Towards a Multi-Hypothesis Road Representation for Automated Driving

Frank Dierkes* and Markus Maurer
Institute of Control Engineering
Technische Universität Braunschweig
{dierkes, maurer}@ifr.ing.tu-bs.de

Marvin Raaijmakers*, Max Theo Schmidt*,
Mohamed Essayed Bouzouraa and Ulrich Hofmann
Pre-development Automated Driving Systems
AUDI AG
marvin@marvin-r.com, {max.schmidt, essayed.bouzouraa, ulrich.hofmann}@audi.de

***Abstract*—The process of perception inevitably involves uncertainty. In the field of automated driving, however, uncertainty about the roadway—and especially about the understanding of the roadway—is yet mostly neglected. In this paper it is argued that this uncertainty should be represented and considered in the process of behavior generation. A road representation capable of representing multiple hypotheses about the roadway is presented. The representation language allows to express where and how hypotheses differ and to infer this information in an efficient way. The focus of this paper is on the qualitative representation of uncertainty.**

I. Motivation

Automated road vehicles have to cope with uncertain information about their environment in the process of behavior generation. Recent work in the field of automated driving that addresses this issue focuses on the uncertainty about the future behavior of other road users, for example to reason whether it is possible and beneficial to perform a lane change maneuver [1] or whether it is possible to perform a turn maneuver [2].

An uncertainty that is mostly neglected is the uncertainty about the roadway itself, especially the uncertainty about the understanding of the roadway. To understand the roadway means to understand what portions of the roadway are considered to be lanes, but also to understand the meaning of the other portions of the roadway. Even if the true physical structure of the roadway is known, uncertainty may arise in the form of ambiguity, for example in construction zones or if old road markings were not removed thoroughly. However, the main source of uncertainty lies in the process of perception, namely false-positive and false-negative detections of lane markings, curb lines, et cetera, as well as false classifications and inappropriate models.

Demonstrations of automated driving so far either relied on highly-accurate a priori maps [3]–[6], which provide the necessary information about the roadway, or took place on controlled-access highways [7], where the perception and the understanding of the roadway is less problematic than in other domains.

*The authors F. Dierkes, M. Raaijmakers and M. T. Schmidt contributed equally to the development of the presented multi-hypothesis road representation.

The understanding of the roadway and its lanes is fundamental for automated road vehicles, since it indicates where the vehicle shall drive. Epistemologically, this question is more difficult to answer than the question where the vehicle can drive, that is where the vehicle does not collide with any obstacles. A hypothesis that a certain path is free of collisions is comparatively easy to reject, for example by using sensors like radar or lidar. A lane, however, is less a physical object, but more a mental model, resultant from interpreting lane markings and other elements of the roadway. The way this shall be done depends on the traffic rules and on the context of the traffic scene.

False detections and classifications lead to multiple, usually conflicting hypotheses about the understanding of the roadway. It is problematic to reject hypotheses about lanes, because there is typically no solid evidence against a hypothesis. One approach might be to test whether a hypothesis fits with the construction guidelines or more generally with the expectation of what a roadway looks like. But a roadway is not necessarily constructed accordingly to the guidelines. The trajectories of other road users might give a hint, but then again other road users do not necessarily follow lanes nor do they necessarily behave compliant to the traffic rules. Ground classification is a way to exclude hypotheses outside of the roadway, but this process itself involves uncertainty. In general, a set of hypotheses about the roadway remains. Its cardinality depends on the quality of the roadway perception.

To understand the roadway and its lanes and to cope with the inevitable uncertainty about it is still one of the major challenges on the way towards autonomous driving. Current approaches focus mostly solely on the detection of lanes [8]. Behavior planning in the approaches mentioned above is based on an unambiguous understanding of the roadway. This assumes, that solely the perception system copes with uncertainty about the roadway, including possible ambiguities that arise from the true state of the roadway. Even though an increasement of the perception capabilities of automated road vehicles may reduce the uncertainty about the roadway significantly, some uncertainty will remain. In general, requiring an unambiguous representation of the roadway for behavior planning forces the perception system to make decisions that may influence the behavior of the vehicle in a crucial way.

Thus, the existing uncertainty about the roadway should be represented and this information should be considered in the process of behavior planning. This requires a representation

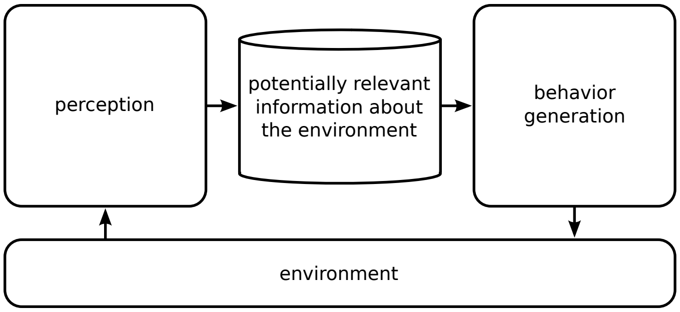

Figure 1. Simplified architecture of the system driving the automated vehicle. The environment perception is intended to be independent of a particular automated driving function. All function-specific interpretation of the environment should be done in the process of behavior generation. The road representation is part of the data interface between both subsystems.

language that is able to express this uncertainty. Although a quantification of the uncertainty is desirable, for some aspects of behavior generation already the qualitative representation of uncertainty is useful. An example are beginning off-ramps, which are easily falsely recognized as a widening lane, for instance because the lane marking that separates both lanes has not accumulated enough evidence at first. Instead of deciding for one hypothesis—the widening lane or the lane bounded by a marking with little evidence—the perception system can represent both hypotheses. The behavior generation then might decide to follow the lane with the constant width, since it is also a portion of the widening lane.

The focus of this paper is on the qualitative representation of uncertainty about the understanding of the roadway. A multi-hypothesis road representation is presented that allows to infer efficiently where and how hypotheses about a roadway differ. It is not discussed, how these hypotheses are constructed in the perception process. Also out of scope are junctions, the paper focuses on the roadway in between junctions.

II. Overview

The multi-hypothesis road representation was designed under consideration of several use cases. The primary one is its use as a part of the data interface between the environment perception and the behavior generation of an automated vehicle (see figure 1). A crucial distinction made here between perception and behavior generation is that the perception system is supposed to generate the environment representation for diverse automated driving functions. No interpretations or abstractions specific to a particular function should be made in the process of perception, it is intended to be function-independent.

Hence, as a first major requirement, the road representation must allow a precise description of the physical structure of the roadway. For example, a lane shall be represented by its boundaries, not by a path that follows the lane, which would be a premature abstraction.

Obviously, the second major requirement is the representation of the existing uncertainty about the roadway, especially the representation of multiple, conflicting hypotheses about the lanes of the roadway. In contrast to [9], the multi-hypothesis representation is not only meant for reasoning about the most probable hypothesis. Instead, it is the basis for behavior planning. Thus, the representation must allow planning algorithms to work efficiently on it. This means for example to infer, which hypotheses exist about a certain region of the roadway and what the difference between these hypotheses is.

This paragraph gives a brief overview of the remaining paper. After the discussion of the related work in section III, section IV introduces the applied model of the roadway. Section V then discusses the representation of multiple hypotheses about the roadway based on this model. Besides a general outlook of the future work, section VI discusses briefly possible approaches to the quantification of uncertainty.

III. Related work

Töpfer *et al.* [9] presented a representation of multiple hypotheses about a roadway as an integral part of a scene understanding approach. It is basically a joint probability distribution over the roadway modeled as an undirected graphical model. Patches, which represent segments of lanes, are defined by their position, orientation, width and length. Lanes are modeled as a set of patches and roads are modeled as a set of lanes. Patches, lanes and roads are nodes in the graphical model. Edges model spatial constraints and the potential between two nodes. The approach aims at infering the most probable hypothesis about the roadway, but not at representing multiple hypotheses to behavior planning. Where and how hypotheses differ is not directly inferable. The spatial topology of lanes is not represented either.

Knaup and Homeier [10] presented a graph-based environment representation. It is intended to be independent of a specific automated driving/driver assistance function. The representation of the road infrastructure is part of the environment representation and basis for the representation of other elements of the environment. For example, other road users are associated to lanes. Uncertainty is considered with respect to the association of road users to lanes, but not with respect to the road infrastructure itself. The representation is not able to represent multiple hypotheses about the roadway. The road model only consists of roadways and lanes.

A similar graph-based context representation was proposed by Ulbrich *et al.* [11]. Lane segments are connected in longitudinal direction via waypoints. In lateral direction, neighboring lane segments are linked via a common boundary. Again, neither uncertainty about the roadway nor the representation of portions of the roadway other than lanes is discussed.

The approach presented by Bender *et al.* [12] aims at representing highly-accurate a priori lane-level maps. Lane segments, which are called lanelets, are defined by two lateral boundaries and are of arbitrary length. The objective is not to represent the physical structure of the road (network), but to represent the logical topology of the network of lanelets. A bidirectional lanelet, for example, is thus represented twice.

IV. Roadway model

The scope of this paper is the roadway between two junctions. While the precise criteria for the decomposition of the road network into roadways and junctions of roadways are out of scope, the crucial distinction made between junctions and

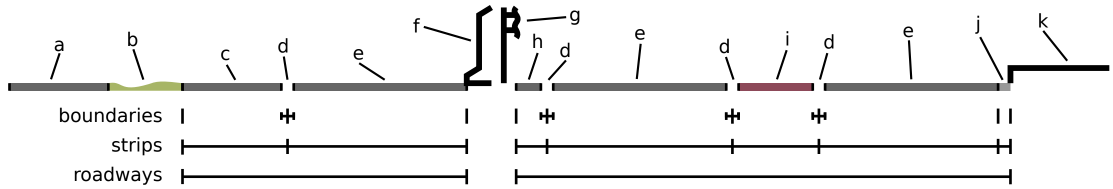

Figure 2. Lateral structure of the roadway model. The roadway is composed of strips. A strip is defined by its lateral boundaries and its type. There are no gaps between strips and no overlaps of strips. A road may have more than one roadway. Here, an example of a rather rare cross-section of a road is pictured. It consists of (a) a side path, (b) a green strip, (c) a shoulder, (d) painted lane markings, (e) travel lanes, (f) a barrier, (g) a guard rail, (h) a separating strip, (i) a bicycle lane, (j) a gutter and (k) a sidewalk.

roadways is, that a roadway is defined to have no overlapping and no intersecting lines of traffic. This section discusses the applied model to represent a single hypothesis about a roadway.

A. Lateral structure

The main objective is to model and to represent that portion of the environmet, where the vehicle is supposed to drive automatedly, that is the space that is potentially considered a legitimate (goal) position by an automated driving function.

This space corresponds to a great extent with the various definitions of the term roadway (ger. *Fahrbahn*). The definitions of the term differ slightly. A roadway sometimes solely comprise the traveled way [13], that is the traffic lanes, sometimes it also comprises contiguous bicycle lanes and the shoulder [14], [15] and sometimes it even comprises sidewalks for example [16]. The flexibility of the term fits well with its use in the presented representation, because an extended functionality might require an extension of the modeled space, for example if the vehicle should be able to use a green strip or a sidewalk contiguous to the pavement to give way to another vehicle.

In the following discussion, a *roadway* is defined as that portion of a road, that is bounded by a shoulder line (edge of the pavement), a curb line or a guard rail or barrier (cf. figure 2). An exception is made for traffic islands, which are allowed to be modeled as a part of a roadway. A road may have more than one roadway, for example in case of a divided highway.

For behavior generation, the most relevant cross-sectional elements of a roadway are its *lanes*. A lane is a portion of the roadway for the movement of a single line of vehicles in one direction [13]–[16]. Understanding the lanes of the roadway requires a step of abstraction. Lane markings indicate, which portions of a roadway shall be interpreted as lanes. They force road users to conceive the same representation/understanding of the roadway. In some domains, for example highways, lane markings are usually present, in other domains, for example residental streets, they are very uncommon. However, also in these domains, it might be useful to apply the model of lanes in form of *virtual lanes* for predicting the behavior of other road users and for planning one's own behavior. It might also be useful for behaving compliant to the traffic rules, for example to decide which vehicle has the right of way or when to activate the turn indicator.

In the following, a lane is defined as a portion of the roadway, that is defined by two lateral boundaries and that is meant for vehicular use. If a lane marking is present on each side, the lane is defined by these markings. If no lane marking is present next to a boundary of the roadway or next to a traffic island, the lane may extent to a curb line, a shoulder line, a guard rail or a barrier. A lane that is separated from neighboring lanes by lane markings is called a marked lane. It is also possible to set virtual boundaries to handle cases where no lane marking is present (or where a lane marking was not detected). A lane that has at least one virtual boundary is called a virtual lane. (Going with the definition that the perception system shall not make any function-specific decisions or abstractions, the creation of virtual lanes may in some cases still be considered a part of the perception process, but in other cases already a part of the process of behavior generation.)

The autonomous system needs not only to understand the lanes, but also the other portions of the roadway. For example, the vehicle might have to cross a bicycle lane to perform a lane change (cf. figure 2). Driving on or crossing other portions of the roadway might require to consider other aspects in the process of behavior generation, for example specific traffic rules. It is mandatory to describe the roadway without any gaps. Therefore, the concept of *strips* is introduced. A strip is a portion of the roadway that is defined by two lateral boundaries and that is of a certain type. It may be a (traffic) lane, a bicycle lane, a shoulder, a traffic island, et cetera. Each portion of the roadway is modeled as a strip, thus there are no gaps in the representation of the roadway.

The above definitions of the cross-sectional elements all base on the description of boundaries. A boundary itself may be extended in lateral direction, for example in case of a lane marking. Other kinds of boundaries are not extended, for example the edge of the pavement. The boundaries carry the geometric information about the roadway.

Common queries from behavior planning are to ask which strip is adjacent to another strip or which strips lie in between two other strips. Thus, it is reasonable to explicitly represent the spatial topology of the strips of a roadway. Two strips are defined to be adjacent, if they share a common boundary. Since no gaps (and no overlaps) are allowed, the complete spatial topology of the lateral structure of the roadway is described by associating adjacent strips via the common boundary or equivalently by ordering the strips from one side of the roadway to the other.

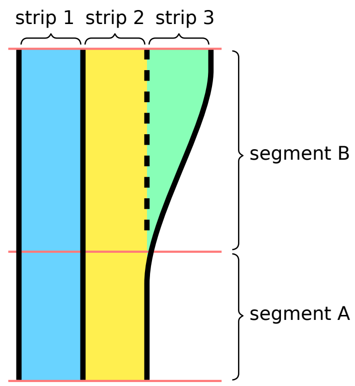

Figure 3. Longitudinal structure of the roadway model. The roadway is divided into two segments according to the segmentation criterion (beginning of strip 3). Strip segments 1A and 2A are connected longitudinally to strip segments 1B and 2B, respectively. Strip segment 3B is only connected laterally to strip segment 2B.

B. Longitudinal structure

In the following, the model of the lateral structure of the roadway is extended in longitudinal direction. The lateral structure is not constant over the length of the roadway. New strips may begin and existing strips may end at some point. For this reason, the roadway is segmented in longitudinal direction (see figure 3). The criterion for a new segment is the beginning or the ending of a strip. Segmenting the roadway and the elements it consists of according to this criterion yields a constant spatial topology in each segment.

The roadway is modeled as a sequence of *roadway segments*. Each roadway segment consists of a sequence of *strip segments*, ordered from one side of the roadway to the other and thereby representing the spatial topology in lateral direction. Each strip segment is defined by its two *boundary segments*. The longitudinal connection of the strip segments of consecutive roadway segments is straightforward. Strip segments are connected, if they share a common portion of the roadway cross-section. In figure 3 strip segments 1A and 2A are connected longitudinally to strip segments 1B and 2B respectively, but strip segment 3B has no predecessor in segment A. A strip segment may also have more than one predecessor or successor. Lanes, for example, often grow wider and then split in front of intersections. For the sake of simplicity, this case it not separately shown in the following discussion of the multi-hypothesis representation, since all mechanisms apply analogously.

The above criterion only yields a representation of the spatial relations between strips. Thinking of lanes, for behavior planning it is not only of interest where a vehicle *can* change to another lane, but also where it is *allowed to* do so. This information is usually encoded in the lane marking. To represent this information, the segmentation criterion can be extended. A new segment then also begins, whenever the type of a lane marking (e. g., solid, broken) changes. In this case, also the type of all lane markings is constant inside a segment. An alternative representation for the same information is as a profile over the length of the boundary, attached as an attribute to the boundary. The query, whether it is allowed to cross the boundary in a certain segment, then requires some more computation. Which form of representation is considered to be more efficient, is a trade-off between the execution time of queries and the number of segments, which may increase significantly in the multi-hypothesis case. In the end, it also depends very much on the actual implementation of behavior planning. In the following discussion of the multi-hypothesis case, the first segmentation criterion is applied, since focusing on one criterion simplifies the discussion. However, there is no loss of generality, the discussed mechanisms work for any segmentation criterion.

V. Multi-hypothesis representation

To support the explanation of the multi-hypothesis road representation in this section we make use of the example illustrated by figure 4. As illustrated on the left side of figure 4, in this example there exists a road with road markings that are partially detected (the coverage range of the detection may be limited by the sensors or occlusions) by the perception system as well as false road marking detections. Since the perception system cannot distinguish between true and false detections of road markings it generates multiple hypothesis about the roadway. Such hypotheses can be generated based on, for example, the quality of the road marking detections, knowledge about road models based on road construction guidelines, or a priori map knowledge. Since the criteria and methods for hypothesis generation are not within the scope of this paper, we present some hypotheses that are generated in our example without discussing the applied criteria in detail or trying to present a complete set of possible hypotheses. The main point of the example is to show that there can be multiple hypotheses about the roadway and how these hypotheses are represented in our representation.

The goal of the multi-hypothesis representation is to represent the set of different hypotheses about the roadway such that it is easily interpretable for its user and is efficient in terms of memory usage. The presented representation is easily interpretable in the sense that it is easy for the user to find out where and how the hypotheses about a roadway differ. There may be some regions (i.e. ranges in the longitudinal direction of the road) on the road where multiple hypotheses about a roadway look the same. The presented representation is memory efficient in the sense that it describes such regions of roadway hypotheses only once instead of describing it multiple times for each roadway hypothesis separately. This kind of memory efficiency and interpretability is achieved by partitioning the roadway into *roadway subsegments*. A roadway subsegment is defined as a range in longitudinal direction of the road with the following properties:

1) There are no two hypotheses about the roadway such that there is one place within the subsegment where the hypotheses differ and another place within the subsegment where they are the same.
2) For each roadway hypothesis the neighbor relations between strips are constant within the subsegment.
3) There exists no range in longitudinal direction that includes the subsegment and is longer than the subsegment.

Property 1 assures the memory efficiency and interpretability requirement described above. Property 2 assures a convenient

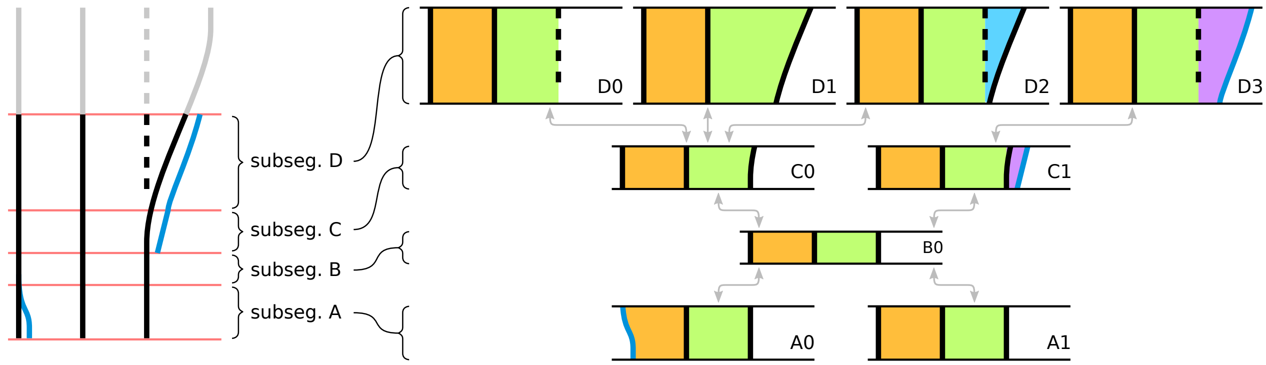

Figure 4. Left side: Road markings as they appear in the real world are visualized by the black/gray lines. The black part of these markings indicate the part that is observed by the perception system and are considered as road marking candidates. The blue lines visualize further candidates for road markings that were detected by the perception system. The red lines indicate the borders of the roadway subsegments in the constructed road representation instance. Right side: For each of the roadway subsegments the roadway subsegment hypotheses are visualized. The arrows indicate the longitudinal connections between these road subsegment hypotheses. If two strip subsegment hypotheses of neighboring subsegments have the same color, then that means that there exists a strip subsegment hypothesis connection between them.

interpretation of strip neighbor relations just as roadway segments do in the single-hypothesis case. Property 3 avoids that roadway subsegments become unnecessarily small and that therefore the number of subsegments is unnecessarily large. As a result of this definition of roadway subsegments, the roadway is split into subsegments at those places where hypotheses about the roadway start or end to differ, or where there is at least one roadway hypothesis for which the neighbor relations of its strips change.

The data structure of our multi-hypothesis representation is illustrated by an UML class diagram in figure 5. Here a roadway is represented by the `Roadway` class and its partitioning into roadway subsegments is represented by a list of `RoadwaySubseg` objects. Each roadway subsegment has a set of *roadway subsegment hypotheses* (represented by class `RoadwaySubsegHyp`) that each describe a hypothesis about the roadway *within* that roadway subsegment.

On the right side of figure 4 we see for our example for each roadway subsegment, the possible roadway subsegment hypotheses. In roadway subsegment A there is an uncertainty about the left border of the left lane because there are two possible candidates for this boundary. Therefore there are two roadway subsegment hypotheses for A: hypothesis A0 and hypothesis A1. Within roadway subsegment B all hypotheses about the roadway are the same and thus there is only one hypothesis B0. Because of this change in uncertainty about the roadway when going from subsegment A to B, property 1 enforces us to split up the roadway into subsegments A and B. Since the perception system does not know whether the falsely detected marking on the right is correct or not, there are two roadway subsegment hypotheses (C0 and C1) for roadway subsegment C. Apart from property 1, a split up between subsegments B and C is also enforced by property 2 because there are roadway hypotheses (namely all roadway hypotheses that contain roadway subsegment hypothesis C1) that get a new lane in subsegment C. When we take a look at the hypothesis for roadway subsegment D we see that there is doubt about the borders of the second and third lane and the existence of a third lane.

The description of a roadway subsegment hypothesis in a `RoadwaySubsegHyp` consists of descriptions of the hypothesis' strips, called *strip subsegment hypotheses*, within the corresponding roadway subsegment and possible connections to preceding and succeeding roadway subsegment hypotheses. The strip subsegment hypotheses of a roadway subsegment hypothesis are represented by an ordered list of references to `StripSubsegHyp` objects. The order of this list is used to encode the neighbor relations between the strips, that is two strips are neighbors if and only if they are neighbors in the list. The list is a list of references since one `StripSubsegHyp` can belong to multiple hypotheses of the same roadway subsegment. This is done because this makes it possible for the user to find out whether different roadway subsegment hypotheses of the same roadway subsegment have strips in common. For example, for roadway subsegment A there are 3 different strip subsegment hypotheses (see also figure 6): roadway subsegment hypotheses A0 and A1 share the strip subsegment hypothesis for the green colored strip, but have different hypotheses for the orange colored strip.

The class `RoadwaySubsegHypConn` models possible connections in longitudinal direction between two roadway subsegment hypotheses. That is, there exists a connection between two road subsegment hypotheses if and only if they belong to consecutive subsegments and there exists a roadway hypothesis that contains both road subsegment hypotheses. For example, there is a connection between roadway subsegment hypotheses C0 and D2 because there exists a hypothesis about the roadway that contains both of them. On the other hand there is no connection between C0 and D3 since for some reason these hypotheses are incompatible with each other (this depends on the method and criteria used to construct hypotheses which is outside the scope of this work).

Similar to connections between roadway subsegment hypotheses, connections in longitudinal direction between strip subsegment hypotheses are modeled by the class `StripSubsegHypConn`. There exists a connection between two strip subsegment hypotheses if and only if they are in consecutive subsegments and there exists a roadway hypothesis in which both strip subsegment hypotheses are part of the same strip. In order to find out how the strip subsegment hypotheses of two given road subsegment hy-

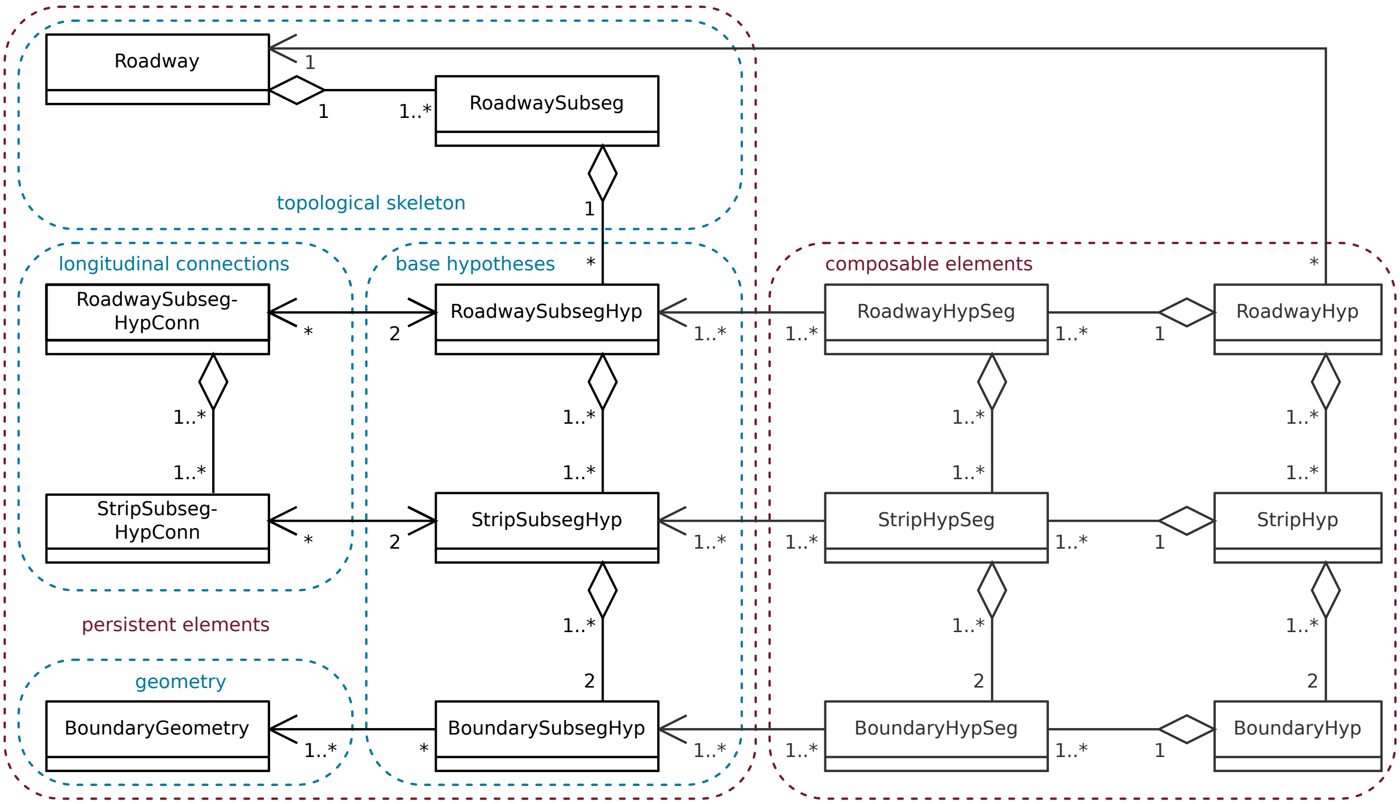

Figure 5. UML class diagram of the multi-hypothesis representation. The persistent elements are mandatory. The roadway (with its location line) and its subsegments are the topological skeleton for the composition of roadway hypotheses. For each subsegment, roadway subsegment hypotheses can be composed of contiguous strip subsegment hypotheses. In longitudinal direction, compatible roadway subsegment hypotheses can be connected. A roadway subsegment hypotheses connection defines which strip subsegment hypotheses of the corresponding roadway subsegment hypotheses are connected. A strip subsegment hypothesis is defined by its boundary subsegment hypotheses. A boundary subsegment hypothesis refers to a subsegment of the geometric description of a detected or virtual boundary. The composable elements can be used to describe a single roadway hypothesis or a continuous lane hypothesis, for example.

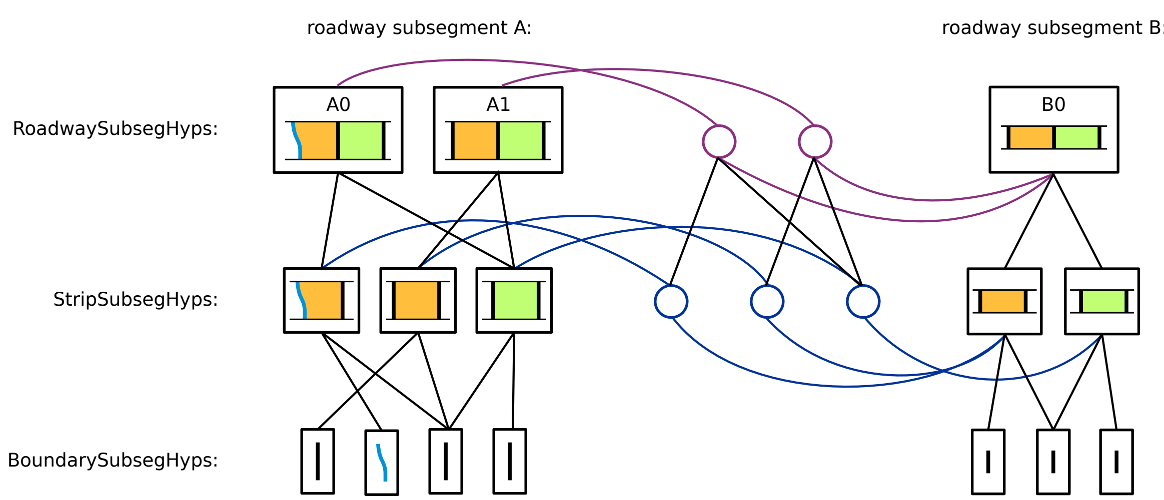

Figure 6. A visual representation of the associations between objects of the classes `RoadwaySubsegHyp`, `StripSubsegHyp`, `BoundarySubsegHyp`, `RoadwaySubsegHypConn` (the purple circles) and `StripSubsegHypConn` (the blue circles), corresponding to the roadway subsegments A and B of the example depicted by figure 4.

potheses are connected, a list with references to the corresponding `StripSubsegHypConn` objects is stored in the `RoadwaySubsegHypConn`. These references are visualized in figure 6 by the lines between the purple and blue circles.

A strip subsegment hypothesis is defined by its two lateral boundaries and its strip type. These boundaries are called *boundary subsegment hypotheses* (modeled by the class `BoundarySubsegHyp`) and describe the hypothesis of a boundary in its corresponding roadway subsegment. A `StripSubsegHyp` refers to the two `BoundarySubsegHyp` objects that correspond to its two boundaries. A `BoundarySubsegHyp` belongs to multiple `StripSubsegHyp` objects, in case two strip subsegment hypotheses of the same roadway subsegment hypothesis are neighbors and therefor have a common boundary, or in case two strip subsegment hypotheses of different roadway subsegment hypotheses share a common boundary. The latter case applies in our example for instance to the right boundary of the orange colored lane in subsegment A: the two different strip subsegment hypotheses of this lane have the right boundary in common as can also be seen in figure 6. A `BoundarySubsegHyp` contains a geometric description of the boundary, the type of the boundary (lane marking, road curb, shoulder, etc.) and other attributes of the boundary that may be relevant for the application.

The elements on the right side of figure 5 can be used to compose elements that describe a single hypothesis. An object of type `RoadwayHyp` describes a single roadway hypothesis, for example.

VI. Future work

Our future work will focus on the integration of junctions, the quantification of uncertainty and the evaluation of strategies to deal with the uncertainty about the roadway.

Based on the approach for the perception of roundabouts presented in [17], a next step is to integrate a multi-hypothesis model of roundabouts into the representation. A future challenge is the integration of arbitrary at-grade intersections. Since there is a wide variety of at-grade intersections, it is difficult to specify a generic model. Regarding a multi-hypothesis representation, it is a challenge to make it easily inferable where and how hypotheses differ. Also the process of perceiving and understanding at-grade intersections is more difficult, because at-grade intersections are basically (large) paved areas with less features (markings, etc.) to detect, but—because of their variety—with much more hypotheses about the desired traffic flow in the intersection.

While for some aspects of behavior generation already the qualitative representation of uncertainty can be useful, a quantification of the uncertainty is very desirable. The presented multi-hypothesis road representation is considered as a platform to study different concepts for the quantification of uncertainty. Quantities that come into consideration are measures of the evidence for a hypothesis, measures of the plausibility of a hypothesis or the probability of a hypothesis.

Evidence arises naturally in the process of perception, since it is usually the first cue for a hypothesis. It is especially relevant in case of observable entities, for example the evidence for a boundary hypothesis might depend on the number of extracted features that support that hypothesis. The plausibility of a hypothesis might measure how well a hypothesis fits with the construction guidelines, that is the plausibility of a hypothesis in itself. But also other information about the environment could be taken into account, for example trajectories of other road users.

The likely most desirable quantity is the probability of a hypothesis, that is the probability that the hypothesis is true given the observations of the environment. It requires to model the dependencies of the elements of the representation within and across the different levels of abstraction. A first approach was presented by Töpfer *et al.* [9]. However, to obtain a probability distribution over multiple hypotheses about the roadway is an open research question.

A further point of our future work is the representation of the uncertainty about the geometry of the roadway. An approach that allows to obtain the uncertainty about the position and orientation of lane boundaries more accurately is presented in [18].

VII. Summary

In this paper it was argued, that a certain degree of uncertainty about the roadway is inevitable. With respect to the understanding of the roadway this yields ambiguity in form of multiple, in general conflicting hypotheses. Yet, this uncertainty is mostly neglected. From our experience, it is necessary to consider the uncertainty about the roadway in the process of behavior generation to achieve a robust behavior of an automated vehicle that solely relies on its onboard sensors, especially in more complex domains than controlled-access highways. For this reason, a multi-hypothesis road representation was presented, that is able to express this uncertainty and that allows an easy interpretation of where and how hypotheses about the roadway differ. To allow to model the roadway without any gaps a strip-based model of the roadway was developed. Future work will focus on the quantification of uncertainty and the evaluation of different strategies to deal with the represented uncertainty in the process of behavior generation.

Acknowledgement

This work was supported by the German Federal Ministry for Economic Affairs and Energy, in the context of the research initiative UR:BAN, on the basis of a decision by the German Bundestag.

References

[1] S. Ulbrich and M. Maurer, "Probabilistic online POMDP decision making for lane changes in fully automated driving," in *Intelligent Transportation Systems (ITSC), 2013 16th International IEEE Conference on*, IEEE, 2013, pp. 2063–2067.

[2] S. Brechtel, T. Gindele, and R. Dillmann, "Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs," in *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, IEEE, 2014, pp. 392–399.

[3] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. Keller, *et al.*, "Making Bertha drive—An autonomous journey on a historic route," *Intelligent Transportation Systems Magazine, IEEE*, vol. 6, no. 2, pp. 8–20, 2014.

[4] J. M. Wille, F. Saust, and M. Maurer, "Stadtpilot: Driving autonomously on Braunschweig's inner ring road," in *Intelligent Vehicles Symposium (IV), 2010 IEEE*, IEEE, 2010, pp. 506–511.

[5] M. Ardelt, C. Coester, and N. Kaempchen, "Highly automated driving on freeways in real traffic using a probabilistic framework," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 13, no. 4, pp. 1576–1585, Dec. 2012.

[6] A. Broggi, P. Cerri, S. Debattisti, M. C. Laghi, P. Medici, M. Panciroli, and A. Prioletti, "PROUD-Public ROad Urban Driverless test: architecture and results," in *Intelligent Vehicles Symposium (IV), 2014 IEEE*, IEEE, 2014, pp. 648–654.

[7] M. Maurer, R. Behringer, S. Furst, F. Thomanek, and E. Dickmanns, "A compact vision system for road vehicle guidance," in *Pattern Recognition, Proceedings of the 13th International Conference on*, IEEE, vol. 3, 1996, pp. 313–317.

[8] A. Bar Hillel, R. Lerner, D. Levi, and G. Raz, "Recent progress in road and lane detection: a survey," *Machine Vision and Applications*, vol. 25, no. 3, pp. 727–745, 2014.

[9] D. Töpfer, J. Spehr, J. Effertz, and C. Stiller, "Efficient scene understanding for intelligent vehicles using a part-based road representation," in *Intelligent Transportation Systems (ITSC), 2013 16th International IEEE Conference on*, 2013, pp. 65–70.

[10] J. Knaup and K. Homeier, "RoadGraph - Graph based environmental modelling and function independent situation analysis for driver assistance systems," in *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on*, IEEE, 2010, pp. 428–432.

[11] S. Ulbrich, T. Nothdurft, M. Maurer, and P. Hecker, "Graph-based context representation, environment modeling and information aggregation for automated driving," in *Intelligent Vehicles Symposium (IV), 2014 IEEE*, IEEE, 2014, pp. 541–547.

[12] P. Bender, J. Ziegler, and C. Stiller, "Lanelets: Efficient map representation for autonomous driving," in *Intelligent Vehicles Symposium (IV), 2014 IEEE*, IEEE, 2014, pp. 420–425.

[13] Forschungsgesellschaft für Straßen- und Verkehrswesen, Ed., *RAA - Richtlinien für die Anlage von Autobahnen*, FGSV-Nr. 202, 2008.

[14] Österreichische Forschungsgesellschaft Straße - Schiene - Verkehr, Ed., *Wörterbuch Verkehrswesen - Begriffsbestimmungen*. 2013.

[15] "Project development and design manual," in, Federal Lands Highway, Ed. 2012, ch. 1, pp. 30–56.

[16] California Department of Transportation, Ed., *Highway design manual*, 2008.

[17] M. Raaijmakers and M. E. Bouzouraa, "Circle detection in single-layer laser scans for roundabout perception," in *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, IEEE, 2014, pp. 2636–2643.

[18] M. T. Schmidt, U. Hofmann, and S. Neumaier, "3D lane boundary tracking using local linear segments," Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on (accepted).