From Data to Information: Die Wahrscheinlichkeitsbasierte Verarbeitung von ungenauen Sensordaten

From Data to Information: Probabilistic Methods for Dealing with Sensor Inaccuracy

A. Gerdes, Braunschweig;

Kurzfassung

Damit Fahrerassistenzsysteme den Fahrer effizient unterstützen können, müssen diese in der Lage sein, die angebotene Assistenz an den Bedarf des Fahrers anzupassen. Mit zunehmender Anzahl und Vielfalt dieser Assistenzsysteme wächst die Menge der Sensordaten über das Fahrzeug, die Fahrumgebung und den Fahrer. Eine Fusion dieser Daten kann den Systemen eine Übersicht über die Fahrsituation zur Verfügung stellen, welche für die situationsabhängige Abstimmung der Assistenz nutzbar ist. Die Rohdaten vieler dieser Sensoren sind intrinsisch unsicher. Statt regelbasierte Datenauswertungen zu benutzen werden deshalb Berechnungen auf Basis der Wahrscheinlichkeitstheorie verwendet, um Schlussfolgerungen aus den unsicheren Daten ziehen zu können. Ein daraus berechneter Wahrscheinlichkeitswert gibt dem Fahrerassistenzsystem einen Hinweis auf die Verlässlichkeit der erhaltenen Informationen.

Abstract

In order to efficiently assist the driver, assistance systems need to be able to tune the offered assistance to the driver’s needs. As the number and variety of assistance systems in the automobile increase, we are increasingly supplied with data from various sensors about the state of the vehicle, the driver and the driving environment. A fusion of the available data can provide the systems with a valuable overview of the driving situation, which can be used for situation-based tuning of the assistance offered. The raw data returned by several of the sensors is, however, intrinsically uncertain, making a simple, rule-based evaluation of the data impractical. Computations based on probability theory are therefore used to draw inferences from the uncertain data available. The probability value assigned to the inference drawn offers an assistance system an indication of the level of confidence that may be associated with the information received.

1. Introduction

With a continuous increase in the traffic density on the roads, driver safety and comfort have gained increasing importance. Advanced Driver Assistance Systems support the driver in various aspects of the driving task, in an effort to optimize the driver's workload. In order to provide a driver with relevant assistance, the assistance system must be able to tune the offered assistance to the driver's needs. Information about the driving situation needs to be made available in order to ensure that systems offer assistance that is relevant to the situation, and do not distract a driver during the performance of a difficult task.

With the increasing number of assistance functions customary today, the automobile is now commonly fitted with several sensors supplying data about the vehicle, the driver and the driving environment. While data is abundant, a major obstacle to the extraction of information is presented by the high degree of uncertainty present in the data delivered by the sensors. The data delivered by the automobile sensors today is sometimes imprecise or inaccurate, and in certain cases, can be false. The aim of this work is to construct a system for the automatic identification of the driving maneuver being performed using a fusion of the uncertain sensor data available.

2. The ViewCar and the Sensors Used

Driver behaviour analysis and testing of the prototype system are performed using the ViewCar. The ViewCar is a research vehicle equipped with several sensors and cameras for the measurement and recording of data about the driver, the vehicle and the driving environment.


Figure 1: The DLR ViewCar

A camera-based lane-detection system is used to measure the type and width of the lane markers, the width of the lane and the position of the vehicle in the lane. The system is relatively robust, but its performance deteriorates when the lane markings are faint or when lighting is poor. As camera sensors for use in the automotive industry have only recently migrated from laboratories with controlled lighting, commercial camera systems available today are often unable to cover the full range of brightness variations experienced in outdoor scenes. Grease on the camera lens can also cause deterioration in performance, as it causes a blurring of the captured image [1]. The lane-detection system was observed to perform best on freeways where lane markings are usually clearly defined. When drives were conducted at dusk, or when bright sunlight fell directly on the camera, its performance was observed to deteriorate.

Figure 2: The camera-based lane-detection system and the laser scanner

A laser scanner is used to measure the positions and speeds of vehicles ahead of the ViewCar. Although the data received is useful in producing an estimate of the traffic around the ViewCar, errors can cause the data to be unreliable at times. Dirt on the surface of the sensor, rain, fog and snow are all possible sources of error. These particles can also reflect the laser pulse, resulting in the production of irrelevant data [2]. The pitching motion of the vehicle, caused by uneven road surfaces, strong brake application or strong acceleration could also result in relevant objects being missed during measurement [3]. In newer sensors, efforts are made to combat these errors by sending laser pulses on multiple planes instead of a single one, and requiring multiple echoes for recognition.

Figure 3: Possible sources of laser scanner error. Figure based on [2] and [3].

A positioning system is used to measure the geographical position of the vehicle. Differential GPS is used to correct errors common to the receiver in the vehicle and a receiver at a fixed base station, such as the errors due to atmospheric refraction of the signals, changes in the orbital path of the satellites, and clock inaccuracies. Multipath and receiver noise errors, however, remain. In order to correct these errors, an inertial platform measures the movement of the vehicle and uses a Kalman filter to estimate the correct position. An odometer is used to correct the drift in the calculated position caused by rounding errors in the measured velocity.

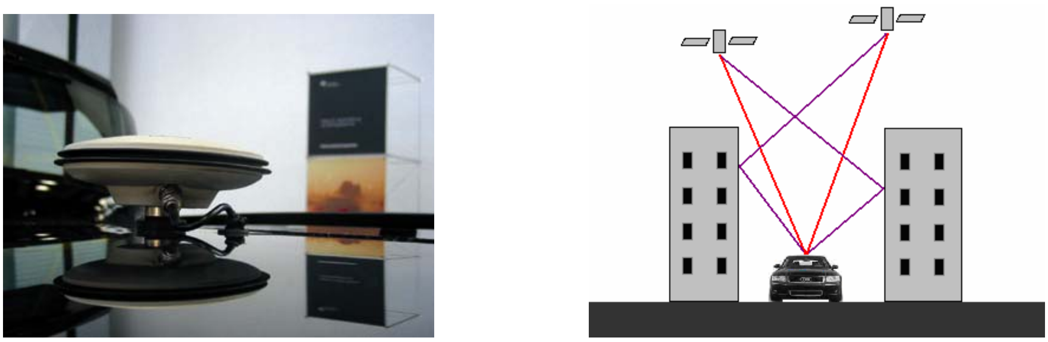

Figure 4: The positioning system and GPS multipath error.

3. Data Fusion

Although the data received from individual sensors contains inaccuracies and uncertainties, a fusion of the data for the extraction of information has several advantages. A combination of the observations of independent sensors measuring the same variable can enable an estimation of the reliability of the data. The fused data from several sensors enables the creation of a joint context in which the data from the individual sensors can then be interpreted. The combined system displays increased robustness, as one sensor can contribute data when another is unavailable or inoperative. The fusion of data from a set of “average-performance” sensors could provide the same level of performance as a single, highly-reliable sensor at a lower cost [4]. Our aim is to fuse the uncertain data from the various sensors in order to infer the driving maneuver being performed.

4. Probabilistic Reasoning

As conclusions may not be drawn from a Boolean combination of the variables involved, we base our computations on probabilities. Bayesian Networks are a useful tool for probabilistic reasoning. Bayesian Networks combine the classical theoretical framework of probability theory with the support of computational constructs in order to draw inferences from uncertain data. A Bayesian Network is a directed, acyclic graph, representing a domain of interest. Nodes of the graph represent the variables of interest in the domain, while the edges of the graph represent the causal dependencies between the nodes. The strengths of these dependencies are described by local conditional probability tables [5].

The core of Bayesian Inference lies in Thomas Bayes' inversion formula. Given a hypothesis H and evidence e, the conditional probability of H being true given e can be calculated as:

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

Bayes' Theorem can be used to compute probability values for uncertain system variables by combining the data provided by several sensors. If a single sensor supplies data with a degree of reliability too low to allow a decision, the combination of supporting data from other sensors can increase our confidence in the data to a degree which renders it useful [6]. If the sensors supply conflicting evidence, the result of the fusion will favor the sensor possessing the better historical output quality, while the conflicting evidence will result in a reduction of the probability associated with the output.

Given a hypothesis H and two pieces of evidence e_1 and e_2 from two independent sensors, Bayes' Rule for the fusion of this evidence states that the probability of the hypothesis H being true given the two independent pieces of evidence e_1 and e_2 can be expressed as [7]:

$$P(H \mid e_1, e_2) = \frac{P(e_1, e_2 \mid H) \times P(H)}{P(e_1, e_2)}$$

As e_1 and e_2 are independent of one another,

$$P(H \mid e_1, e_2) = \frac{P(e_1, \mid H) \times P(e_2, \mid H) \times P(H)}{P(e_1) \times P(e_2)}$$

The importance of this equation is that it enables us to update our belief in a hypothesis in response to evidence received from multiple, independent sources. It allows us to calculate the value $P(H \mid e_1, e_2)$, a term usually difficult to compute, in terms of quantities that may be directly determined from experiments. $P(H)$, $P(e_1)$ and $P(e_2)$ are the prior probabilities of the hypothesis and each piece of evidence, while $P(e_1 \mid H)$ and $P(e_2 \mid H)$ represent the probability of observing each piece of evidence when the hypothesis is true.

An important feature of Bayesian Networks is their support of bidirectional inference. Inferences in a Bayesian Network can be derived by propagating information in any direction.

Given evidence about any subset of its variables, the network supports the computation of the probabilities of any other subset [8].

If the lane sensor, for example, returns uncertain evidence about the width of the lane, independent evidence that the current road is a freeway can be used to infer the probable width of the lane.

5. The Prototype Maneuver-Recognition System

A prototype maneuver-recognition system is in development at the DLR. Data from the automobile data bus, the positioning system, the lane detection system and the laser scanner is input to the system. A reference test stretch of road was defined for the test drives. Static environmental information about the stretch was recorded. Input from the positioning system is mapped to road information about the speed limit, the number of lanes and the presence of a ramp at the identified location. The input information is used to build an internal representation in the system of the driving situation and the driver’s actions. A Bayesian model of the domain of driving maneuvers is used to map the acquired evidence to the most likely driving maneuver. The output of the system is the recognized maneuver along with its probability.

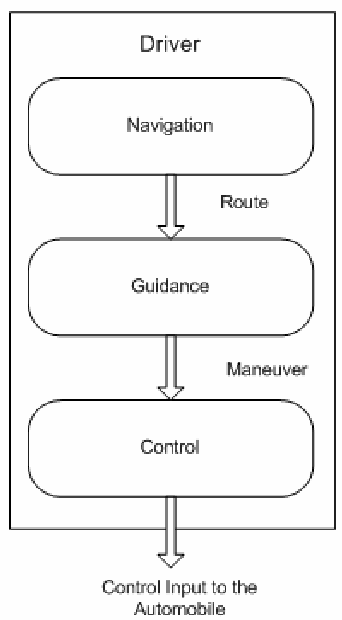

Figure 5: Three-level hierarchical driver model

The Bayesian Network is a simplified model of the domain of driving maneuvers. The driving task is considered to be a hierarchical task with three levels. Navigation is the high-level task of route-planning, and the correlation of driving directions to signs and landmarks on the road. Guidance is the selection of 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 [9].

In the initial prototype, ten driving maneuvers were selected for recognition: stopped, follow road, follow car, turn left, turn right, merge left, merge right, overtake, enter freeway and exit freeway. A Bayesian model of the control and guidance levels of the driver model was used as the basis of inference. The results of early evaluations of this prototype, while promising, indicated a need for improvement [10]. In an evaluation of 40 minutes of driving data in the Braunschweig area, 92.1 % of the maneuvers driven were recognized by the system. The correct maneuver was recognized during 77% of the drive duration. An unrecognized maneuver was returned during 9.2% of the drive duration, and a false maneuver was recognized during 13.7% of the drive duration.

In order to improve the recognition accuracy of the system, a more detailed model of the domain is in development, which includes the top-down inference of the expected driving maneuver. The model used in the first prototype was restricted to the guidance and control tasks, and maneuver recognition was performed using a bottom-up approach. Evidence about the driver's actions was used to infer the probable maneuver being performed.

In the current phase of development, top-down inference is being added to the system. Navigation information in the form of the planned route has been added to the system. The model has been extended to supply it with the ability to infer expected maneuevers and their probabilities, given information about the state of the driving situation and the driver's planned route. As such inferences are extremely dependent upon characteristics of the driver, care is taken to include a high level of variance in the probability distributions used. The goal is to support the bottom-up inference, not to override it.

6. Lane Identification

In order to be able to use the information about the driver's planned route and the traffic situation to predict lane changes, information about the lane on which the driver is traveling is required. While the lane-detection sensor returns information about the width and markings on the current lane, it does not tell us which lane of the road the vehicle is currently using.

In order to support the top-down inference of the performance of a lane change maneuver, a system has been developed to infer the current lane upon which the vehicle is traveling. The system uses vehicle data from the automobile data bus, data about the lane markings and the width of the current lane from the lane-detection sensor, along with information about the number of lanes on the road from the geographic data in order to infer the current lane. A probabilistic fusion of the data is performed here as well, as uncertainties are present both in the geographical data and in the lane data. The system was designed for use on rural roads and freeways. Lane markings in the cities are often faint, and the resulting probabilities are too low to be used as meaningful information.

The current lane was classified into one of the following categories: the only lane, the left lane, the middle lane, the right lane or the right lane with a neighboring exit. The lane-identification system was tested with data from a 35-minute trial drive performed with the ViewCar on roads in the Braunschweig area. 7.4 minutes of the trial drive took place on suburban and/or rural roads, while 27.6 minutes were driven on the freeway. The recognition threshold was set to 60%. The correct lane was identified during 88% of the duration of the drive. A wrongly identified lane was returned during 8.6% of the drive. The recognition probability lay below the 60% threshold during 3.4% of the drive.

On the freeway, where the performance of the lane-recognition system was fairly robust, data regarding the lane markings and widths was clear, and supported the inference of the lane being used. On the smaller, single-lane roads where lane markings were less prominent and sometimes unavailable, geographic data supplied the information that the vehicle was traveling on the road’s only available lane. Errors in lane identification were commonly observed during the execution of a lane change maneuver.

7. Perspectives

The positive results observed with early prototypes of both the maneuver-recognition system and the lane-identification system support the viability of the Bayesian approach to information extraction. Early evaluations of the prototypes provide good indications of the areas where improvement is required.

The models must be tuned in order to achieve higher recognition rates with the systems. Several of the probability distributions used in the prototype models are approximations based on the observed frequencies of the values. A detailed statistical analysis of the variables must be performed, and the probability distributions in the models need to reflect these results, in order to improve the recognition accuracy.

The experiments with maneuver recognition highlighted the importance of maneuver classification. In the initial classification, the ten maneuvers selected were defined as being mutually exclusive. Evaluation of the data illustrated the impracticality of the approach. Despite definitions that attempted to clearly demarcate the differences between the maneuvers, both human observers and the recognition system alike were at times unable to unambiguously classify certain maneuvers. A clear, well-defined classification of the driving task is required, wherein basic, mutually exclusive maneuvers are defined as such, while others are treated as special cases of one or more of these basic maneuvers, and identified based on the attributes of the driving situation.

In the case of the lane-identification system, the lane-change state must be explicitly modeled so as to eliminate the observed errors in lane identification during the performance of the maneuver.

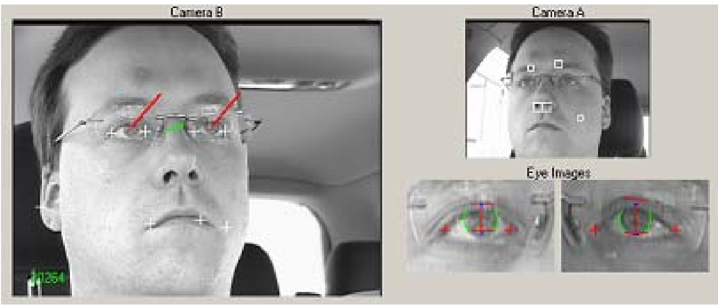

Figure 6: Driver gaze behavior: an indicator for driver intent

Researchers have described the usefulness of driver-gaze behavior in predicting driver intent [11], [12]. A camera-based gaze-detection system in the ViewCar measures driver gaze behavior. Concurrent research at the DLR focuses on the analysis of the driver’s gaze behavior, and the possibility of using the driver’s gaze information to infer her intentions. The incorporation of this information into the system paves the way for the prediction of maneuvers prior to their execution.

[1] C. Koch, J. Yoon, and N. Lii, “Evaluation of Vision Based In-vehicle Applications”, in *Proceedings of the 3rd International Workshop on Intelligent Transportation*, Hamburg, 2006.

[2] "Ibeo: Technology: Bad Weather Performance", Internet. http://www.ibeo-as.de/technology_d_badweather.asp [August 3, 2006]

[3] "Ibeo: Technology: Multilayer technology", Internet. http://www.ibeo-as.de/technology_d_multilayer.asp [August 3, 2006]

[4] D. J. Dailey, P. Ham and P. Lin, "ITS Data Fusion", Final Research Report, Research Project T9903, Task 9, ATIS/ATMS Regional IVHS Demonstration, prepared for the Washington State Transport Commission, Seattle, 1996.

[5] J. Pearl, "Bayesian Networks", UCLA Cognitive Systems Laboratory, Technical Report (R-216), Revision 1, In M. Arbib (Ed.) Handbook of Brain Theory and Neural Networks, Cambridge: MIT Press, 1995.

[6] D. Koks, S. Challa, "An Introduction to Bayesian and Dempster-Schafer Data Fusion", *Sadhana – Academy Proceedings in Engineering Sciences*, vol. 29, pp 145 – 174, 2004.

[7] B. Yu, K. Sycara, J. Giampapa, S. Owens, "Uncertain Information Fusion for Force Aggregation and Classification in Airborne Sensor Networks", in *AAAI Workshop on Sensor Networks*, AAAI Press, July 2004.

[8] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, San Mateo: Morgan Kaufmann, 1998.

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

[10] A. Gerdes, "Automatic Maneuver Recognition in the Automobile: the Use of Bayesian Models for the Fusion of Uncertain Sensor Values", in *Proceedings of the 3rd International Workshop on Intelligent Transportation*, Hamburg, 2006.

[11] N. Oliver and A. P. Pentland, "Graphical Models for Driver Behavior Recognition in a Smart Car", in *Proceedings of the IEEE Conference on Intelligent Vehicles*, Detroit, 2000.

[12] D.D. Salvucci, E.R. Boer and A. Liu, "Toward an Integrated Model of Driver Behavior in a Cognitive Architecture", *Transportation Research Record*, No. 1779, 2004.