

Generic Driver Intent Inference based on Parametric Models

Martin Liebner, Christian Ruhhammer, Felix Klanner and Christoph Stiller

Abstract—Reasoning about the driver intent is fundamental both to advanced driver assistance systems as well as to highly automated driving. In contrast to the vast majority of preceding work, we investigate an architecture that can deal with arbitrary combinations of subsequent maneuvers as well as a varying set of available features. Detailed parametric models are given for the indicator, velocity and gaze direction features, all of which are parametrized from the results of extensive user studies. Evaluation is carried out for continuous right-turn prediction on a separate data set. Assuming conditional independence between the individual feature likelihoods, we investigate the contribution of each feature to the overall classification result separately. In particular, the approach is shown to work well even when faced with implausible observations of the indicator feature.

Index Terms—Driver intent inference, active safety.

I. INTRODUCTION

The aim of active safety systems is to react early enough to prevent accidents from happening while still maintaining a sufficiently low false alarm rate. Actively intervening systems such as automatic braking or evasive steering enjoy the benefit of very low reaction times that enable them to be activated only when the conflict with another traffic participant is almost certain. On the negative side, however, they pose very strong demands on the car's sensors as it must be ascertained that the evasive maneuver will not make the situation worse. Therefore, an alternative approach is to warn the driver of potential conflicts early enough to allow him solve the situation on his own. Once he has had time to react, the driver's ability to assess the traffic situation is assumed to be superior to that of the car. In addition, the driver can verify any information provided by the sensors before taking action, so the requirements on the reliability of the sensor data are considerably lower than for actively intervening systems.

One major challenge that arises from the driver's reaction time is the need to predict the current traffic situation up to several seconds into the future. Especially at inner city intersections, this cannot be done without knowledge of the driver's intent as well as that of the other traffic participants. While today's systems are still limited by insufficient knowledge of the vehicle's environment, such predictions might soon become feasible as research projects such as Ko-PER [1] and sim^{TD} [2] are currently investigating methods to augment the vehicle's onboard sensors with information

received via Car2X-communication. Also, advances in sensor equipment are to be expected as more research is carried out on autonomous and highly-automated driving. As such systems grow more sophisticated, better sensors will eventually make the car see as much or even more than the driver himself.

A. Related work

Motivated by this prospect, driver intent inference has been an important research issue for more than a decade by now. Common approaches include discriminative methods such as support vector machines [3], relevance vector machines [4], conditional random fields [5] and neural networks [6] as well as generative approaches such as hidden Markov models [7] and dynamic Bayesian networks [8], [9]. Discriminative approaches usually aim to predict a single type of maneuver while generative models are more often applied if the probability distribution over a set of available maneuvers is to be inferred [10]. In simple-structured environments such as on highways, this set may be predefined whereas more complex inner-city scenarios usually require a digital map representation of the environment [11], [12], [13].

B. Problem addressed

In contrast to the vast majority of preceding work, we investigate an architecture that creates hypotheses about possible future paths rather than single maneuvers. Hypotheses can therefore include arbitrary combinations of subsequent maneuvers as well as several instances of the same type of maneuver. This allows us to predict combinations such as "lane-change right and then turn right" as well as to infer not only the probability of a lane change itself, but also the distance at which it is most likely to occur. While the former might be helpful for inner-city scenarios, the latter is crucial for risk-assessment on highways. In order to describe the expected driver behavior with respect to each out of an arbitrary set of features, we use simple parametric models that make use of contextual information and therefore generalize well to arbitrary situations.

The remainder of the paper is organized as follows: In Section II, we describe the general architecture of our driver intent inference system. The parametric likelihood functions for the indicator, velocity and driver gaze direction feature are given in Section III, IV and V. In Section VI, we evaluate the performance of our system for continuous right-turn prediction and investigate the contribution of each of the features to the overall classification result. Finally, Section VII concludes this paper.

M. Liebner, C. Ruhhammer and F. Klanner are with BMW Group, Research and Technology, D-80788 Munich, Germany. firstname.lastname@bmw.de

Christoph Stiller is with Karlsruhe Institute of Technology, Department of Measurement and Control, D-76131 Karlsruhe, Germany. stiller@kit.edu

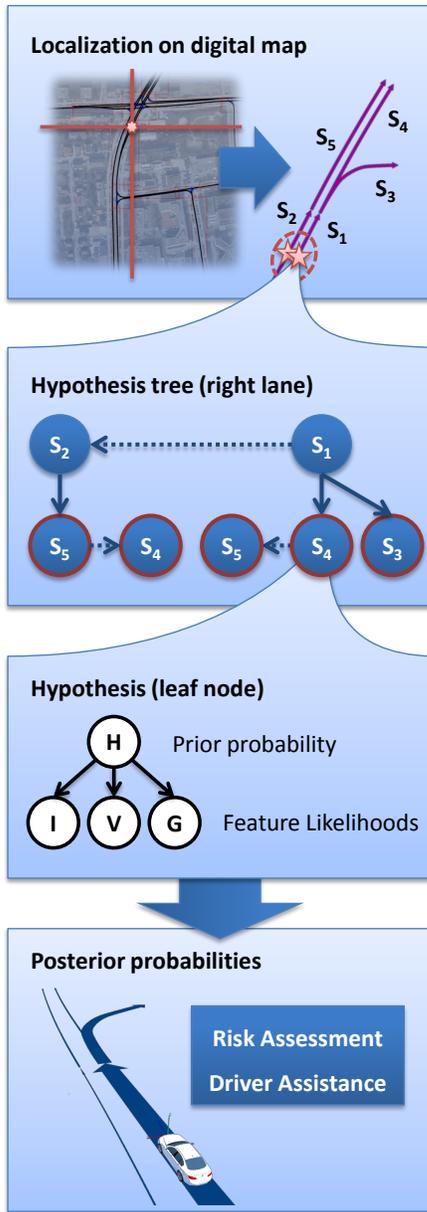


Fig. 1. System overview.

II. GENERAL ARCHITECTURE

Figure 1 shows the general architecture of our approach as well as the steps needed to obtain the posterior probabilities for each possible driver intent.

A. Self localization

The first step is to match our current position to a high precision digital map that represents each lane by the average path of vehicles driving on that lane [14]. In Ko-PER, several self-localization methods are investigated, including cooperative sensor technology, laser scanner landmarks and tightly-coupled GPS. The accuracies range from a few centimeters up to 2 or 3 meters, so the current lane cannot always be uniquely identified. Instead, we need to map the normal

distributed position probability density function $f_X(x|\mu, \Sigma)$ to probabilistic lane assignments $P(L|\mu, \Sigma)$.

To do so, we use discrete particles q_{ij} to approximate f_X by probabilities $P(q_{ij}|\mu, \Sigma)$ that are equal to the integral of f_X over the area assigned to q_{ij} . Choosing rectangular areas aligned with the eigenvectors and scaled with the eigenvalues of Σ , $P(q_{ij}|\mu, \Sigma)$ can be made independent of μ and Σ and thus be stored in a look-up table for performance. For each particle, we then determine the longitudinal as well as orthogonal distances with respect to the nearest lane segments and assign probabilities $P(q_{ij}|L)$ based on the assumption that the position of vehicles driving on a lane is normal distributed around its center. Therefore, we have

$$P(L|\mu, \Sigma) = \sum_{i,j} \frac{P(q_{ij}|L) P(L)}{P(q_{ij})} P(q_{ij}|\mu, \Sigma)$$

with $P(q_{ij})$ the prior probability of each particle, which is the same for all particles iff they all have the same size, and $P(L)$ a uniform prior for the lane assignment. For each lane, a single most likely current distance is calculated from the weighted mean of the longitudinal distances of each particle q_{ij} with respect to that lane. In Figure 1, the localization results are visualized as red stars on the two neighboring lanes.

B. Hypothesis tree

For each localization result, we then construct a hypothesis tree by recursively extracting neighboring and connecting lane segments from the digital map up to a predefined distance from our localization result. Each node in the tree refers to exactly one lane segment S_i , but one lane segment can be referred to by several nodes as there might be several ways to reach it. Connecting lane segments are represented as children of the current node, whereas horizontal neighbours represent lane change maneuvers. We assume that only a single lane change is conducted during the length of each lane segment. Therefore, the hypothesis tree is a directed acyclic graph and each of its leaf nodes represents a distinct hypothesis on the vehicle's future path.

C. Prior probabilities

Prior probabilities are propagated top to bottom throughout the hypothesis tree. Starting with the tree root's value of $P(L|\mu, \Sigma)$, the prior probability of each parent node is uniformly distributed to its children. Children have to give part of their prior probability to their horizontal neighbors so as to account for the probability of a lane change. Assuming sufficiently short segments, the lane change probability can be approximated as being proportional to the remaining length of the corresponding lane segment with respect to the current position. For our experiments, we assume a statistical lane change rate of 1/500 m.

D. Posterior probabilities

In order to calculate the posterior probability of each hypothesis, we rely on three different features: The current status of the indicator signal as well as the time since its last

activation (I), the velocity profile of the past 1.4 s (V) and the driver's gaze direction measured by his head heading angle within the last 1.0 s (G). Assuming conditional independence given the hypothesis, the joint probability distribution over all possible hypotheses H and observations $O = \{I, V, G\}$ can be modeled as a naive Bayes classifier. Hence, the posterior probability distribution over H calculates to

$$P(H|O) = \frac{\prod_i P(O_i|H)}{P(O)} P(H). \quad (1)$$

The denominator is the same for all hypotheses and can therefore be seen as a normalization constant. Our observations are actually from the continuous domain, so in the following we denote the likelihoods $P(O_i|H)$ as probability density functions $f_{\text{Ind}}^{(h)}$, $f_{\text{Vel}}^{(h)}$ and $f_{\text{Gaze}}^{(h)}$ with $h \in H$.

E. Path probabilities

After calculating the posterior probability for each leaf node based on the parametric likelihood functions described in the following sections, the posterior probability of all other nodes can be obtained as the sum over the posterior probabilities of their children and their children's horizontal neighbours. A subsequent risk assessment algorithm might use the approach described in [14] to simulate the future velocity profile along each path so as to obtain both probability and time of potential conflicts with other traffic participants.

III. INDICATOR MODEL

While the indicator is probably the most obvious means to predict lane changes and turn maneuvers, some researchers argue that it should not be used as a feature at all since accidents occur especially in those cases when the indicator is not activated although it should [11]. Our perspective is that even though this is true for some scenarios, in others it may help us predict and avoid an accident that could not have been predicted at reasonable false alarm rates otherwise. We agree, however, on that the indicator likelihood function should explicitly account for the indicator's misuse.

In addition to its mere status, the indicator signal carries the information about the time and distance since its last activation. In the following, we will make use of this information to create a feature that can help to predict the exact time of a lane change, to distinguish between several possibilities to turn right and even tell intentional and random indicator activation apart.

A. Input variables

To be more robust in stop and go situations as well as at traffic lights, we chose to use distances rather than time differences for our modeling. Beside the current indicator status, we thus have the following input to our indicator feature likelihood calculation:

- s_C Current distance along the path.
- s_A Distance of last indicator activation.
- s_T Distance to the turn's crotch point.
- s_0 Start of lane change lane segment or, if $s_0 < s_C$, the current distance s_C .
- s_1 End of lane change lane segment.

We assume that only the next oncoming lane change or turn maneuver is relevant for the indicator feature. Within the hypothesis tree, nodes that represent turn maneuvers feature a so called crotch point which represents the distance along the lane segment at which it first reaches a distance of 1.5 m from the lane segment for going straight. Crotch points serve both as a reference for the indicator activation distance as well as a means to identify turn maneuvers along the path. Note that all distances may be provided with respect to an arbitrary reference point along the path.

B. Model for random indicator activation

In (1) we multiply individual feature likelihoods in order to obtain the overall probability of each hypothesis. To avoid this probability to calculate to zero on account of just a single feature, our model must explain every possible observation however unlikely it might be. This includes both unintentional as well as – based on our limited knowledge of the environment – inexplicable indicator activations.

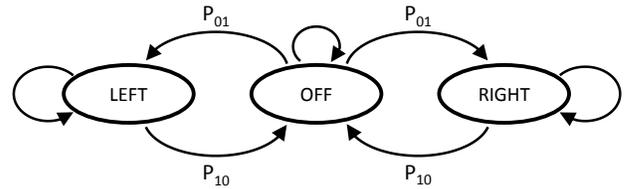


Fig. 2. Markov process of random indicator activation. In each time step, the indicator can be either activated, deactivated or left at the current status.

For the case of drivers setting the indicator although going straight, we assume that they do so on account of the Markov process shown in Figure 2. We assume that the chance of a random indicator activation within the interval $[s, s + \Delta s]$ is represented by the same probability P_{01} in both directions, and P_{10} for its deactivation respectively. The probability of observing an indicator activation at distance s relative to our current location s_C is hence given by

$$f_R(s, s_C) = P(\text{OFF}) P_{01} (1 - P_{10})^{(s_C - s)/\Delta s}, \quad (2)$$

with $P(\text{OFF})$ being the stationary probability of state OFF for the given process. As (2) resembles an exponential distribution, it can be rewritten as

$$f_R(s, s_C) = P_R \lambda e^{\lambda(s - s_C)} \quad (3)$$

with $\lambda = -\log(1 - P_{10})/\Delta s$ and P_R the overall probability of the driver activating the indicator in a particular direction when going straight. For our experiments, we assumed $P_R = 0.02$, $P_{10} = 1/200$ and $\Delta s = 1$ m. The resulting probability density function is visualized in Figure 3.

C. Turn related indicator activation

In order to capture typical distances of indicator activation before turning, we collected more than 200 right turn maneuvers conducted by 6 subjects on 5 intersections. We found the distance normal distributed with $\mu_T = -55.6$ m and $\sigma_T = 25.3$ m. The results are shown in Figure 4.

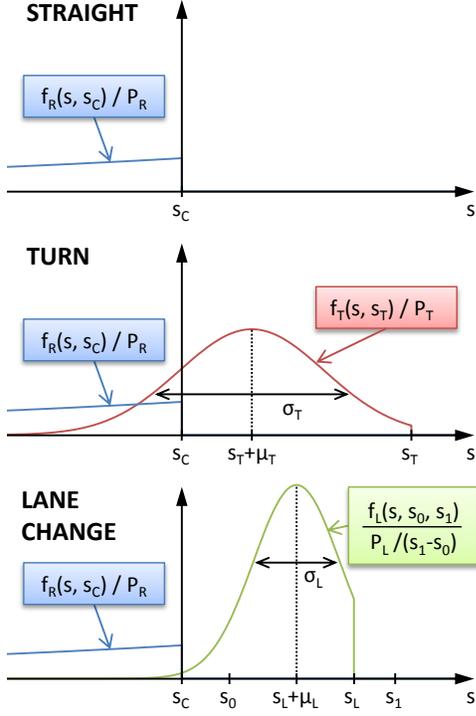


Fig. 3. Likelihood functions and their input parameters.

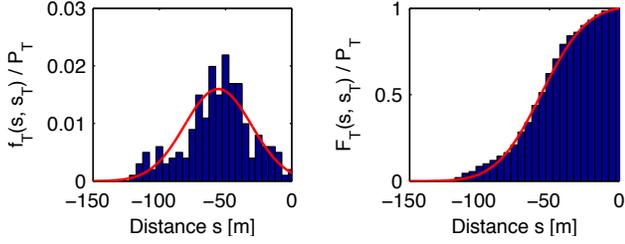


Fig. 4. Probability density and cumulative distribution for turn related indicator activations as observed in the data (bars) and according to our model (line).

As to be expected in a supervised user study, we did not observe any right turns without indicator activation. Assuming that they account for about 20% of all right turn maneuvers as reported in [15], and given that some of the observed indicator observations have to be accounted to the random process described in the previous section, we have an overall probability of a turn related indicator activation of $P_T = 1 - 0.20 - P_R = 0.78$. For $s < s_T$, the indicator activation likelihood is therefore given by

$$f_T(s, s_T) = \frac{P_T q_T}{\sqrt{2\pi}\sigma_T} \exp\left[-\frac{1}{2}\left(\frac{s - s_T - \mu_T}{\sigma_T}\right)^2\right] \quad (4)$$

with

$$q_T^{-1} = \int_{-\infty}^{s_T} \frac{f_T(s, s_T)}{P_T q_T} ds = \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{-\mu_T}{\sqrt{2}\sigma_T}\right) \right] \quad (5)$$

to normalize the original distribution as there are no turn-

related indicator activations for $s \geq s_T$ and

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \quad (6)$$

the error function for which there are efficient numerical implementations. The cumulative distribution function, used to determine the likelihood of not having observed an indicator activation at the current distance s_C relative to the turn distance s_T , hence calculates to

$$F_T(s, s_T) = \frac{P_T q_T}{2} \left[1 + \operatorname{erf}\left(\frac{s - s_T - \mu_T}{\sqrt{2}\sigma_T}\right) \right]. \quad (7)$$

In this paper, we assume that f_T and F_T are valid for left as well as right turn maneuvers. The combined likelihood functions are given in Table I. Their logarithms are visualized in Figure 5. Our model guarantees that for high values of s_A and s_C the indicator likelihoods for all hypotheses become the same.

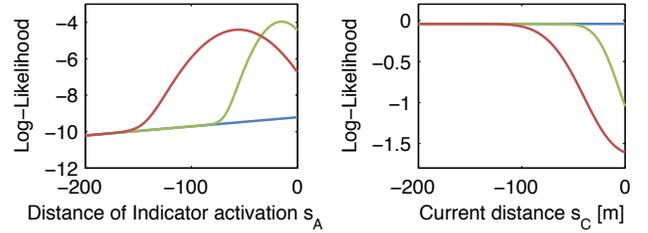


Fig. 5. Example likelihoods for the case that the indicator is activated in the direction of travel (left), and for the case that it is not (right). The curve rising and falling first is that of a turn hypothesis, the second that of a possible lane change. The last corresponds to going straight.

D. Indicator activation due to lane changes

While indicator activations due to lane changes can be modeled in a similar fashion as those caused by turn maneuvers in principle, our investigations show that they are motivated by the time rather than by the distance to an oncoming lane change. Based on an overall of more than 500 lane changes, the time of indicator activation is visualized in Figure 6. Apparently, the mean indicator activation time is more or less independent of the current velocity. We

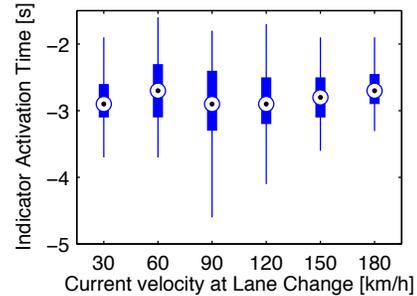


Fig. 6. Indicator activation time relative to the line crossing. Boxes represent the range between the first and the third quartile. The dot within the box is marking the median. Whiskers are drawn up to 1.5 times the box size so as to represent the 99.3% interval for normal distributed data.

TABLE I
LIKELIHOOD $f_{\text{Ind}}^{(h)}$ OF THE INDICATOR SIGNAL OBSERVATION

	LEFT	OFF	RIGHT
Go Straight	$f_R(s_A, s_C)$	$1 - 2 P_R$	$f_R(s_A, s_C)$
Turn Right	$f_R(s_A, s_C)$	$1 - 2 P_R - F_T(s_C, s_T)$	$f_T(s_A, s_T) + f_R(s_A, s_C)$
Turn Left	$f_R(s_A, s_C) + f_T(s_A, s_T)$	$1 - 2 P_R - F_T(s_C, s_T)$	$f_R(s_A, s_C)$
Lane Change Right	$f_R(s_A, s_C)$	$1 - 2 P_R - F_L(s_C, s_0, s_1)$	$f_L(s_A, s_0, s_1) + f_R(s_A, s_C)$
Lane Change Left	$f_R(s_A, s_C) + f_L(s_A, s_0, s_1)$	$1 - 2 P_R - F_L(s_C, s_0, s_1)$	$f_R(s_A, s_C)$

therefore model the indicator activation time as a single normal distribution with $\mu_L^{(t)} = -2.83\text{ s}$ and $\sigma_L^{(t)} = 0.61\text{ s}$. For each current velocity v_C , the distance is thus normal distributed with $\mu_L = v_C \mu_L^{(t)}$ and $\sigma_L = v_C \sigma_L^{(t)}$.

In contrast to the model for turn maneuvers, we now aim at estimating the indicator activation likelihood given that the driver intends to do a lane change within a whole lane segment s_0 to s_1 rather than at a single point. Assuming a uniform prior, the probability density function calculates to

$$f_L(s, s_0, s_1) = \frac{1}{s_1 - s_0} \int_{s_0}^{s_1} \tilde{f}_L(s, s_L) ds_L \quad (8)$$

with $\tilde{f}_L(s, s_L)$ and q_L analogue to (4) and (5). The overall probability of a lane change related indicator activation is estimated by $P_L = 1 - 0.30 - P_R = 0.68$ in accordance to the figures reported in [15]. For numerical evaluation, we use the error function (6) to transform (8) into

$$f_L(s, s_0, s_1) = \frac{P_L q_L}{2(s_1 - s_0)} \operatorname{erf} \left(\frac{s - s' - \mu_L}{\sqrt{2} \sigma_L} \right) \Big|_{s'=s_1}^{s'=s_0}. \quad (9)$$

After integration and some simplifications, we obtain

$$F_L(s, s_0, s_1) = \frac{P_L q_L}{2} \left[1 + \frac{\sqrt{2} \sigma_L}{s_1 - s_0} H \left(\frac{s - s' - \mu_L}{\sqrt{2} \sigma_L} \right) \Big|_{s'=s_1}^{s'=s_0} \right] \quad (10)$$

with the indefinite integral

$$H(x) = \int \operatorname{erf}(x) dx = x \operatorname{erf}(x) + \frac{1}{\sqrt{\pi}} e^{-x^2} + C. \quad (11)$$

IV. VELOCITY MODEL

Beside the indicator, another important feature for inferring the driver's intent has been shown to be his velocity profile during the intersection approach. The underlying idea is that drivers who are about to do a left or right turn will have to slow down before the turn whereas those who are about to go straight can maintain their speed. In the presence of preceding vehicles, however, the driver might be required to adjust his velocity regardless of his intents.

In [14], we proposed to account for car-following behavior by means of the Intelligent Driver Model [16]. In [17], the approach has been extended to make use of the curvature of the path lying ahead in order to calculate the expected turn related deceleration. Both defensive and sporty driving styles are captured by a total of nine different driver profiles. For each hypothesis $h \in H$ and driver profile $d \in D$, the expected acceleration $\hat{a}_{dk}^{(h)}$ of time step k is calculated from

the velocity of our own vehicle as well as the distance and relative velocity of the preceding vehicle, if present.

Assuming normal distributed deviations, the likelihood of the observed acceleration a_k with respect to hypothesis h can thus be obtained from

$$f_{A k}^{(h)}(a_k) = \frac{P_{A0}}{20 \text{ m/s}^2} + \sum_d P(d) f_{A d k}^{(h)}(a_k) \quad (12)$$

with

$$f_{A d k}^{(h)}(a_k) = \frac{1 - P_{A0}}{\sqrt{2\pi} \sigma_A} \exp \left[-\frac{1}{2} \left(\frac{a_k - \hat{a}_{dk}^{(h)}}{\sigma_A} \right)^2 \right] \quad (13)$$

and $P(d)$ the prior of driver profile d . The standard deviation σ_A is a measure for the remaining variance within each driver profile, whereas P_{A0} defines the probability that the observed acceleration is not at all to be explained by our model. In this case, we assume a to be uniformly distributed within -10 m/s^2 and 10 m/s^2 . For our experiments, we assumed $\sigma_A = 1.2 \text{ m/s}^2$ and $P_{A0} = 0.01$.

In order to obtain a smooth estimate of the driver's intent, it makes sense to include both current and past observations in the overall feature likelihood. The challenge here lies in that the individual observations are likely to be strongly correlated with each other, so by assuming conditional independence we would run at risk to include the same piece of information several times and thereby overweight our feature. Instead, we chose to average the log-likelihoods of the individual observations over $N_{\text{Vel}} = 14$ time steps of length 100 ms so as to obtain the likelihood of a single virtual observation:

$$f_{\text{Vel}}^{(h)} = \exp \left[\frac{1}{N_{\text{Vel}}} \sum_{i=0}^{N_{\text{Vel}}-1} \log \left(f_{A k-i}^{(h)}(a_{k-i}) \right) \right]. \quad (14)$$

V. DRIVER GAZE MODEL

Visual input is known to be the driver's main source of information. By observing his gaze direction, we can obtain clues about the driver's intents. We distinguish two main causes for the driver's search for visual information: The need to monitor the car's heading with respect to the planned path along the street, and the need to make sure that there are no conflicts with other traffic participants.

The latter results in quick glances in the direction of possible threats. In our previous work [18], we used such glances to predict the driver's wish to commit lane-change maneuvers on highways and right-turn maneuvers at urban

intersections. In order to predict the probability that such maneuvers are actually carried out, however, the driver's situation awareness needs to be taken into account. This again depends on the history of the driver's gaze direction, so the corresponding probabilistic model is rather complicated and quite out of the scope of this paper.

Instead, we aim at inferring the driver's intent based on the portion of his gaze behavior that is due to his need to monitor the car's heading with respect to the planned path along the street. Participants in professional driver trainings are often told not to look towards the obstacle they want to evade but always in the direction of intended travel, as the car is most likely to go where the driver is looking. Conversely, a literature survey reveals that drivers are believed to look either at a curve's inner tangent point [19] or somewhere along the path that they are about to follow [20].

In practice, the head pose of the driver can be captured much more reliably than the actual gaze direction. We therefore propose to use the head heading angle for driver intent inference. In order to approximate the driver's gaze behavior observed in [19], we define an expected gaze point that lies on the path of intended travel defined by the corresponding hypothesis h . The relative distance $\Delta s = a_0 + a_1 v$ at which the expected gaze point is located along this path may depend on the car's current velocity v . In addition, we allow for a constant lateral displacement $\Delta v = b_0$. A plausible set of expected gaze points is visualized in Figure 7.

Given the heading angle $\hat{\varphi}_k^{(h)}$ of the expected gaze point of hypothesis h and time step k relative to the car's coordinate system, we define the likelihood function for the observed head heading deviation $\Delta\varphi_k^{(h)} = \varphi_k - \hat{\varphi}_k^{(h)}$ by

$$f_{\Phi}(\Delta\varphi_k^{(h)}) = \frac{1 - P_{\Phi 0}}{\sqrt{2\pi}\sigma_{\Phi}} \exp\left[-\frac{1}{2}\left(\frac{\Delta\varphi_k^{(h)}}{\sigma_{\Phi}}\right)^2\right] + \frac{P_{\Phi 0}}{2\pi}, \quad (15)$$

where $P_{\Phi 0}$ is the probability that the observed head heading angle is not to be explained by our expected gaze point model but by the need to check for conflicting traffic participants, for instance. In this case, we assume a uniform distribution over all possible head heading angles.

The gaze point parameters a_0, a_1 and b_0 as well as the likelihood parameters σ_{Φ} and $P_{\Phi 0}$ have been obtained by maximizing the overall likelihood of observed head heading angles of more than 200 intersection crossings conducted by 6 subjects on 3 intersections. The results are given in Table II. Figure 8 shows the observed head heading angles for going straight and turning right as well as the expected head heading angles corresponding to our model. The deviation likelihood function $f_{\Phi}(\Delta\varphi_k^{(h)})$ is visualized in Figure 9.

For a smoother estimate of the driver's intent, we average the log-likelihoods over $N_{\text{Gaze}} = 10$ time steps:

$$f_{\text{Gaze}}^{(h)} = \exp\left[\frac{1}{N_{\text{Gaze}}}\sum_{i=0}^{N_{\text{Gaze}}-1}\log\left(f_{\Phi}(\varphi_{k-i} - \hat{\varphi}_{k-i}^{(h)})\right)\right]. \quad (16)$$

TABLE II
MAXIMUM LIKELIHOOD ESTIMATE OF GAZE PARAMETERS

Parameter	Value	Unit
Look-ahead distance a_0	6.07	m
Look-ahead time a_1	1.05	s
Lateral deviation b_0	-0.81	m
Standard deviation σ_{Φ}	0.13	—
Wrong model probability $P_{\Phi 0}$	0.11	—

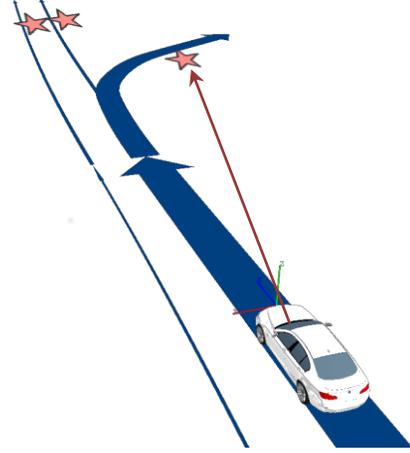


Fig. 7. Expected gaze points for right-turn situation.

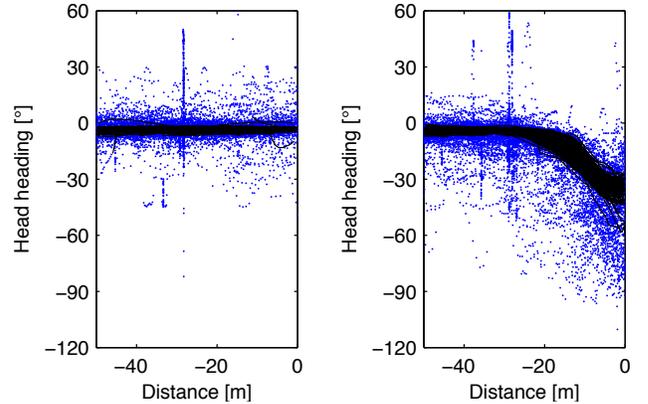


Fig. 8. Actual driver head heading angle (blue dots) and expected head heading according to our model (black lines) for straight intersection crossings (left) and right-turn maneuvers (right).

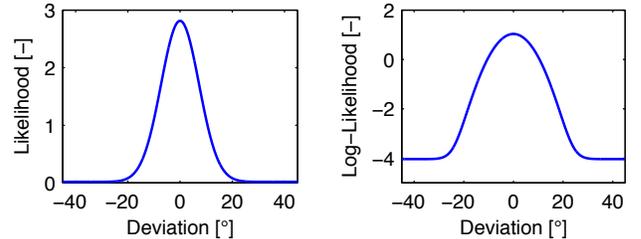


Fig. 9. Likelihood $f_{\Phi}(\Delta\varphi_k^{(h)})$ of the head heading deviation.

VI. EXPERIMENTAL RESULTS

After describing the parametric models of each of our features, we are now going to evaluate them based on how well they perform by themselves and how well they can be combined to improve the overall classification result. To this end, the indicator, velocity and gaze direction feature must all be able to contribute to the respective hypotheses, and good classification performance must not be possible based on the vehicle’s localization result alone. Also, to allow for quantitative results, it must be feasible to collect a large amount of relevant maneuvers in real traffic. For these reasons as well as for the ability to compare the results with those of preceding work, we chose to predict simple right-turn maneuvers for our evaluation even though the architecture described in Section II allows for more complex combinations of subsequent maneuvers in principle.

The evaluation has been carried out on a dataset containing 15 hours of driving data collected by 12 subjects. Each subject was to drive along a predefined route that repeatedly crossed 5 intersections that have not been used in the process of model parametrization. At these intersections, a total of 155 right turns and 244 straight intersection crossings have been collected. We manually defined a conflict point at the pedestrian crossing directly after each right turn. The remaining time to reach that conflict point (*TTC*) given that the driver was to turn right has been continuously estimated according to the method presented in [14]. We used a high-precision GPS/INS platform for our evaluation, so self-localization errors can be neglected.

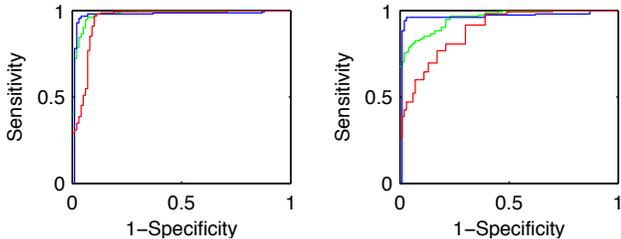


Fig. 10. Single-feature right-turn prediction results at *TTC* = 2 s (left) and *TTC* = 3 s (right) for the indicator feature (best), the velocity feature (medium) and the driver gaze direction feature (worst).

The classification performance of the individual features is shown in Figure 10. The indicator feature turns out to be superior, though this might be due to the fact that drivers tend to use the indicator rather diligently when taking part in a user study. The velocity feature is comparably strong as well, with misclassifications happening only at very low speeds, such as in situations with preceding vehicles or directly after a stop at red light. A drawback of our current head tracking system is that it frequently loses track of the head direction at fast movements such that occur when the driver is looking over his shoulder to check for bicycles. Also, the driver gaze feature does not yet model such shoulder and mirror glances. Both might explain the poor classification performance at high specificity values at *TTC* = 2 s. Still, the classification performance is quite good, but less so for *TTC* = 3 s.

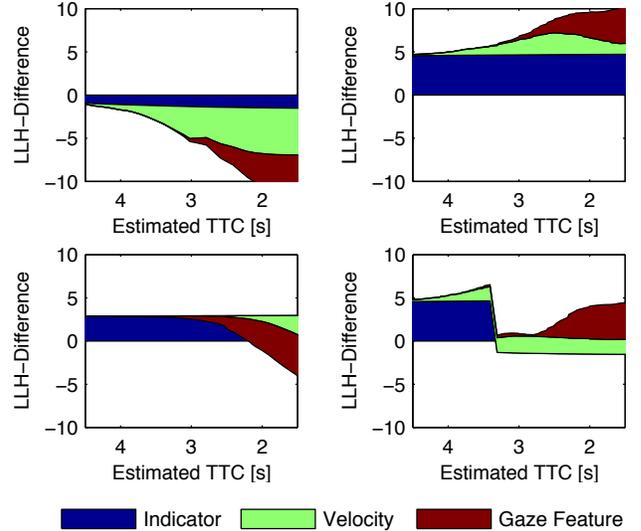


Fig. 11. Contribution of each of the three features’ log-likelihood difference $\Delta LLH_i(h_1, h_2)$ to the overall classification result. Starting from zero, the indicator feature’s contribution is drawn first, then the velocity and gaze feature’s contributions are stacked on top. Upper left: Normal straight intersection crossing. Upper right: Normal right turn maneuver. Lower Left: Straight intersection crossing with indicator activated to the right. Lower right: Right turn maneuver with indicator turned off at *TTC* \approx 3.3 s.

According to (1), the posterior probability based on all three features is proportional to the product of their individual likelihoods. Therefore, we have

$$\log \frac{P(h_1|O)}{P(h_2|O)} = \sum_i \Delta LLH_i(h_1, h_2) + C \quad (17)$$

with

$$\Delta LLH_i(h_1, h_2) = \log P(O_i|h_1) - \log P(O_i|h_2) \quad (18)$$

and $C = 0$ for any two hypotheses h_1 and h_2 as long as their prior probabilities are the same. For the case of $h_1 = \text{Turn}$ and $h_2 = \text{Straight}$, the contribution of the three feature’s log-likelihood differences to the classification result for some particularly interesting situations is shown in Figure 11.

The examples show that the features are well balanced, which means that any two features could overrule the third as long as they are confident enough. This works well in the lower right situation, where the indicator is accidentally deactivated just before a right turn maneuver. It did not work for the lower left scenario though, apparently because it was a car-following situation or a stop at red light so the velocity feature was not confident at all. The examples confirm that our driver gaze feature works best for *TTC* < 2.5 s.

A quantitative evaluation of different feature combinations is provided by Table III, assuming that a right turn is predicted iff its probability is greater than 0.5. We distinguish between three different cases:

- 1) The indicator signal is not available (-), which is the case when we try to reason about other traffic participants that we observe with our laser scanner or radar sensors. In this case, both the velocity and gaze feature as well as their combination show excellent sensitivity but only moderate specificity at *TTC* = 3 s.

TABLE III
CLASSIFICATION RESULTS AT $TTC = 3$ s.

Ind	Vel	Gaze	Sensitivity	Specificity
-	-	-	0.56	0.69
-	-	+	0.92	0.71
-	+	-	0.97	0.70
-	+	+	0.98	0.74
+	-	-	0.96	0.98
+	-	+	0.96	0.98
+	+	-	0.97	0.98
+	+	+	0.97	0.97
0	-	-	0.15	1.00
0	-	+	0.19	1.00
0	+	-	0.67	1.00
0	+	+	0.75	0.99

- The indicator signal is available (+), which is the case when we reason about our own driver's intent or about that of a vehicle that communicates its indicator status via Car2X-communication. As long as the indicator is near to always set, its feature can hardly be improved.
- The indicator signal is available but not activated (0), which, in practice, has been shown to be the case for 20% of all turn maneuvers [15]. For this evaluation, we post-processed the dataset to have the indicator signal always turned off. The results show that the specificity is always near to 100% now, as any feature or combination thereof must be quite confident before it can overrule the indicator feature. It turns out that 15% of all turn maneuvers can be predicted based on the lane assignment itself, which is the case at some intersections if the car is moving very slowly, e.g. because of preceding vehicles. While the gaze feature by itself is seldom confident enough to overrule the indicator, it increases the ratio of detected right-turn maneuvers from 2/3 to 3/4 if combined with the velocity feature.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have introduced a novel architecture for generic driver intent inference that allows to reason about arbitrary combinations of subsequent maneuvers. Starting with the velocity feature of our previous work [17], we have provided new parametric models for the indicator and driver gaze direction feature. Both have been parametrized from the results of extensive user studies. Finally, we evaluated the right-turn classification performance of our approach based on a separate study. By combining the individual feature likelihoods by a naive Bayes classifier, we were able to investigate the contribution of each feature to the overall classification result separately. In particular, our approach has been shown to work well even when faced with implausible observations of the indicator feature.

Future work will be concerned with interaction between traffic participants and the implementation of right-of-way rules. Aiming at risk assessment, we will need to take the driver's awareness of the situation into account. For

this purpose, we will extend our approach to do exact or approximate inference in a more general class of probabilistic graphical models.

ACKNOWLEDGMENTS

This work was funded in part by the Federal Ministry of Economics and Technology of the Federal Republic of Germany under grant no. 19 S 9022.

REFERENCES

- R. Wertheimer and F. Klanner, "Cooperative Perception to Promote Driver Assistance and Preventive Safety," in *8th International Workshop on Intelligent Transportation*, no. ii, 2011.
- H. Stübing, M. Bechler, D. Heussner, T. May, I. Radusch, H. Rechner, and P. Vogel, "sim TD: a car-to-x system architecture for field operational tests," *IEEE Communications*, no. 5, pp. 148–154, 2010.
- G. S. Aoude, V. R. Desaraju, L. H. Stephens, and J. P. How, "Behavior Classification Algorithms at Intersections and Validation using Naturalistic Data," in *2011 IEEE Intelligent Vehicles Symposium*, 2011, pp. 601–606.
- B. Morris, A. Doshi, and M. Trivedi, "Lane Change Intent Prediction for Driver Assistance : On-Road Design and Evaluation," in *Intelligent Vehicles Symposium, 2011 IEEE*, 2011, pp. 895–901.
- Q. Tran and J. Firl, "A probabilistic discriminative approach for situation recognition in traffic scenarios," in *2012 IEEE Intelligent Vehicles Symposium*, 2012, pp. 147–152.
- G. Ortiz, J. Fritsch, F. Kummert, and A. Geppert, "Behavior prediction at multiple time-scales in inner-city scenarios," in *2011 IEEE Intelligent Vehicles Symposium*, 2011, pp. 1066–1071.
- J. Firl, H. Stübing, S. Huss, and C. Stiller, "Predictive maneuver evaluation for enhancement of car-to-x mobility data," in *IEEE Intelligent Vehicles Symposium*, Spain, June 2012, pp. 558–564.
- T. Gindele, S. Brechtel, and R. Dillmann, "A Probabilistic Model for Estimating Driver Behaviors and Vehicle Trajectories in Traffic Environments," in *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on*, 2013, pp. 1066–1071.
- S. Levèvre, C. Laugier, and J. Ibanez-Guzmán, "Risk assessment at road intersections: Comparing intention and expectation," in *2012 IEEE Intelligent Vehicles Symposium*, 2012, pp. 165 – 171.
- A. Doshi and M. M. Trivedi, "Tactical driver behavior prediction and intent inference: A review," in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Oct. 2011, pp. 1892–1897.
- H. Berndt and K. Dietmayer, "Driver intention inference with vehicle onboard sensors," in *Vehicular Electronics and Safety (ICVES), 2009 IEEE International Conference on*. IEEE, 2009, pp. 102–107.
- J. Zhang and B. Roessler, "Situation analysis and adaptive risk assessment for intersection safety systems in advanced assisted driving," in *Autonome Mobile Systeme 2009*, ser. Informatik aktuell, R. Dillmann, J. Beyerer, C. Stiller, J. Zöllner, and T. Gindele, Eds. Springer Berlin Heidelberg, 2009, pp. 249–258.
- S. Lefèvre and J. Ibanez-Guzmán, "Context-based Estimation of Driver Intent at Road Intersections," *Intelligence in Vehicles*, 2011.
- M. Liebner, M. Baumann, F. Klanner, and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model," in *IEEE Intelligent Vehicles Symposium*, Alcalá de Henares, Spain, June 2012, pp. 1162–1167, (best paper award).
- Auto Club Europa, "Reviere der Blinkmuffel," 2008. [Online]. Available: <http://www.ace-online.de/grafiken>
- M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E* 62, 2000.
- M. Liebner, F. Klanner, M. Baumann, C. Ruhhammer, and C. Stiller, "Velocity-based driver intent inference at urban intersections in the presence of preceding vehicles," *IEEE Intelligent Transportation Systems Magazine*, vol. 5, no. 2, pp. 10–21, May 2013.
- M. Liebner, F. Klanner, and C. Stiller, "Der Fahrer im Mittelpunkt – Eye-Tracking als Schlüssel zum mitdenkenden Fahrzeug?" in *8. Workshop Fahrerassistenzsysteme Waling*, 2012, pp. 87–96.
- M. F. Land and D. N. Lee, "Where we look when we steer," *Nature*, vol. 369, pp. 742–744, 1994.
- J. P. Wann and D. K. Swapp, "Where do we look when we steer and does it matter?" *Journal of Vision*, vol. 1, no. 3, p. 185, 2001.