DRIVING MANOEUVRE RECOGNITION

Arati Gerdes
German Aerospace Center,
Lilienthalplatz 7, 38108 Braunschweig, Germany
Arati.Gerdes@dlr.de
(0049) 531-2953479

ABSTRACT

The variety of driving styles in which driving manoeuvres are performed and the uncertainty present in the available sensor data are major obstacles to automatic manoeuvre recognition in the automobile. A Bayesian model of the driving task is used for the probabilistic inference of the manoeuvre being performed from the uncertain evidence available. Top-down inference about the driver's likely intentions is combined with bottom-up inference of the manoeuvre from evidence about control actions and vehicle behaviour. A prototype manoeuvre-recognition system identified the correct manoeuvre during 86.5% of the duration of two trial drives on public roads.

KEYWORDS

Driving manoeuvre recognition, probabilistic reasoning, Bayesian modelling

INTRODUCTION

Information about the driver's intended manoeuvre would enable Advanced Driver Assistance Systems to tune their offered assistance to the driver's real-time needs.

Although automobiles today are equipped with several sensors providing a wide range of data about the driving situation, the task of ascertaining the driver's intended manoeuvre from this data faces two major obstacles. The first obstacle is presented by the sensor data, which though abundant, is often uncertain. Sources of error inherent in today's sensors often render the data unreliable. The second obstacle is presented by the wide range of styles in which manoeuvres are performed, which make the development of a set of rules describing these manoeuvres infeasible.

In order to tackle these problems, a probabilistic model of the driving situation is in development. The model is used to identify the driving manoeuvre being performed, and is intended as the basis for manoeuvre-specific assistance in the automobile. A prototype manoeuvre-recognition system has been built, which uses the model for the inference of the driving manoeuvre. The system is tested with data from drives conducted using a research vehicle on public roads in the Braunschweig area.

An early prototype of the manoeuvre-recognition system used a bottom-up approach to address the problem. Evidence from the driver's control activities and the vehicle behaviour was used to infer the most likely manoeuvre being performed. Although the system performed well, errors still existed [1]. In order to improve the recognition accuracy of the system, top-down inference is added. Given information about the driving situation, and knowledge about driver behaviour, it is possible to infer the manoeuvres that a driver is most likely to perform. The manoeuvre-recognition system is extended to include this inference. In addition to the bottom-up inference of the manoeuvre from supporting evidence, an analysis of the driving situation and a model of driver behaviour in the selection of manoeuvres is used to infer the driver's most likely choice of manoeuvre.

The uncertain data from the sensors is fused in order to create a model of the probable structure of the driving situation [2], and driver behaviour in both the selection and the execution of driving manoeuvres is modelled. The use of probabilistic modelling enables us to make allowances for the variations observed in driver behaviour.

THE TEST BED

Experiments with the automatic manoeuvre recognition system were performed using the ViewCar. The ViewCar is a research vehicle equipped for the synchronised measurement and recording of data from the automobile CAN-Bus and several sensors.

Figure 1 – The ViewCar and the sensors used

The position of the vehicle is measured using a high-precision positioning system. A laser scanner measures the positions and speeds of objects around the ViewCar. A camera-based lane-detection system measures the position of the vehicle in the lane, along with data about the lane width and markings. The driver's gaze direction is measured using a camera-based gaze-detection system. Physiological data about the driver, such as his heartbeat and skin resistance, is measured using sensors on the driver's body. The values returned by these sensors vary in their accuracy. The gaze-detection system and the lane-detection system supply estimates of the momentary accuracy of their output. Cameras focused on the driver and the environment record visual information during the drive, which is manually evaluated after the drive in order to check the performance of the manoeuvre-recognition system.

PROBABILISTIC MODELLING

The problem with modelling human behaviour is that even after years of research, we do not completely understand it. A human passenger in the automobile, observing the driving environment and the actions of the driver, could make an educated guess about the driver's likely intentions, based on his own experience with driving. He would, however, be hard pressed to develop a set of rules that describe the behaviour of the driver, or the set of actions that constitute a driving manoeuvre.

Savage argued for the case of using human knowledge in the form of subjective probabilities as a tool for statistical inference [3]. Combining human knowledge gained from experience, training or hearsay with empirical data about a domain is a key to intelligent behaviour [4]. Bayesian Networks offer us the possibility to perform such inference. A Bayesian Network is a tool for probabilistic reasoning, which allows us to combine data gained from experiments with our prior knowledge about the domain of driver behaviour.

Bayes' Theorem allows us to compute the probability of a hypothesis being true, by combining evidence supporting the hypothesis with our prior knowledge of the domain. Given a hypothesis H and evidence e, the probability of H being true given e can be calculated as:

$$P(H \mid e) = \frac{P(e \mid H) \times P(H)}{P(e)}$$

A Bayesian Network is a directed, acyclic graph. Variables of interest in the domain are represented by nodes. Dependencies between these variables are represented by edges connecting these nodes. The strengths of the dependencies between variables are represented by local conditional probability tables. Given evidence about any subset of the variables in the network, the probabilities of any other subset can be computed [5]. Bayesian Networks are useful not only in modelling the uncertainties present in driver behaviour, but also as a tool for the extraction of information from the uncertain data delivered by the automobile sensors. The inaccuracies present in sensor data render the extraction of information through a Boolean combination of the available values ineffective. Bayesian inference is therefore used for the fusion of the uncertain sensor data from several sensors [2]. The domain of driving manoeuvre selection and execution is represented as a Bayesian Network.

MODELLING THE DRIVING TASK

The three-level driver model

The driving task is viewed as a hierarchical task with three levels: navigation, guidance and control [6].

Navigation refers to the high-level task of route-planning, and the correlation of driving directions to signs and landmarks on the road.

Guidance refers to the selection of the speed and path in response to the road geometry, hazards, traffic and the physical environment.

Control refers to the exchange of information and control inputs between the driver and the vehicle. Most control activities are skill-based, and performed with little conscious effort.

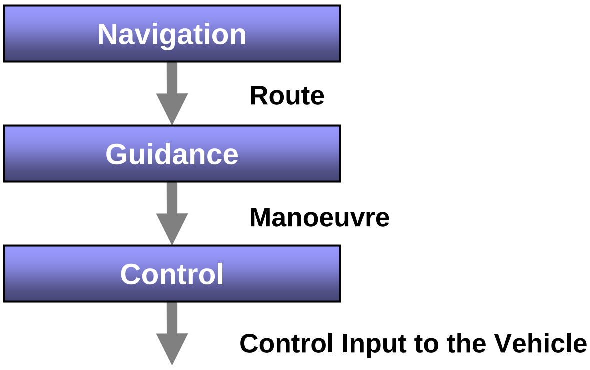

Figure 2 – The Three-Level Driver Model

A Bayesian model of the domain of driving manoeuvres has been developed based upon this three-level model. As evidence from assistance systems and sensors offer us partial evidence about aspects of the navigation and the control level of driving, we combine top-down inference of the expected manoeuvre from navigation information and information about the driving situation with bottom-up inference of the manoeuvre being performed from evidence about the driver's control input and the vehicle behaviour.

Navigation

For the purposes of our model, we assume the use of a navigation system in the automobile. We do not attempt to model the route-planning level of driving, but make the assumption that the desired route is known. Participants in our experiments drove a fixed route. A simulated navigation application was developed which received as input the geographical position of the vehicle as received from the positioning system, and returned the next upcoming turn, freeway entrance or exit along with the distance to the same. In real-life driving situations, the directions suggested by the navigation system are not always followed. In order to account for this variability, the driver's response to such navigation information is handled probabilistically within the model.

Guidance – the selection of a manoeuvre

The selection of a manoeuvre to be performed is taken at the guidance level. Much research has been done regarding an optimal classification of driving manoeuvres. Previous experiments with the classification of the driving task into ten manoeuvres highlighted the need for a hierarchical definition of manoeuvres, wherein complex manoeuvres such as overtaking or entering a freeway are derived from a set of basic, mutually-exclusive manoeuvres [1]. Our model considers the following set of mutually exclusive manoeuvres for classification: stopped, follow road, follow car, merge left, merge right, turn left and turn right. More complex manoeuvres will be derived from these basic manoeuvres in later studies.

Sukthankar describes the tactical level of the task of driving as a “constant battle between long-term goals and real-time constraints” [7]. Manoeuvres such as car following or lane changing are selected so as to satisfy both the long-term goal of reaching a desired destination as well as the real-time constraints such as staying safe and optimizing speed. As the priorities that drivers attach to these differing goals show strong variations depending upon a number of factors including the driver’s mood, a model of driver behaviour must allow for these variations. Probabilistic reasoning allows us to infer the manoeuvres that a driver is most likely to perform in the given situation, allowing for several possibilities with differing likelihoods.

In modelling the driver’s selection of manoeuvre, each of the following goals influences the manoeuvre decision:

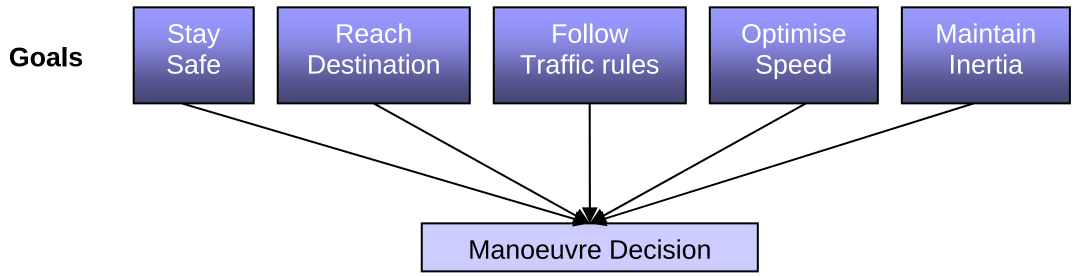

Figure 3 – Competing driver goals

Reach destination

The driver’s long-term goal is to reach her destination. The navigation application provides information about upcoming turns, freeway entrances and exits that the driver needs to take in order to reach the desired destination. Manoeuvres which facilitate the execution of these instructions are selected with a higher probability than the manoeuvres which impede it. In considering the degree of influence exerted by the navigation information on the driver’s decision to change lanes, we consider three driving zones, based on the classification by Gipps [9]:

1. Far away from the turn: The turn is far away, and has no effect on the driver’s behaviour.
2. Approaching the turn: As the turn nears, lane changes to the turning lane are favoured. If the driver is not currently in the turning lane, lane changes to the turning lane will be performed with an increasingly higher probability.
3. Immediately preceding the turn: If the driver is not currently in the turning lane, the probability of a lane change to the turning lane is extremely high. If the driver is in the turning lane, she remains in the lane until the turn is made, and ignores speed considerations.

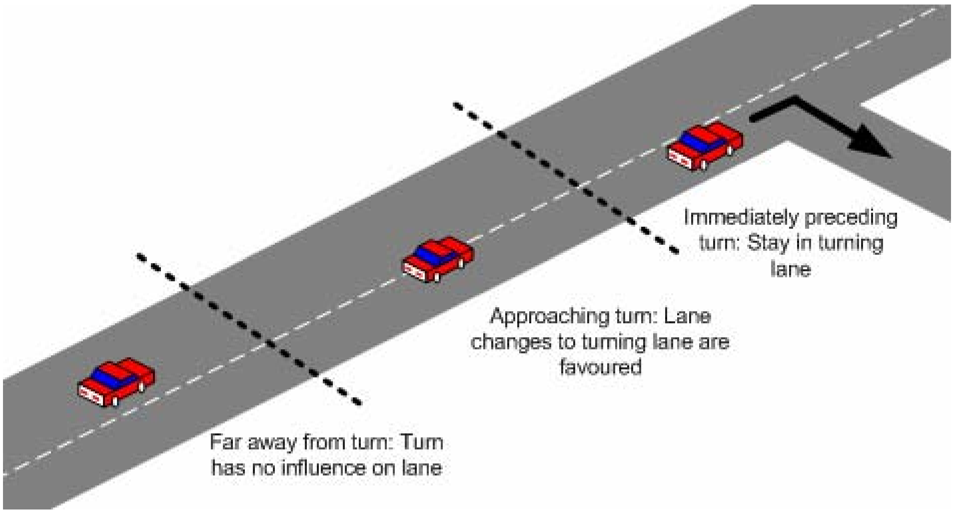

Figure 4 – Three driving zones for mandatory lane changes

Stay safe

Vehicle safety depends upon the vehicle dynamics and the positions and dynamics of other elements in the environment. Certain actions involve more inherent risk than others. Driving into oncoming traffic to overtake, or following a vehicle too closely, are examples of manoeuvres which are inherently risky. While following his goal to remain safe, manoeuvres which involve extremely high levels of risk are performed with lower probabilities than safer manoeuvres. It is important to note that drivers are willing to take some level of risk while driving. Wilde's Risk Homeostasis Theory maintains that drivers continuously evaluate the level of risk to which they believe they are exposed, compare this with the level of risk that they are prepared to accept, and attempt to minimize the difference between these two levels [8]. This implies that drivers do not attempt to minimize their perceived risk, but to maintain this perceived risk at an approximately constant level with which they are comfortable.

Follow traffic rules

The driver has the goal of following the traffic rules. Allowances are made to take into account the drivers who choose to break the rules in order to remain safe, or, in some circumstances, to optimise their speed or to reach their destination.

Optimise speed

The driver has the goal of optimising the speed at which she drives. In addition to the mandatory lane changes demanded by the navigation information, she may also make discretionary lane changes for the purpose of optimising her driving speed [10]. A driver's preferred speed depends upon the speed limit for the road stretch as well as other factors such as the driver's inherent preference for speed and the driver's current mood. A driver's satisfaction with his vehicle's speed is assumed to fall the longer that he remains following a vehicle whose speed is significantly lower than his own preferred speed for the road stretch. Decreasing driver satisfaction increases the probability that the driver decides to change lanes and/or overtake the vehicle ahead. Care is taken to model the differences between lane change behaviour from the left lane to the right lane and vice versa.

Maintain inertia

Inertia refers to the tendency of the modelled driver to remain in execution of the current manoeuvre. As the decision to perform a manoeuvre involves conflict between various goals, the absence of a sense of inertia could leave a driver vacillating between two or more equally attractive possibilities. The goal of maintaining inertia ensures that in the presence of a number of equally advantageous manoeuvres, the driver selects with the highest probability the manoeuvre which he has already begun to execute [7].

Control

Depending upon the manoeuvre selected, the driver communicates his desired actions to the vehicle by means of control inputs via the steering wheel, the gas, brake and clutch pedals, the gears and the indicators. Evidence about the driver's control input to the vehicle offers us a means of inference of the probable manoeuvre being executed.

EXPERIMENTAL RESULTS

A preliminary evaluation of the system was performed with 32 minutes of driving data from two drives by different drivers on public roads in the Braunschweig area. The drives were conducted principally on freeways, with only a small portion being conducted on urban and suburban roads. In total, two left turns, two right turns, 20 lane changes to the left and 22 lane changes to the right were performed, in addition to 51 instances of road following and 38 instances of car following. 132 of the 135 manoeuvres performed (97.7%) were correctly identified by the system. All turns and lane changes were correctly identified, while one instance of car-following and two instances of road-following were falsely classified. With a threshold recognition probability set to 70%, the correct manoeuvre was identified during approximately 86.5% of the test drives. During the drive, the system returned a wrong manoeuvre during 12.7% of the drive duration and an unrecognized manoeuvre during 0.8% of the drive duration.

A common source of error was the failure to detect the vehicle ahead, which resulted in the false classification of the car-following manoeuvre as an instance of road-following. The use of long-range radar, and a fusion of video data with the laser data is planned in order to improve the quality of vehicle-detection. Another source of error observed was the continued indicator signal after an exit into a curved exit ramp, which resulted in a false classification of the following of the curved road as a continuation of the previous lane change performed. The model must be amended to take this situation into consideration.

While the results of the trial freeway drive were promising, it is to be noted that freeway conditions represent an optimal test case for the system. Lane markings on urban and rural roads are less clear and are identified with lower reliability than those on the freeway, and the variations in driver actions possible on urban roads are much larger in comparison with the possibilities on freeways. Testing of the system off the freeway in the preliminary evaluation was minimal. An in-depth evaluation is required in order to tune the system parameters for optimal performance and to evaluate the performance of the system in varying environments.

The results of the preliminary evaluation were positive. The strengths of the method were confirmed, and areas requiring improvement were highlighted. Further research will focus on testing and tuning the system for varying environments and developing methods for improving recognition accuracy. The system will be extended to allow for the detection of more complex manoeuvres, such as the overtake manoeuvre and freeway entrances and exits as combinations of the basic manoeuvres mentioned above with data about the environment.

REFERENCES

[1] Gerdes, A. (2006). Automatic Maneuver Recognition in the Automobile: the Fusion of Uncertain Sensor Values using Bayesian Models. In *Proceedings of the 3rd International Workshop on Intelligent Transportation*, Hamburg.

[2] Gerdes, A. (2006). From Data to Information: Probabilistic Methods for Dealing with Uncertain Sensor Data. To be published in *Proceedings of the 22nd Internationale VDI/VW Gemeinschaftstagung Integrierte Sicherheit und Fahrerassistenzsysteme*, Wolfsburg: VDI Verlag.

[3] Savage, L.J. (1962). *The Foundations of Statistical Inference*, London: Methuen & Co. Ltd.

[4] Pearl, J. (2001). Bayesianism and Causality, or, Why I am only a Half-Bayesian, in D. Corfield and J. Williamson (eds.), *Foundations of Bayesianism*, Applied Logic Series Volume 24, The Netherlands: Kluwer Academic Publishers.

[5] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Belief*, San Mateo: Morgan Kaufman.

[6] Koppa, R.J. (1990). Human Factors. In Gartner, N.H.; Messer, C.J. and Rathi, A. (eds.), *Revised Monograph on Traffic Flow Theory*. TRB Special Report, Transportation Research Board, National Research Council, Washington, D.C.

[7] Sukthankar, R. (1997). Situation Awareness for Tactical Driving. PhD thesis, Robotics Institute, Carnegie Mellon University.

[8] Wilde, G.J.S. (1994). *Target Risk: Dealing with the Danger of Death, Disease and Damage in Everyday Decisions*, Toronto: PDE Publications.

[9] Gipps, P.G. (1986). A Model for the Structure of Lane-Changing Decisions, *Transportation Research* vol. 20B, no. 5, pp. 403-414.

[10] FHWA (1998). CORSIM User Manual Version 1.01, Mc Lean: Federal Highway Administration, US Department of Transportation.