is multi-modal: four video signals capture the surrounding traffic, the driver's head position and the driver's viewpoint; and a real-time data acquisition system records the car's brake, gear, steering wheel angle, speed and acceleration throttle signals. We have carried out driving experiments with the car over a period of 2 months. Over 70 drivers have driven the SmartCar for 1.25 hours in the greater Boston area. Dynamical Graphical models, HMMs and potentially extensions (CHMMs), have been trained using the experimental driving data to create models of seven different driver maneuvers: passing, changing lanes right and left, turning right and left, starting and stopping. These models are essential to build more realistic automated cars in car simulators, to improve the human-machine interface in driver assistance systems, to prevent potential dangerous situations and to create more realistic automated cars in car simulators. Our behavior models let us correctly classify the maneuver on average 1 **second before** any significant change in the car signals takes place.

The paper is structured as follows: first, the most relevant previous work is described in section 2; section 3 presents an overview of the system; the perceptual input of the SmartCar testbed is described in section 4; the statistical models used for behavior modeling and recognition are described in section 5. Section 6 contains the description of our experiments and reports the recognition results in real driving situations. Finally, section 7 summarizes the main conclusions and outlines our future lines of research.

Further author information: (Send correspondence to Nuria Oliver)
Nuria Oliver: E-mail: nuria@media.mit.edu
Alex Pentland.: E-mail: sandy@media.mit.edu

2. BACKGROUND AND PREVIOUS WORK

Human driver modeling is an interdisciplinary endeavor involving a number of fields including robotics, psychology, control theory and statistics. Driving in a real-life traffic situation is a very difficult task because good decisions need to be made given only incomplete information in real time. Traditional AI techniques such as search-based planning are infeasible for at least two reasons: most of these methods cannot function under noisy, uncertain conditions, and the state-space is extremely large if realistic maneuvers such as aborted lane changes are taken into account.

One critical issue in machine-human interface systems are the transitions between manual and automated operation. They should be as seemless and smooth as possible. Such transitions would occur, for example, when the system encounters non-supported situations, when it fails, returning to manual mode; or when initiated by the driver. In any case, it is very important not to interfere with the driver's intended maneuver, specially in emergency situations, and to avoid discontinuities in the system, inducing feelings of incongruity while driving. Therefore, developing systems for predicting the driver's next maneuver or inferring driver's intentions is imperative to facilitate smooth and appropriate control mode transitions.

Building effective driver behavior recognition methods requires a thorough understanding of driver behavior and the construction of a model capable of both generating and explaining the drivers' behavioral characteristics. The task of driving has traditionally been characterized as consisting in three different levels: strategic, tactical and operational.[6] At the highest (strategic) level, a route is planned and goals are determined; at the intermediate (tactical) level, maneuvers are selected to achieve short-term objectives –such as deciding whether to pass a blocking vehicle–; and at the lowest (operational) level, those maneuvers are translated into control operations. In this paper we focus on recognizing driving maneuvers at a tactical level. Namely, we have built models of passing, changing lanes right and left, turning right and left, starting, and stopping.

Previous studies in psychology have found that driver behavior can be characterized as a sequence of basic actions each associated with a particular state of the driver-vehicle-environment system and characterized by a set of observable features.[4]

The closest work to ours is that of Pentland and Liu,[8,13] and that of Kuge at al.[7] In[13] Pentland and Liu develop a computational state-based model of driver behavior. They model the driver's internal state as a four-state Hidden Markov Model (HMM). Once the HMM has been trained the system is able to predict when the driver is about to brake or turn. This knowledge may then be used by a smart vehicle to optimize its behavior for the expected maneuver –in some sense, the situation awareness is shared over the vehicle-driver system. In a similar way, Kuge et al. present an HMM method that characterizes and detects lane changing maneuvers. The authors focus on information processing models of human driver behavior generation and utilize them to adopt a model based approach in the development of a lane change detection and recognition model. The primary components are skilled low level maneuvers whose initiation is managed by higher level decision making components. Perceptual models can be used to gain some insights into this area. Recent research[12] shows that drivers' eye fixation patterns are strongly correlated with their current mental state. Other more constrained but certainly important aspects of driver behavior were estimated by few early methods, such as, for example, lane change intention.[9] However, none of these methods was human model-based.

Pentland and Liu validated their model in an experiment conducted in a driving simulator. The objective of that validation test was to recognize different driving maneuvers at a tactical level, such as a right turn, a left turn or stopping. In order to apply such a model to a driver assistance system, it is necessary to assess to what degree the HMM based behavior recognition model also provides a plausible model for human behavior generation. This knowledge may not only offer better insight into selecting a particular HMM structure but also provide better insight into potential limitations of the characterization in situations that were not part of the training set used to fit the HMM parameters.

None of these previous systems, however, incorporates contextual information when modeling driver behavior. Nonetheless, knowledge of the context is necessary to properly make decisions in complex dynamic environments such as driving. Psychologists attribute this competence to a task-specific understanding of the situation, termed *situation awareness*. I this paper we develop a machine models of driver behavior that incorporate elements of situational awareness for tactical driving.

There is today strong research efforts invested in developing partially or fully automated driver assistance systems. For example, headway distance control or lane keeping control systems, which make use of Intelligent Transportation

System (ITS) technologies.[2,5] To achieve such assistive systems, it is important to adopt approaches aimed at improving the performance of the driver-vehicle-context cooperative system by regarding driving as an interaction between the driver, the vehicle and the surrounding road information and traffic.

Finally, it has also been argued that laboratory research of SA should be conducted under conditions that afford as much realistic behavior as possible. Due to the simplicity of most car simulators, specially the lack of realism of the computer generated automated cars, the experiments carried out in this paper took place in a real car while driving in the greater Boston area.

To summarize, this paper extends Pentland and Liu's framework[8,13] in several ways: (1) we model a larger number of maneuvers at a tactical lever –namely seven–; (2) we show that contextual information is critical for the accurate recognition of some maneuvers; (3) we use real data collected in an instrumented car, as opposed to using a car simulator.

3. SYSTEM'S ARCHITECTURE

The system's architecture is depicted in figure 1. In the proposed architecture, there is a *bottom-up* stream of

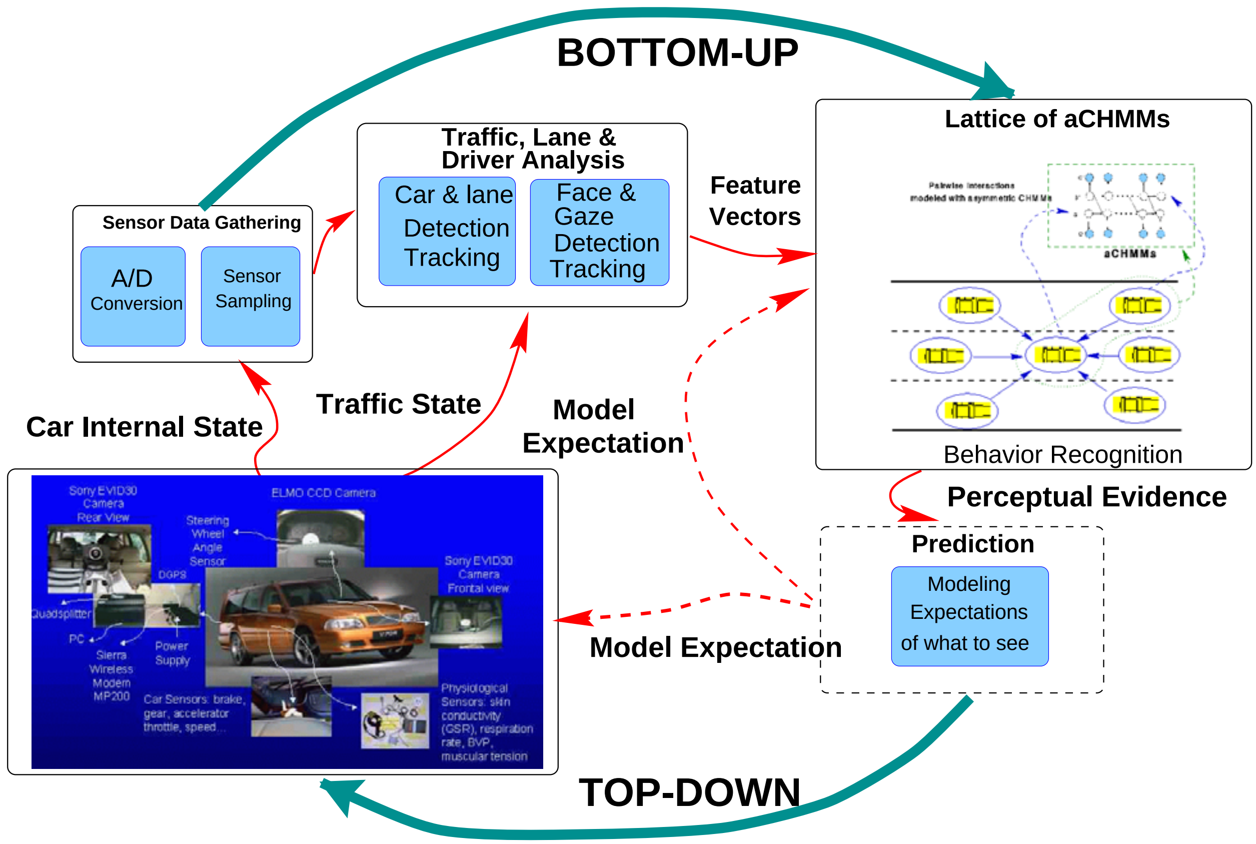

Figure 1. SmartCar architecture

information gathered with the various sensors, and a *top-down* information flow through the predictions provided by the behavior models. Consequently a Bayesian approach offers a mathematical framework for both combining the observations (bottom-up) with complex behavioral priors (top-down) to provide expectations that could eventually be fed back to the perceptual system.

4. PERCEPTUAL INPUT

There are at least three different aspects beyond the driver that are relevant when driving at a tactical-level driving:

1. SmartCar physical self-state: information sensed from the speedometer, acceleration throttle, steering wheel angle sensor (optical encoder), brake pedal, gear and GPS unit.
2. Road state: including road geometry and exit information.
3. Traffic state: speeds and relative positions of the surrounding traffic.

The sensors installed in the SmartCar provide information about: (1) the car's internal state : brake, acceleration throttle, steering wheel angle, gear and speed; (2) surrounding traffic and lanes: via two Sony EVI-D30 cameras with wide field of view lenses, mounted on the front dashboard and on a tripod in the trunk, to record frontal and rear traffic; and (3) driver's face position, orientation and viewpoint: by use of an ELMO CCD camera mounted on the steering wheel, recording the driver's face, and another ELMO CCD camera mounted on a pair of glasses to record the driver's viewpoint. All the video signals are combined in a quadsplitter whose output is recorded using a Sony GV-A500 Hi8 Video Walkman recorder.

We have developed the hardware and software for acquiring in real-time car state data. All the signals but the steering wheel are available directly from the car electronics system. We have designed a steering wheel angle sensor and mounted it on the car's steering wheel. The hardware obtains its inputs from sources of three different nature as shown in table 1. All the car signals are connected to a PCMCIA Data Acquisition Card by National Instruments. The analog signals are digitized and sampled at 150 scans/sec. The digital signals are sampled using the same card at the same sampling rate (150 scans/sec). All the signals can be directly connected to one of these boards, except for the speed, given that it consists of a 12 pulse-per-revolution signal. Therefore for this signal a frequency-to-voltage converter is used to convert it to analog.

Table 1. Sensor signals in the Smart Car.

Signal	**Nature**	**Description**
Speed	Analog	12 pulse/wheel rev. square wave
Acc	Analog	Linear 0-12 V
Brake	Digital	Boolean (0=off, 1=on)
Gear	Digital	2-bit
Steering angle	Analog	Up to 3 revolutions

The software for data acquisition and playback has been developed in LabVIEW. LabVIEW is a powerful programming environment used in engineering and scientific environments. LabVIEW is based on a functional programming language known as *G*, developed by National Instruments. We have developed a graphical user interface for calibrating the car signals, triggering the acquisition, and annotating the driving maneuvers as they take place.

5. DRIVER BEHAVIOR MODELS

Statistical directed acyclic graphs (DAGs) or probabilistic inference networks (PINs)[3,15] can provide a computationally efficient solution to the problem of time series analysis and modeling. HMMs and some of their extensions, in particular CHMMs,[1,11] can be viewed as a particular and simple case of temporal PIN or DAG. Graphically Markov Models are often depicted 'rolled-out in time' as Probabilistic Inference Networks, such as in figure 2.

PINs present important advantages that are relevant to our problem: they can handle incomplete data as well as uncertainty; they are trainable and easier to avoid overfitting; they encode causality in a natural way; there are algorithms for both doing prediction and probabilistic inference; they offer a framework for combining prior knowledge and data; and finally they are modular and parallelizable.

Traditional HMMs offer a probabilistic framework for modeling processes that have structure in time. They offer clear Bayesian semantics, efficient algorithms for state and parameter estimation, and they automatically perform dynamic time warping. An HMM is essentially a quantization of a system's configuration space into a small number of discrete states, together with probabilities for transitions between states. A single finite discrete variable indexes the current state of the system. Any information about the history of the process needed for future inferences must be reflected in the current value of this state variable.

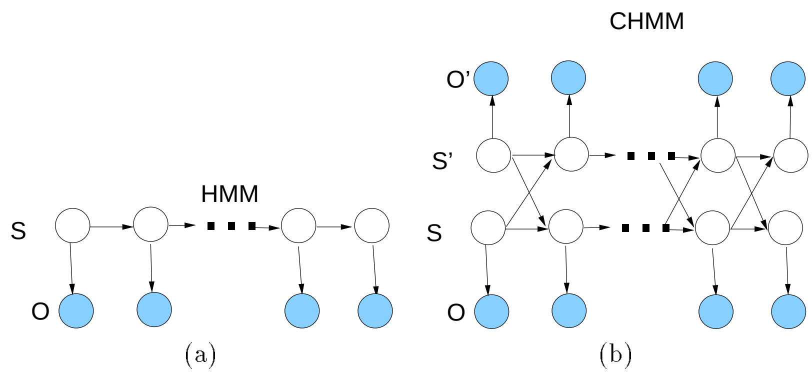

Figure 2. Graphical model representation of a HMM (a) and a CHMM (b) rolled-out in time

However many interesting real-life problems are composed of multiple interacting processes, and thus merit a compositional representation of two or more variables. This is typically the case for systems that have structure both in time and space. With a single state variable, Markov models are ill-suited to these problems. In order to model these interactions a more complex architecture is needed. We have developed a new architecture called Coupled Hidden Markov Models (CHMMs) for modeling interacting processes.[1,11]

The posterior state sequence probability $P(S|O)$ for HMMs is given by

$$P(S|O) \quad = \quad P_{s_1} p_{s_1}(o_1) \prod_{t=2}^{T} p_{s_t}(o_t) P_{s_t|s_{t-1}} \tag{1}$$

where $S = \{a_1, \ldots, a_N\}$ is the set of discrete states, $s_t \in S$ corresponds to the state at time t. $P_{i|j} \doteq P_{s_t=a_i|s_{t-1}=a_j}$ is the state-to-state transition probability (i.e. probability of being in state a_i at time t given that the system was in state a_j at time $t-1$). In the following we will write them as $P_{s_t|s_{t-1}}$. $P_i \doteq P_{s_1=a_i} = P_{s_1}$ are the prior probabilities for the initial state. Finally $p_i(o_t) \doteq p_{s_t=a_i}(o_t) = p_{s_t}(o_t)$ are the output probabilities for each state*. The well-known Baum-Welch algorithm efficiently estimates –using dynamic programming– the state posterior probability in an HMM (inference problem). The MAP identification problem in the context of HMMs involves identifying the most likely hidden state sequence given the observed evidence. Just as with the inference problem, the Viterbi algorithm provides an efficient, locally recursive method for solving this problem with complexity TN^2.

6. EXPERIMENTS

Apparatus A self-instrumented automatic Volvo V70XC (1998) was used to measure driver behavior data. The car sensors have been described in section 4.

The procedure The driving task took place in the greater Boston area. Over 70 drivers drove both in the city and in different highway sections, for about 1.25 hours. The drivers were asked to sign an consent form before starting the experiment. They were rewarded $20 for participating.

A driving instructor was with the driver throughout the experiment. The instructor gave directions to the driver about where to go and labeled the driving maneuvers as they took place using the laptop computer and the LabVIEW GUI. Because our focus in on predicting what is the most likely maneuver to take place next, the driver was requested to verbally report his/her **next intended action before** carrying it out. The four video signals were recorded for the entire route. The car signals, however, were only recorded when a maneuver was about to happen. A time window of 2 seconds was used, i.e. the car signals were recorded starting 2 seconds **before** the driver reported his intentionality to perform a maneuver. Both the video and car data was time stamped (the VCR and the laptop clocks were synchronized before every session). The maneuvers that we collected data for are: passing another car, turning right and left, changing lanes right and left, starting, stopping and merging.

*The output probability is the probability of observing o_t given state a_i at time t

Figures 3 and 4 show typical car and context signals in one example of a 'passing' and a 'turning' maneuvers collected in the experiments. Note how, in the case of passing, the car signals contain little information about the maneuver type, whereas the gaze and lane are much more relevant features.

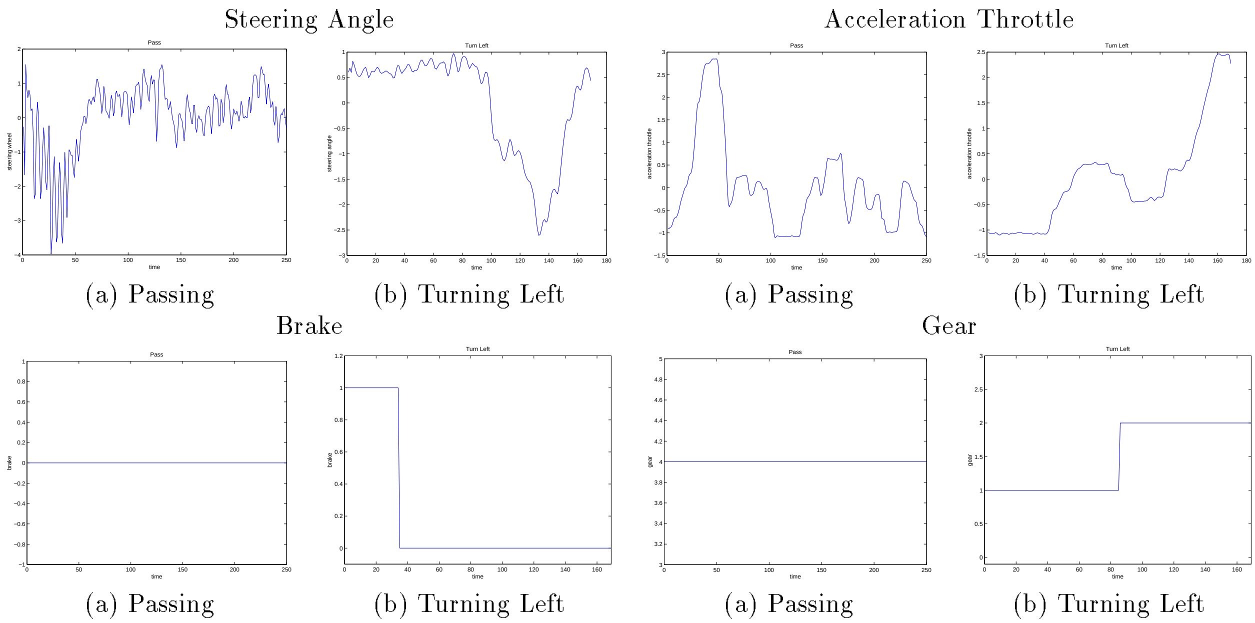

Figure 3. Typical car signals for passing and turning left maneuvers

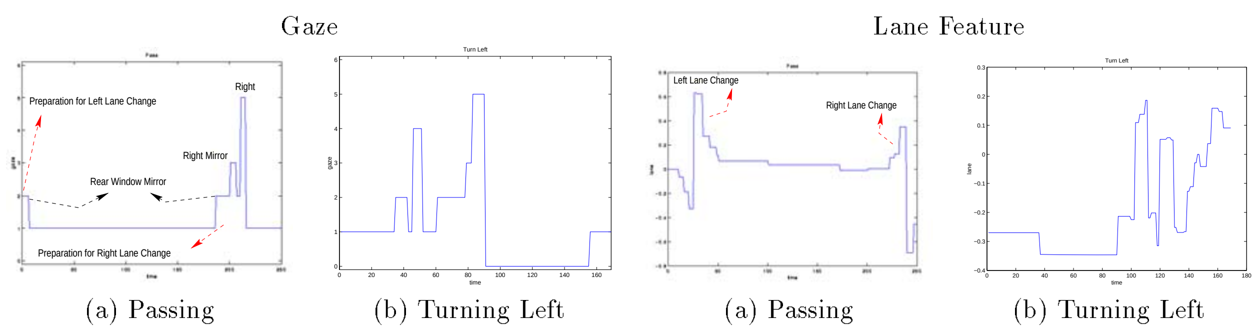

Figure 4. Typical contextual (gaze and lane) signals for a passing and turning left maneuvers

After the driving task was completed, the drivers were asked to fill in a questionnaire with basic questions about their driving experience, skills and the experiment.

Data post-processing and Driver Models The contextual information was acquired via the video signals. We have developed a video processing graphical environment that let's the user record, playback and annotate the video signals coming from the front, rear and face driver cameras. Contextual information –such as the driver's gaze, the relative position of the road lanes or the relative position, velocities and direction of the surrounding traffic– was manually annotated for each frame and maneuver.

Due to the different sampling rate on the car and video signals, we subsampled the car data to match the video frame rate. The final sampling rate was of approximately 30 samples/s. All the continuous signals were low-pass filtered using Butterworth filters.

Using the car, driver's gaze and road lane data, we built HMMs for each of the maneuvers to be recognized. We evaluated the performance on recognition (accuracy) of the best HMMs trained with different feature vectors:

1. Only car signal data: brake, steering wheel angle, gear, and acceleration throttle.
2. Car data and lane position information (front and back lane positions).
3. Car data and driver gaze information.
4. Car data, lane and driver information.

The gaze was a discrete signal with 6 possible values: (1) front road, (2) rear window mirror, (3) right mirror, (4) left mirror, (5) right and (6) left.

In the case of the lanes, a single value was computed from the (x, y) image coordinates of the road lanes:

$$lane_i = \mathrm{atan2}(|\mathrm{y}_2 - \mathrm{y}_1|, |\mathrm{x}_2 - \mathrm{x}_1|) \tag{2}$$

$$i \in \{\text{front left (fl)},\ \text{front right (fr)},\ \text{back left (bl)},\ \text{back right (br)}\} \tag{3}$$

$$lane_{feat} = \frac{lane_{fr} + lane_{br} - (lane_{fl} + lane_{bl})}{4.0} \tag{4}$$

The best models (best number of states and feature vector) were selected using 10-fold cross-validation. The training data set was about 80% of the total amount of data. The testing data set consisted of the rest of the data that had not been used for training.

The number of examples collected in the driving experiments is summarized in table 2. The table contains also the average length of each maneuver in number of samples and in seconds.

Table 2. Number of driving examples and average length per maneuver in number of samples

	Number of driving examples		Average Length #samples (s)	
	Car data	Traffic data	Car data	Traffic data
Passing	710	40	517 (17.2 s)	341 (11.6 s)
turning right	257	37	258 (8.6 s)	159 (5.3 s)
turning left	260	31	258 (8.6 s)	158 (5.3 s)
changing lane right	663	81	159 (5.3 s)	106 (3.5 s)
changing lane left	711	87	165 (5.5 s)	115 (3.8 s)
starting	401	30	174 (5.8 s)	103 (3.4 s)
stopping	404	26	199 (6.6 s)	123 (4.1 s)

The results on recognizing the previous driving maneuvers are depicted in table 3. Some interesting conclusions to be drawn from the experimental results are:

1. There is a plateau of accuracy that can be reached using car information only. Certain maneuvers –such as passing and changing lanes left– cannot be accurately distinguished using car information only.
2. The context is crucial for recognizing maneuvers such as turnings and lane changes.
3. As shown in[12] in a car simulator, the driver's gaze seems to be strongly correlated with the driver's mental state in real life driving. It is, thus, a relevant feature for driver maneuver prediction, specially in lane changes, passings and turnings.

Table 3. Accuracy for HMMs car only, car and lane and car and gaze data

Accuracy (%)			
	Car	Car and Lane	Car and Gaze
passing	100.0	100.0	100.0
turning right	71.4	85.7	85.7
turning left	0.0	33.3	66.7
changing lane right	0.0	12.5	6.3
changing lane left	29.4	17.6	23.5
starting	100.0	66.7	83.3
stopping	100.0	100.0	100.0

4. **Predictive Power**: The models are able to recognize the maneuver on average 1 **second before** any significant (20% deviation) change in the car or contextual signals take place. Table 4 contains the average prediction power for each of the maneuvers, and figure 5 illustrates through an example what this *predictive power* means. It depicts, frame by frame, the lane feature and the -log(likelihood) of the different models for a passing maneuver. There is no significant change in the lane position until frame 26. However, the models are able to recognize the passing from frame 4 on. In consequence, our driver behavior models are able to anticipate that the passing is going to take place about 2/3 seconds before any significant, perceivable change takes place.

Table 4. Predictive power of the models in frames and secods

Maneuver	Average Predictive Power in Frames (seconds)
passing	37.7 (1.26 s)
stopping	70.7 (2.4 s)
changing lane left	2.1 (.1 s)
turning left	23.0 (.8 s)
changing lane right	20.3 (.7 s)
turning right	15.1 (.5 s)
starting	41.7 (1.4 s)

7. CONCLUSIONS AND FUTURE WORK

In this paper we have described our SmartCar testbed: a real-time data acquisition system in a real car and a machine learning framework for modeling and recognizing driver maneuvers at a tactical level, with special emphasis on how does the context affect the driver's performance. We have validated our theoretical framework with real driving data of over 70 subjects that drove for 1.5 hours in the greater Boston area. We have shown the predictive power of our modeling framework: on average, each of the seven driving maneuvers can accurately be recognized 1 second **before** any significant change in the car signals takes place. We believe that these models would be essential to build more realistic automated cars in car simulators, to improve the human-machine interface in driver assistance systems, to prevent potential dangerous situations and to create more realistic automated cars in car simulators.

We are building more complex models of driver behavior following our mathematical framework. More specifically we have developed extensions of HMMs. Extensions to the basic Markov model generally increase the memory of the system (durational modeling), providing it with compositional state in time. We are interested in systems that have compositional state in *space*, e.g., more than one simultaneous state variable. It is well known that the exact solution of extensions of the basic HMM to 3 or more chains is intractable. In those cases approximation techniques are needed.[15] However, it is also known that there exists an exact solution for the case of 2 interacting chains.[14] We propose using Coupled Hidden Markov Models (CHMMs[1,10]) for modeling two interacting processes (humans[11]

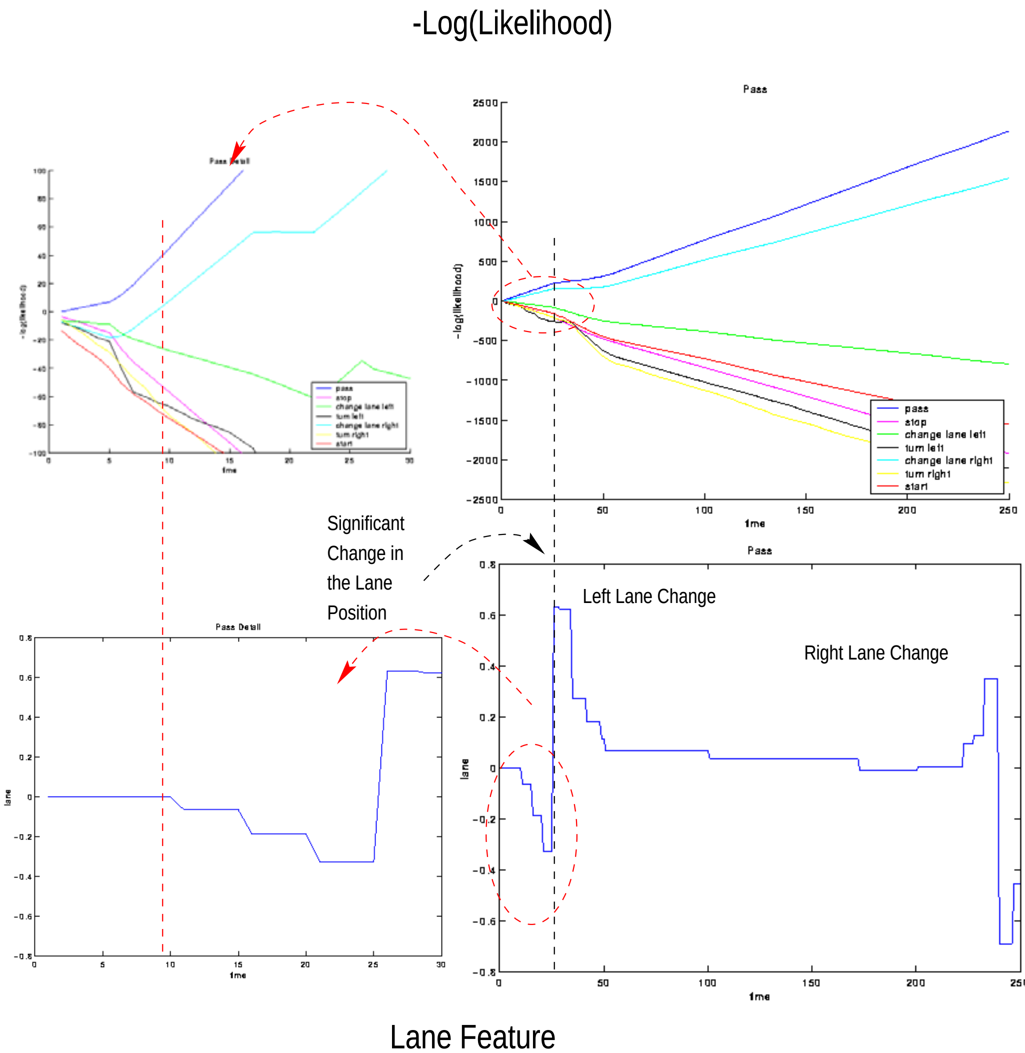

Figure 5. Prediction of a passing maneuver about 2/3 seconds before any significant lane change takes place.

or cars). In this architecture state chains are coupled via matrices of conditional probabilities modeling causal (temporal) influences between their hidden state variables. Figure 2 (b) depicts the graphical model architecture for CHMMs.

In the case of CHMMs, the state posterior probability is given by 5. Note that we need to introduce another set of probabilities, $P_{s_t|s'_{t-1}}$, which correspond to the probability of state s_t at time t in one chain given that the other chain –denoted hereafter by superscript $'$ – was in state s'_{t-1} at time $t-1$. These new probabilities express the causal influence (coupling) of one chain to the other. The posterior state probability for CHMMs is expressed as

$$P(S|O) = \frac{P_{s_1} p_{s_1}(o_1) P_{s'_1} p_{s'_1}(o'_1)}{P(O)} \times \prod_{t=2}^{T} P_{s_t|s_{t-1}} P_{s'_t|s'_{t-1}} P_{s'_t|s_{t-1}} P_{s_t|s'_{t-1}} p_{s_t}(o_t) p_{s'_t}(o'_t) \tag{5}$$

where $s_t, s'_t; o_t, o'_t$ denote states and observations for each of the Markov chains that compose the CHMMs.

In[1,10] a deterministic approximation for maximum *a posterior* (MAP) state estimation is introduced. It enables fast classification and parameter estimation via EM, and also obtains an upper bound on the cross entropy with the full (combinatoric) posterior which can be minimized using a subspace that is linear in the number of state variables. An "N-heads" dynamic programming algorithm samples from the $O(N)$ highest probability paths through a compacted state trellis, with complexity $O(T(CN)^2)$ for C chains of N states apiece observing T data points. The Cartesian product equivalent HMM would involve a combinatoric number of states, typically requiring $O(TN^{2C})$ computations. We are particularly interested in efficient, compact algorithms that can perform in real-time.

Extending the CHMM framework, we propose a graphical model architecture for modeling driver behavior (see figure 6): instead of a symmetric CHMM, asymmetric CHMMs are connected in a lattice structure, where the surrounding traffic affects the behavior of the driver, but not vice-versa. This is just an approximation to the more realistic situation of mutual interactions. The main justification of such an approximation comes from the fact that in our experiments, the driver did indeed modify his/her behavior depending on the surrounding traffic, but not vice-versa.

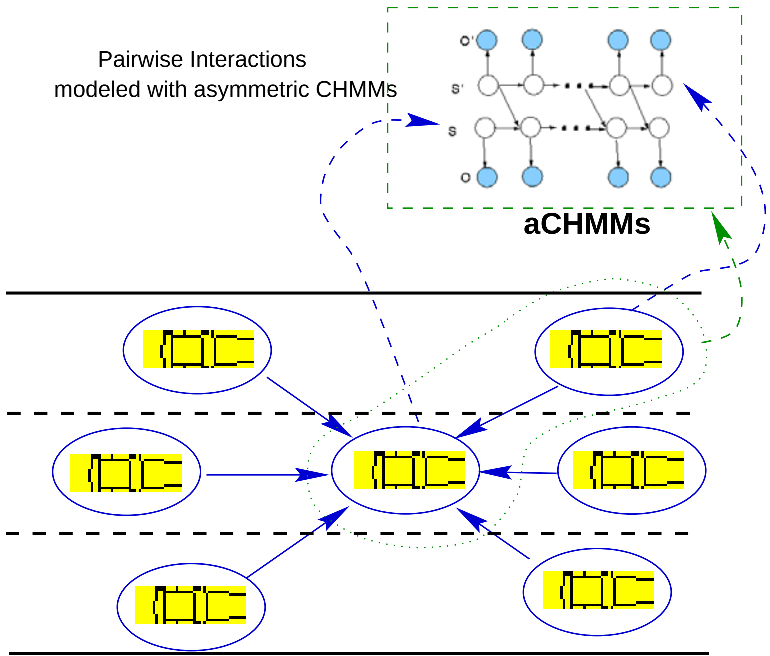

Figure 6. Representation of the asymmetric CHMMs lattice for modeling car interactions

Our preliminary results using CHMMs for driver behavior recognition show that: (1) the performance of CHMMs the same as that of HMMs in the worst case; (2) there are many situations in driving where a single HMM cannot capture the interactions between the driver and the surrounding traffic. It is specially in these cases when CHMMs would offer the greatest advantage. We expect to have statistics on the performance of CHMMs in the next weeks.

7.1. Acknowledgments

Our most sincere thanks to Volvo for donating the car, to Ms. Betty Lou McClanahan for her inestimable help and constant support, and to the instructors and all the MIT SmartCar team members for their help with the experimental part of this work. This research was sponsored by the MIT Media Lab's *CC++* research interest group.

REFERENCES

1. M. Brand. Coupled hidden markov models for modeling interacting processes. *Submitted to Neural Computation*, November 1996.
2. P.S. Fancher, R. D. Ervin, J. R. Sayer, M. Hagan, S. Bogard, Z. Bareket, M. L. Mefford, and J. Haugen. Intelligent cruise control field operational test. *Report No. UMTRI-97-4. The University of Michigan Transportation Research Institute, Ann Arbor, MI.*, 1997.
3. David Heckerman. A tutorial on learning with bayesian networks. Technical Report MSR-TR-95-06, Microsoft Research, Redmond, Washington, 1995. Revised June 96.

4. M. Helander. Applicability of driver's electrodermal response to the design of the traffic environment. *Journal of Applied Psychology*, 63:481–488, 1978.
5. M. Ikegaya, N. Asanuma, S. Ishida, and S. Kondo. Development of a lane following assistance system. In *AVEC 98*, page 9837102, 1998.
6. Michon J.A. *Human Behavior and Traffic Safety*, chapter A critical view of driver behavior models: what do we know, what should we do? Plenum, 1985.
7. N. Kuge, T. Yamamura, and O. Shimoyama. A driver behavior recognition method based on a driver model framework. *Society of Automotive Engineers Publication*, 1998.
8. A. Liu and A. Pentland. Towards real-time recognition of driver intentions. In *Proc. of the 1997 IEEE Intelligent Transportation Systems Conference*, 1997.
9. S. Nakano, Katsutoshi, and Maeda. Development of an intelligent steering system for automated highway systems. In *Proceedings of FISITA 98*, 1998.
10. N. Oliver. *Towards Perceptual Intelligence: Statistical Modeling of Human Individual and Interactive Behaviors*. PhD thesis, Massachusetts Institute of Technology (MIT), Media Lab, Cambridge, MA, 2000.
11. N. Oliver, B. Rosario, and A. Pentland. A bayesian computer vision system for modeling human interactions. In *Proceedings of ICVS99*, Gran Canaria, Spain, 1999.
12. A. Pentland and A. Liu. Towards augmented control systems. In *Proceedings of IEEE Intelligent Vehicles*, 1995.
13. A. Pentland and A. Liu. Modeling and prediction of human behavior. *Neural Computation*, 11:229–242, 1999.
14. Lawrence K. Saul and Michael I. Jordan. Boltzmann chains and hidden Markov models. In Gary Tesauro, David S. Touretzky, and T.K. Leen, editors, *NIPS*, volume 7, Cambridge, MA, 1995. MITP.
15. Padhraic Smyth, David Heckerman, and Michael Jordan. Probabilistic independence networks for hidden Markov probability models. AI memo 1565, MIT, Cambridge, MA, February 1996.