

Scene Understanding of Urban Road Intersections with Description Logic

Britta Hummel¹, Werner Thiemann², and Irina Lulcheva¹

¹ Institut für Mess- und Regelungstechnik, Universität Karlsruhe (TH),
Engler-Bunte-Ring 21, D-76131 Karlsruhe, Germany

hummel|lulcheva@mrt.uka.de

² BlueAge, www.blue-age.de

D-76137 Karlsruhe

thiemann@blue-age.de

Abstract. Road recognition from video sequences has been solved robustly only for small, often simplified subsets of possible road configurations. This contribution argues for a massive augmentation of the amount of prior knowledge to enable the development of more generally applicable estimators. Description Logic is introduced as an expressive knowledge representation formalism for scene understanding.

A Description Logic knowledge base for a wide class of road and intersection configurations is set up. Sensor data from an in-vehicle vision sensor and from a digital map provide evidence for a particular intersection. Partial observability and different abstraction layers of the input data are naturally handled by the formalism.

Deductive inference services are used to narrow down the intersection hypothesis space based on the evidence and the background knowledge, and to retrieve intersection information relevant to a user, i.e. a human or a driver assistance system.

Keywords. Description Logic, Scene Understanding, Formal Reasoning, Road Recognition

1 Motivation

Building on the term Image Understanding ([13]), we define *Intersection Understanding* as the subtask of interpreting an image of a road intersection that enables (at least) the generation of a human-readable, qualitative scene description and an autonomous navigation through the intersection according to traffic rules. The vast majority of current road or intersection recognition systems solely deal with geometric reconstruction. Moreover, the algorithms are restricted to highly specialized domains, e.g. highways. The rare works on intersections focus on one particular, non-complex type of intersection ([8]).

Typically, these methods first extract contour and/or region based cues (edges, their aggregation to lane markings, road texture, ...) from images of an on-board vision sensor. Based on these cues a generic road geometry model of low

dimensionality¹ is instantiated. An additionally available model of the vehicle dynamics can be used for tracking and smoothing the parameter estimates over time (cf. [9] and [11] for an overview).




Fig. 1. Inner city intersection 1(a) image taken by an onboard camera with 50° opening angle, 1(b) map from land surveying office.

Despite more than 20 years of research, these approaches have not yet scaled to less restricted domains. One possible explanation is that an ill-posed estimation problem would then arise. This argument is based on the following observations about inner-city intersections (cf. Fig. 1):

- The abundance of existing intersection geometries necessitates a high-dimensional parameter space.
- A large part of the intersection does not enter the field of view of a standard onboard camera during traversal.
- Dense traffic and inner-city infrastructure lead to a massive amount of occlusion of relevant image clues.
- Frequently omitted markings on the intersection lead to a lack of image cues.
- The presence of an abundance of unmodelled objects feed as noise into the estimation process.
- The image features are of inferior quality, due to an – on average – worse road quality, more variations in marking shape and more rapidly changing lighting conditions.

In brief, a reduced and noisy amount of features contrasts with the necessity of a high dimensional parameter space.

The latter problem can only be addressed by appropriately constraining a high-dimensional parameter space. Such constraints can be derived from general domain knowledge and from specific information about a particular intersection.

¹ The still popular clothoid representation is often approximated by a second or third order polynomial.

Due to the complexity of the domain – thus for reasons of understandability, maintainability and extendability – an explicit formulation of this considerable amount of constraints might be preferable over an implicit hard-wiring in source code.

1.1 Knowledge Representation Formalisms

We introduce *Description Logic* (DL, [1]) as a knowledge representation formalism for intersection understanding. Description Logic is a 2-variable fragment of First Order Logic. It provides several advantages compared to other formalisms:

Most DLs are decidable, which means that sound, complete, and terminating algorithms exist. This is a clear advantage over theorem provers for full first-order logic or Horn clauses with function symbols (e.g. PROLOG). The more recently added expressivity on so-called concrete domains (like the reals), allows for a more natural integration of quantitative constraints than earlier logic formalisms.

DL axioms are similar to human language which –after a training period– allows for understandable and thus maintainable knowledge bases. Due to their nevertheless rigid formal framework, the chance of semantic ambiguities is reduced with respect to human-to-machine and machine-to-machine communication. The integration of several DL knowledge bases (desirable domains for road and intersection understanding include marking types, traffic signs, traffic participants, ...) is a common task, understood better for DL (e.g. [2]) than for maybe any other representation formalism.

In contrast to purely geometric lane recognition algorithms, information of different type and abstraction layer can be fused within one coherent framework, as will be shown in this contribution by fusing digital map and video data.

Whereas *Bayesian Belief Nets* (BBN, [15]), the most prominent representation from the probabilistic world, can capture only propositional, i.e. variable-free, statements, DL provides a clear-cut separation between general knowledge ("A man with a child is a father.") and the individuals in the domain ("John is a man. Emily is John's child."). This allows for modular and thus reusable knowledge bases, as well as for more efficient coding of knowledge ([17], [3]). Some DL systems allow to formulate complex queries on the knowledge base (e.g.: "Retrieve all of John's children!"), which is impossible in purely propositional knowledge bases. Additionally, in contrast to BBN, constraints involving lots of input variables can be formulated without jeopardizing performance.

In contrast to *databases*, it can deal naturally with incomplete information due to its open world semantics².

Open challenges in Description logics involve the representation of spatial relations among individuals, how to deal with limited inference power due to the monotonicity requirement, and how to incorporate probabilistic information.

² Open World semantics denotes that if something cannot be proven to be true, then it is not automatically assumed false.

1.2 Outline

After an introduction into Description Logic (Section 2) we develop a knowledge base for arbitrary roads and intersections (Section 3). Geometric as well as semantical properties are covered. Using input data from a commercially available digital map and from a video sensor, we deductively perform instance classification for missing information (e.g.: "Is this lane a right turn lane?"), query the knowledge base for further entailed information ("Which lane is the vehicle on?"), and show how inconsistencies in the knowledge base and in the sensor data can be detected (Section 4). Current limitations of DL based reasoning and possible remedies for them are discussed in Section 5.

2 Road Network Knowledge Base

In the sequel, a knowledge base \mathcal{KB} for road networks is introduced, which is formalized in Description Logic. Its TBox \mathcal{T} describes general knowledge about road networks. The ABox \mathcal{A} initially captures partial information about a particular road or intersection acquired with onboard vehicle sensors. It is later enhanced by new assertions obtained through ABox inference. The description of the rule base \mathcal{R} is omitted here for brevity.

2.1 The TBox

The TBox of the road network knowledge base introduces relevant concepts of a road network – namely roads, lanes, dividers between roads, road markings and junctions – and the relations that must hold between them.

Taxonomy All atomic concepts are introduced and arranged in a specialization hierarchy called a taxonomy, which is visualized as a UML diagram in Fig. 2(a). Arrows denote inclusion axioms, as e.g.

$$\text{Highway} \sqsubseteq \text{Road} \quad // \quad \textit{All highways are roads.} \quad (1)$$

To denote that individuals of the superclass have to be a member of at least one subclass, one can use a *covering axiom* instead:

$$\begin{aligned} & // \quad \textit{A road is either an autobahn,} \\ & // \quad \textit{highway or an urban road.} \\ \text{Road} & \equiv \text{Autobahn} \sqcup \text{Highway} \sqcup \text{UrbanRoad} \end{aligned} \quad (2)$$

Disjointness between subclasses can optionally be stated, too:

$$\begin{aligned} & // \quad \textit{The set of highways is disjoint from} \\ & // \quad \textit{the set of urban roads and autobahns.} \\ \text{Highway} & \sqsubseteq \neg \text{UrbanRoad} \sqcap \neg \text{Autobahn} \end{aligned} \quad (3)$$




Fig. 2. RoadNetwork TBox 2(a) Taxonomy, 2(b) partonomy. For clarity, only direct parts are visualized, although part-of is transitive.

Mapping atomic concepts to geometric primitives Fig. 2(a) shows that nearly all RoadNetworkElements are also descendents of GeometricEntity, from which three generic types of geometric primitives – GE1, GE2 and GE3 – inherit. The three types are visualized in Fig. 3. Their free parameters are elements of the concrete domain \mathbb{R} . These types are pairwise disjoint but no covering axiom is used. This way, a concept can be of no type, which means that it is composed out of several geometric entities, i.e. complex shapes can be built out of the primitive ones. A more exhaustive description of a preliminary version of this geometry model can be found in [8].




Fig. 3. Geometric entities The three generic shapes that are used to construct the road network.

Spatial representation To capture the relative spatial arrangement of scene elements, three types of spatial relations are introduced (see Fig. 4): The *degree of overlap* of two individuals is described using the common RCC8 representation ([18]). Their *relative orientation* is coarsely discretized into three intervals of the unit circle, namely parallel, perpendicular, and oblique. Their *relative position* is discretized using the eight cardinal directions. The types are introduced as roles with the possible values being their subroles. A description of the spatial arrangement of one individual with respect to another comprises three ABox role assertions, one of each type. The set of subroles of each type is *jointly exhaustive and pairwise disjoint* (JEPD). Unfortunately, this property cannot be modelled in $\mathcal{ALCQHI}_{\mathcal{R}^+}(\mathcal{D}^-)$. As a workaround, the complement of each subrole is explicitly introduced (e.g. `notNTPP`). Each individual pair is required to have either the subrole or its complement stated, and to have exactly one non-complement subrole of each type in total.

Constraints Using the concept names and the spatial roles, *constraints* are formulated on the set of possible models. Without constraints, arbitrary concept and role assertions could be asserted. The types of constraints roughly fall into two categories: *geometric constraints* and *road building regulations*. A subset of geometric constraints are part-of constraints (NTPP, TPP and their inverse) which are arranged in a *partonomy*. The partonomy is visualized in Fig. 2(b), and written in DL as follows:

$$\begin{aligned}
 & // \textit{From all road network elements, only lanes} \\
 & // \textit{dividers and arrows can be part of a Road.} \\
 \text{Road} \sqsubseteq \forall \text{NTPP.}(\text{Lane} \sqcup \text{Divider} \sqcup \text{ArrowMarking})
 \end{aligned} \tag{4}$$

Cardinality constraints can be imposed on the parts as well:

$$\text{Road} \sqsubseteq \exists_{\geq 1} \text{NTPP.Lane} \sqcap \exists_{\leq 6} \text{NTPP.Lane} \tag{5}$$

Compositions are parts that cannot exist without their wholes, denoted by a black diamond ended arrow, written as:

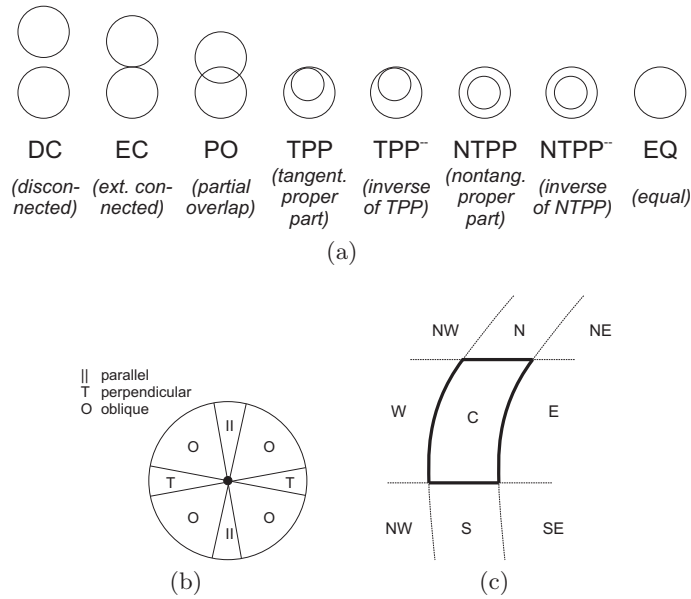


Fig. 4. Spatial Relations 4(a) degreeOfOverlap using RCC8 relations, 4(b) relativeOrientation, 4(c) relativePosition. 4(b) is defined wrt the first individual’s coordinate system. For 4(c), each GeometricEntity has to supply its own ”defintion of perspective” as to what it considers north etc. of it. This is depicted here for GE1.

$$\begin{aligned}
 & // \text{A lane is part of exactly one road.} \\
 \text{Lane} & \sqsubseteq \exists_{=1} \text{NTPP}^- . \text{Road}
 \end{aligned}
 \tag{6}$$

Besides the very basic geometric constraints captured in the partonomy many more geometric constraints of often much more complex nature must hold, which have to be left out for brevity. One example is:

$$\begin{aligned}
 & // \text{An exit lane is connected longitudinally} \\
 & // \text{only to right or left turn lanes.} \\
 \text{ExitLane} & \sqsubseteq \forall \text{isLongitudinallyConnectedTo.} \\
 & (\text{RightTurnLane} \sqcup \text{LeftTurnLane})
 \end{aligned}
 \tag{7}$$

The role isLongitudinallyConnectedTo is asserted iff its parents EC, II and LON hold. The latter is asserted if either N or S are true (cf. Fig. 4). The role isLaterallyConnectedTo is defined analogously. From these, is[east|west|north|south]ConnectedTo are derived with the obvious semantics.

The road building regulations comprise many rules that human drivers have internalized and that are used when approaching an unknown intersection. Few examples include:

```

// Only right turn lanes can be right of right turn lanes.
RightTurnLane  $\sqsubseteq$   $\forall$  hasEastNeighbor.RightTurnLane
// All autobahns and highways are one way roads.
Autobahn  $\sqsubseteq$  Highway  $\sqsubseteq$  OneWayRoad
// A one way road is defined as a road
// which has only one way lanes.
OneWayRoad  $\equiv$  Road  $\sqcap$  ( $\forall$  NTPP.OneWayN  $\sqsubseteq$ 
 $\forall$  NTPP.OneWayS)
// A one way road does not have a u-turn lane.
OneWayRoad  $\sqsubseteq$   $\forall$  NTPP.  $\neg$ UTurnLane

```

(8)

The role `hasNeighbor` is asserted iff `II` and `LAT` holds, and if both individuals are of the same type. A connection `EC` is not necessary. The subroles are derived analogously to `isConnectedTo`.

The hypothesis space that emerges from this combined conceptual-geometrical intersection representation is by orders of magnitude smaller compared to a traditional, purely geometric representation, as the vast majority of hypotheses are elegantly ruled out on the conceptual level using constraints. At the same time, complex geometries are representable with sufficient accuracy, as shown exemplarily for an inner-city area in [8].

2.2 The ABox

A non-stationary video camera and a commercially available digital map along with positioning devices are used as input devices. All are readily available in our experimental vehicles ([20], [7]). However, arbitrary informative (in terms of the terminology introduced by the TBox) sensors can be used, provided they operate in a common coordinate frame. After a brief description of the data registration, the mapping from map and video data to ABox axioms is described.

Coordinate systems and data registration All computations are done within a cartesian, vehicle-centered, two-dimensional coordinate system in the road plane (assuming a locally *flat earth*). A common frame is needed to enable the computation of spatial relations between objects as introduced in sec. 2.1. Map data is transformed using the common UTM projection from geographical to cartesian – and thus length- and angle-preserving – coordinates, and then applying a map matching algorithm as described in [6]. As this version did not deliver lateral position estimates within the road, the ego lane has been manually assigned. However, lane precise estimates will become available in the future.

Video input data is transformed by using a calibrated camera wrt the vehicle coordinate system, and by additionally knowing the camera’s height above ground and yaw angle.

Map input data Digital maps are produced for navigation purposes only, which means that the road network topology is captured correctly, but not the geometry. Roads are represented in the form of coarsely digitized straight line segments whose start and end coordinates coincide with junctions. Junctions are merely represented by a coordinate pair ([6]). Few more attributes are available: the road class, ranging from freeway to pedestrian mall, the allowed driving directions, and few more, which are not of importance here. Current work on the map provider side involves including the number of lanes per road.

This data is used to generate geometry patches outside of the logical framework (data not reported). Each generated patch l_i and d_j is then included in the ABox by stating l_i :Lane and d_j :Divider, respectively, and by adding their types of geometric entity, e.g.: l_i :GE3, and the respective concrete domain fillers.

Each road r_k is added by stating

$$\begin{aligned} r_k & : \text{Autobahn} & // \in \text{Autobahn, Highway, UrbanRoad} \\ r_k & : \text{TwoWayRoad} & // \in \text{TwoWayRoad, OneWayRoad} \\ (r_k, l) & : \text{NTPP} & // \text{for each lane } l \text{ that is part of } r_k \end{aligned} \quad (9)$$

For the transition roads $tr_{i,k}$ an axiom $(tr_{i,k}, j) : \text{NTPP}^-$ is added additionally, where j is a generated instance of type Junction.

To state that our map knowledge is assumed complete we add the following *local closed world assumptions*:

$$\begin{aligned} rn & : \exists_{\leq i} \text{NTPP.Road} & // \text{with } i = |R| \\ rn & : \exists_{\leq j} \text{NTPP.Lane} & // \text{with } j = |L| \\ rn & : \exists_{\leq k} \text{NTPP.Divider} & // \text{with } k = |D| \\ rn & : \exists_{\leq l} \text{NTPP.Junction} & // \text{with } l = |J| \end{aligned} \quad (10)$$

The spatial relations of all GeometryEntity individuals with respect to each other, as introduced in Section 2.1, which have been previously computed, are added to the DL via statements of the form $(l_{4,i}, l_{4,j} : \text{II})$ for all i, j , for example.

Video input data Video-based object detections are straightforwardly integrated in the ABox with axioms like, e.g.

$$\begin{aligned} a_1 & : \text{StraightAheadArrow} & // a_1 \text{ is newly introduced.} \\ d_3 & : \text{SingleDashedDivider} \end{aligned} \quad (11)$$

and additional axioms stating the spatial relations to all other GeometricEntities, like, e.g.

$$\begin{aligned} (a_1, d_3) & : \text{DC} & // a_1 \text{ and } d_3 \text{ are disconnected,} \\ (a_1, d_3) & : \text{II} & // \text{parallel,} \\ (a_1, d_3) & : \text{E} & // \text{and } d_3 \text{ is east of } a_1. \end{aligned} \quad (12)$$

which have to be computed (automatically in a straightforward manner) outside of the DL system. As few reasoning as possible should be done *within* an object detection algorithm. For example, don't hastily conclude **Autobahn** from

a detected GuardRail. Instead, exclusively state the detection result of specialised video object detectors and leave the inference to the KB, where these topics are dealt with more thoroughly. In this example, a guard rail can occur on highways and even on urban roads.

It is worth mentioning that the ABox can be built incrementally online, just as usually more data about the intersection will become available during approach. Instead of recomputation of the ABox just an additional axiom has to be added.

3 Experiments

The knowledge base is denoted with $\mathcal{KB} = (\mathcal{T}, \mathcal{A}, \mathcal{R})$. \mathcal{T} is of the form described in sec. 2.1, \mathcal{A} is initially empty. The description of the rule base \mathcal{R} has been omitted for brevity. After a description of the features detected by the video sensor and the map, classical deduction is used to draw conclusions about the properties of the intersection. At first, the types of several lanes are deduced using instance classification. Then it is shown how the ego lane of the vehicle is inferred using entailment. Eventually it is demonstrated how both the sensor data and the domain knowledge can be tested for inconsistencies.

3.1 Evidence

We are approaching the intersection depicted in Figures 1 and 5(a). The digital map is automatically processed to generate ABox axioms as described in sec 2.2.

The map matching correctly yields that the vehicle is on r_4 :

$$\begin{aligned}
 \text{egovehicle} & : \text{Vehicle} && // \text{ create new vehicle} \\
 (\text{egovehicle}, r_4) & : \text{NTPP}^- && // \text{ is inverse proper part of road } r_4 \\
 (\text{egovehicle}, r_4) & : \parallel && // \text{ is parallel to road } r_4
 \end{aligned} \tag{13}$$

The object detectors described in [4] and [21] have processed an image taken one second before the one shown in Fig. 1(a) and detected the two dividers and the arrow shown in Fig. 5(b). The following statements are thus additionally automatically added to the ABox:

$$\begin{aligned}
 \text{arrow}_1 & : \text{StraightAheadArrow} \\
 \text{divider}_1 & : \text{RoadCurb} \\
 \text{divider}_2 & : \text{Marking20-50} \\
 & // \text{ known spatial relations (Fig. 4)} \\
 & // \text{ wrt all other geometric entities:} \\
 (\text{arrow}_1, \text{egovehicle}) & : \parallel && // \text{ parallel,} \\
 (\text{arrow}_1, \text{egovehicle}) & : \text{S} && // \text{ south of,} \\
 (\text{arrow}_1, \text{egovehicle}) & : \text{DC} && // \text{ and disconnected.} \\
 & \dots
 \end{aligned} \tag{14}$$




Fig. 5. Sensor Input 5(a) Digital map, plus the number of lanes. 5(b) Illustration of video-based object detectors.

3.2 Task 1: Determining lane types (Instance Classification)

Several assertions are immediately available after *classifying* the individuals. Classification can be seen as a special case of *entailment* where the entailed axiom is a concept assertion.

Results 1-4:

$$\begin{aligned}
 \mathcal{KB} \models l_{4.1} : \text{OneWayN} \sqcap \text{BicycleLane} \sqcap \\
 \text{StraightAheadLane} \\
 l_{4.2} : \text{OneWayN} \sqcap \text{CarLane} \sqcap \\
 \text{StraightAheadLane} \\
 l_{4.3} : \text{OneWayS} \\
 l_{4.4} : \text{OneWayS}
 \end{aligned} \tag{15}$$

All classification results can be automatically added to the ABox. Although only the classification results for Lane individuals of road r_4 are given in Eq. 15, all ABox individuals, e.g. all dividers d_j , are classified this way.

Results were given here without proof. To provide a better understanding on how additional assertions are inferred, a proof sketch for a particular entailment query is given in the next subsection.

3.3 Task 2: Determining the vehicle's ego lane (Entailment)

The ego lane of our vehicle is not determined by the positioning device. However, the ego lane can be deduced from the available domain knowledge and the sensor data. Querying for assertions that are *entailed* in this knowledge base yields:

Result5:

$$\begin{aligned}
 // \text{ The vehicle is on lane } l_{4.2}. \\
 \mathcal{KB} \models (\text{egovehicle}, l_{4.i}) : \text{NTPP}^- , \text{ iff } i = 2 .
 \end{aligned} \tag{16}$$

Proof sketch:

The TBox of the \mathcal{KB} contains the following statements:

$$\begin{aligned}
& // \textit{Marking50-20 is a divider for bicycle lanes.} \\
& \textit{Marking50-20} \sqsubseteq \exists \textit{isLaterallyConnectedTo}.\textit{BicycleLane} \\
& // \textit{This type of arrow only occurs on car lanes.} \\
& \textit{Arrow} \sqsubseteq \forall \textit{NTPP}^-. \textit{CarLane} \\
& // \textit{Bicycle lanes are not next to each other.} \\
& \textit{BicycleLane} \sqsubseteq \forall \textit{hasNeighbor}.\textit{CarLane}
\end{aligned} \tag{17}$$

These deductively lead to:

Result 5a:

$$\begin{aligned}
& // \textit{The driver's lane is a car lane (fortunately :)}, \\
& // \textit{and right of it, there is a bicycle lane.} \\
& \mathcal{KB} \models \textit{egovehicle} : \exists \textit{NTPP}^-. (\textit{CarLane} \sqcap \\
& \quad \exists \textit{hasEastNeighbor}.\textit{BicycleLane})
\end{aligned} \tag{18}$$

From the third axiom in Eq. 17 we know as well:

Result 5b:

$$\begin{aligned}
& // \textit{Right of the bicycle lane there can only be a car lane.} \\
& \mathcal{KB} \models \textit{egovehicle} : \exists \textit{NTPP}^-. (\textit{CarLane} \sqcap \\
& \quad \exists \textit{hasEastNeighbor}.\textit{BicycleLane} \sqcap \\
& \quad \forall \textit{hasEastNeighbor}.\textit{CarLane})
\end{aligned} \tag{19}$$

The remaining necessary TBox axioms are described only textually for brevity: A lane with a straight ahead arrow is a straight ahead lane. Right neighbors of straight ahead lanes are only straight ahead or right turn lanes. Bicycle lanes do not occur between lanes of the same turning lane type. Therefore,

Result 5c:

$$\begin{aligned}
& // \textit{Right of the bicycle lane can only be a right turn lane.} \\
& \mathcal{KB} \models \textit{egovehicle} : \exists \textit{NTPP}^-. (\textit{CarLane} \sqcap \\
& \quad \exists \textit{hasEastNeighbor}.\textit{BicycleLane} \sqcap \\
& \quad \forall \textit{hasEastNeighbor}.\textit{RightTurnLane})
\end{aligned} \tag{20}$$

Road r_1 is a one way road towards the junction. Therefore, r_4 cannot have any right turn lane at all. Therefore,

Result 5d:

$$\begin{aligned}
& // \textit{There is no more lane right of the bicycle lane.} \\
& \mathcal{KB} \models \textit{egovehicle} : \exists \textit{NTPP}^-. (\textit{CarLane} \sqcap \\
& \quad \exists \textit{hasEastNeighbor}.\textit{BicycleLane} \sqcap \\
& \quad \neg \exists \textit{hasEastNeighbor})
\end{aligned} \tag{21}$$

From this, together with the closure assumption from Eq. 10 we can deductively infer that the ego vehicle is a proper part of (geometrically, concerning its projection on the road plane (see sec. 2.2)) lane $l_{4.2}$. \square

This consequence is confirmed by looking at the corresponding land surveying office map form Fig. 1(b). As entailment is a standard reasoning task this fact is obtained instantly with an appropriate reasoner.

Likewise, many further assertions about the intersection are entailed in the knowledge base. They can be viewed as inferred constraints on the intersection hypothesis space. Those hypotheses that have not been ruled out can then be tested using variants of the classical lane recognition methods described in the introduction. An initial hypothesis testing algorithm has been developed in [8].

3.4 Task 3: Detecting Inconsistencies (Satisfiability)

Inconsistencies in both TBox and ABox are detected straightforwardly by checking for TBox satisfiability and ABox consistency, respectively. Defining a new TBox concept by stating

$$\begin{aligned} \text{TwoWayAutobahnLane} &\equiv \\ &\text{TwoWayLane} \sqcap \exists \text{NTPP}^- . \text{Autobahn} \end{aligned} \quad (22)$$

or, alternatively, stating in the ABox, that

$$\text{lane}_1 : \text{TwoWayLane} \sqcap \exists \text{NTPP}^- . \text{Autobahn} \quad (23)$$

will immediately lead to an unsatisfiable TBox in the former case as `TwoWayAutobahnLane` will never have any instances, and to an inconsistent ABox in the latter.

4 Towards non-monotonic reasoning

Classical logic is confined to purely *deductive reasoning*, which poses limits to its expressivity. First, once a conclusion is sustained by a valid argument, this argument can never be invalidated, no matter which new assertions are added. This is known as *monotonicity*. This contrasts the modern understanding of vision, that hypotheses are generated based on partially missing evidence via *jumping to conclusions* ([5]), which implies that the arrival of new information will oftentimes lead to the withdrawal of hypotheses. This process is also known as *belief revision*. Second, deduction cannot create new individuals. Transferred to scene interpretation, this amounts to delivering a complete low-level a priori segmentation of the scene. However, the past 50 years of research in computer and biological vision have shown that a purely data-driven segmentation is not feasible, as already low-level segmentation is crucially dependent on top-down input from higher processing levels ([10]).

In accordance with [14], [16], [19], we conclude that classical deductive reasoning is not sufficient in general for real-world scene interpretation. Instead,

deductive and hypothetical reasoning must be combined. This way, the soundness of deduction, as demonstrated in this contribution, can be united with the more far-reaching conclusions possible in non-monotonic reasoning. *Logical model construction* can be seen as an instance of hypo-deductive reasoning. From the model construction perspective, scene interpretation amounts to incrementally constructing a (partial) logical model of the TBox axioms that is consistent with the ABox axioms ([14], [19]).

The development of a model construction algorithm that is suitable for scene interpretation is the focus of ongoing work. The model construction algorithm enhances the reasoning capabilities described in this chapter by *incrementally* hypothesizing new individuals, i.e. lane patches in this example, which have not been asserted in the ABox yet, based on the evidence and on previous hypotheses. Additionally, the *set of all models*, i.e. the set of all plausible intersection hypotheses in this example, can be constructed if desired.

In general, a model construction algorithm proceeds in close analogy to the *Tableau Calculus* algorithms that are used for satisfiability testing of TBoxes and consistency testing of ABoxes (cf. [1]). They apply a set of so-called consistency-preserving completion rules to the original ABox. The presence of \sqcup and $\exists_{\leq n}$ symbols in the TBox triggers so-called non-deterministic rules which split the ABox in a depth-first-search way. The process stops when no more rules can be applied or when an obvious contradiction occurs. In the first case the generated model satisfies all TBox and all ABox axioms. The latter proves that the knowledge base is inconsistent.

Although tableau calculi are suitable to scene interpretation as new individuals are generated and as – in principle – the set of all models of the knowledge base can be generated, some modifications are necessary. Implementations of tableau calculi do not output the model but a mere yes/no answer, they are highly optimised and therefore produce rather ”canonical models”, and they stop after having found one model, preferring simpler ones. In scene interpretation, a preference for a simple model makes no sense, and in many cases it is desirable to return all models instead of just one. [12] have shown an initial sketch of how a model construction could be implemented on top of the RACER reasoner, using its available inference services and query language.

5 Conclusion

This contribution argues for a general paradigm shift towards a stronger acknowledgement of the role of knowledge engineering in real world high-level scene interpretation tasks. Up to date, no satisfying knowledge representation and reasoning framework for such estimation tasks exists.

We introduced Description Logic extended by rules as a knowledge representation formalism for the sensor-based understanding of complex roads and intersections. It was demonstrated how highly incomplete sensor data, coming on various abstraction layers, can be fused within one coherent and semantically sound framework. Data can be processed iteratively on arrival, not requiring

a recomputation of the hypothesis space. The stated domain knowledge is extendable to other domains like traffic signs, traffic participants, etc. Deductive reasoning was used to constrain and query the intersection hypothesis space, and to detect inconsistencies in the sensor data as well as in the stated domain knowledge.

In summary, the more recent developments in logic provide a highly promising framework for developing, constraining and querying the large and complex hypothesis spaces that are typical for image understanding tasks.

Future work includes refinement of the spatial relations, and embedding the existing knowledge base and reasoning services in a model construction framework, which is capable of hypothesizing new individuals and of non-deterministic assertions, to incrementally build the set of all logical models of an intersection.

6 Acknowledgement

This work was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) within the scope of the Transregional Collaborative Research Centre on Cognitive Automobiles (SFB/Tr 28).

References

1. F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.
2. A. Borgida and L. Serafini. Distributed description logics: Assimilating information from peer sources. *J. Data Semantics*, 1:153–184, 2003.
3. P. Domingos. What’s missing in AI: The interface layer. In P. Cohen, editor, *Artificial Intelligence: The First Hundred Years*. AAAI Press, to appear.
4. C. Duchow. A novel, signal model based approach to lane detection for use in intersection assistance. In *IEEE Intelligent Transportation Systems Conference*, 2006.
5. R. L. Gregory. Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society B (Biological Sciences)*, 352(1358):1121–1128, August 1997.
6. B. Hummel. *Dynamic and Mobile GIS: Investigating Changes in Space and Time*, chapter Map Matching for Vehicle Guidance. CRC Press, 2006.
7. B. Hummel, S. Kammel, T. Dang, C. Duchow, and C. Stiller. Vision-based path-planning in unstructured environments. In *Proceedings of the IEEE Intelligent Vehicles Symposium*, pages 176–181, 2006.
8. B. Hummel, Z. Yang, and C. Duchow. Kreuzungsverstehen – ein wissensbasierter Ansatz. *IT-Schwerpunktheft Fahrerassistenzsysteme*, 1:5–16, 2007.
9. V. Kastinaki, M. Zervakis, and K. Kalaitzakis. A survey of video processing techniques for traffic applications. *Image and Vision Computing*, 21:359–381, 2003.
10. T. Lee, D. Mumford, R. Romero, and V. Lamme. The role of the primary visual cortex in higher level vision. *Vision Research*, 38:2429–2454, 1998.
11. J. C. McCall and M. M. Trivedi. Video-based lane estimation and tracking for driver assistance: Survey, system, and evaluation. *IEEE Transactions on Intelligent Transportation Systems*, 7:1:20–37, 2006.

12. R. Möller and B. Neumann. Ontology-based reasoning techniques for multimedia interpretation and retrieval. In Y. Kompatsiaris P. Hobson, editor, *Semantic Multimedia and Ontologies : Theory and Applications*. to appear, 2007.
13. B. Neumann. Bildverstehen – ein Überblick. In G. Görz, editor, *Einführung in die künstliche Intelligenz*, pages 559–588. Addison-Wesley, 1993.
14. B. Neumann and R. Möller. On scene interpretation with description logics. In H. I. Christensen and H.-H. Nagel, editors, *Cognitive Vision Systems: Sampling the Spectrum of Approaches*, volume 3948 of *Lecture Notes in Computer Science*, pages 247–278. Springer Berlin / Heidelberg, 2006.
15. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
16. D. Poole. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.
17. D. Poole. Logic, knowledge representation and bayesian decision theory. In *International Conference on Computational Logic (CL)*, 2000.
18. D. A. Randell, Z. Cui, and A. Cohn. *A Spatial Logic Based on Regions and Connection*, pages 165–176. Morgan Kaufmann, San Mateo, California, 1992.
19. C. Schröder. *Bildinterpretation durch Modellkonstruktion: Eine Theorie zur rechnergestützten Analyse von Bildern*. DISKI 196, Infix, 1999.
20. C. Stiller, G. Färber, and S. Kammel. Cooperative cognitive automobiles. In *Proc. Intelligent Vehicles Symposium*, 2007.
21. Z. Yang. *Fahrbahngeometriemodellierung und videobasierte Pfeilmarkierungserkennung*. Master’s thesis, Universität Karlsruhe, Institut für Mess- und Regelungstechnik, 2006.