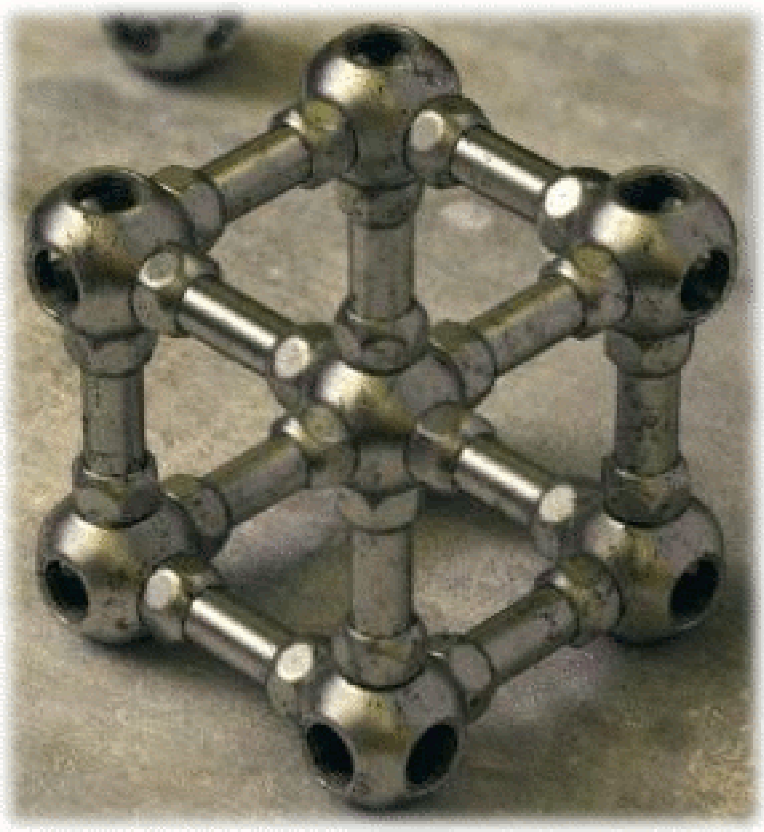

Encoding Spatial Domains with Relational Bayesian Networks [1]

Author(s):

Valquiria Fenelon
Britta Hummel
Paulo E. Santos
Fabio G. Cozman

[1]This work was supported by Fapesp Project LogProb, grant 2008/03995-5, São Paulo, Brazil.

Encoding Spatial Domains with Relational Bayesian Networks

Valquiria Fenelon[1] and **Britta Hummel**[2] and **Paulo E. Santos**[3] and **Fabio G. Cozman**[4]

Abstract. This paper investigates an encoding of a probabilistic TBox using relational Bayesian networks that are specified through a probabilistic description logic. The probabilistic description logic extends the popular $\mathcal{ALC}$ language; on top of this language we add a few operations that are needed to represent the cardinal direction calculus. Using such resources we model roads containing lanes, and vehicles containing digital maps, GPS and video cameras.

1 Introduction

In this paper we study the combination of description logics and relational Bayesian networks as a language to encode qualitative spatial reasoning (QSR). Bayesian networks can be used to represent uncertainty in propositional domains [26], while relational Bayesian networks lift the representation to first-order. Our strategy is to restrict the full generality (and full complexity) of relational Bayesian networks by focusing on a class of relational Bayesian networks that can be specified using a probabilistic description logic called CR$\mathcal{ALC}$ [7, 6]. In a recent publication [30] we have discussed the use of CR$\mathcal{ALC}$ to encode a subset of a cardinal direction calculus [14, 19], but we did so by restricting some features of this calculus. In this paper we remove some of these restrictions by adding a few elements to our relational Bayesian networks (elements that cannot be directly expressed by CR$\mathcal{ALC}$ but that do not introduce substantial complexity). We then investigate the use of the resulting qualitative spatial reasoning formalism to handle queries about a traffic scenario.

We focus on *lane recognition* tasks. Lane recognition research has traditionally focused on estimating the geometric properties of the lane in front of the vehicle using on-board imaging devices (McCall and Trivedi [22] provide an overview of this area). Only a few attempts have been made at inferring *functional* properties of lanes [17], such as the permitted driving directions (e.g. *going up / going down* the road), the permitted turning directions (e.g. *right turn/straight ahead/left turn*), or the permitted traffic participants (e.g. *motor vehicles/emergency vehicles/cyclists*). Existing on-board sensors can only provide a highly incomplete picture of the functional road configuration, with substantial uncertainty.

This paper is organized as follows. After a literature review in the next section, Sections 3 and 4 present, respectively, a formalization of the chosen application scenario and its implementation, with emphasis on features that we have added to the CR$\mathcal{ALC}$ specification and that go beyond our previous publication [30]. We note that the present paper revisits material from this previous publication, contributing mainly on the extensions of CR$\mathcal{ALC}$ that handle role hierarchies and disjoint concepts needed in the cardinal direction calculus. Conclusions are left to Section 5.

[1] Valquiria is with Escola Politecnica, USP, S. Paulo, Brazil
[2] Britta is with Karlsruhe University, Germany
[3] Paulo is with FEI, S. Paulo, Brazil
[4] Fabio is with Escola Politecnica, USP, S. Paulo, Brazil

2 Literature Overview

This section reviews relevant literature on QSR, cardinal direction calculus, probabilistic description logics, and CR$\mathcal{ALC}$; this material is mostly lifted from our previous publication [30].

2.1 Qualitative spatial reasoning

The aim of Qualitative Spatial Reasoning (QSR) is the logical formalisation of knowledge from elementary spatial entities, such as spatial regions, line segments, cardinal directions, and so forth, as surveyed in [4, 5]. Relevant to the present work are the developments of spatial formalisms for computer vision and robotics. The first proposal for a logic-based interpretation of images is described in [28], where image interpretation is reduced to a constraint satisfaction problem on a set of axioms. Inspired by these ideas, [21] proposes a system that generates descriptions of aerial images, which more recently received a descriptive logic enhancement [24].

A spatial system based on spatio-temporal histories for scene interpretation was investigated in [15], which was inspired on an earlier proposal for learning event models from visual information [13]. More recently, [3] proposes a system that uses multiple spatio-temporal histories in order to evaluate an image sequence. A logic formalisation of the viewpoint of a mobile agent was presented in [27], and was further explored in the interpretation of scenes within a mobile robotics scenario in [29]. In [17], functional and geometric properties of roads and intersections could be inferred using an expressive, yet deterministic, description logic in combination with on-board vehicle sensors.

These approaches do not handle uncertainty, which is either left for the low-level processing [3] or simply ignored [29].

2.2 Cardinal direction calculus

The cardinal direction calculus (CDC) [14] is a formalism for reasoning about cardinal directions between spatial objects. The major reasoning task that CDC is concerned with is to infer the direction between two objects A and C, from the known directions between A and (another object) B and between B and C. The basic part of the calculus has nine relations: equal (eq), north (n), east (e), west (w), south (s), northwest (nw), northeast (ne), southeast (se) and southwest (sw).

This paper defines a CDC inspired on the formulation given in [19], where spatial objects are points in a two-dimensional space and the cardinal directions between two objects A and B are defined as the two projections of the straight line from A to B: one on the axis South-North and the other on the axis West-East.

In this paper we assume that each road defines its local cardinal direction system, whereby the direction *"South-North"* goes from the origin of the road towards its end, following the road's centre line. In other words, the *"South-North"* direction between two objects A and B *on the road* are defined as the projection of the line from A to B on the road's centre line. The *"East-West"* direction is defined at every point of the road as the continuous orthogonal line to the tangent of the centre line at that point. Figure 1 shows an example of this local CDC.

In order to make clear that we are not dealing with global cardinal directions (while also taking inspiration of the dynamic nature of a traffic scenes), this paper refers to the directions *going down* and *going up* (the road), instead of resp. *"South"* and *"North"*, and *right–left* instead of *"East"–"West"*.

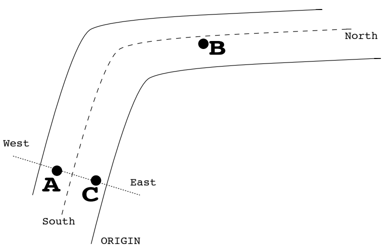

Figure 1. The local cardinal system for roads: A is to the south of B and to the west of C.

2.3 Probabilistic description logics

Description logics (DLs) are fragments of first-order logics originated in the 1970s as a means to provide a formal account of frames and semantic networks. In general terms, description logics are based on *concepts* , which represent sets of individuals (such as Plant or Animal); and *roles*, which denote binary relations between individuals, such as fatherOf or friendOf. Set intersection, union and complement are usual operators found in DLs, as well as some constrained forms of quantification. A key feature of most description logics is that their inference is decidable [1].

In recent years there have been an increasing interest in the combination of probabilistic reasoning and logics (and with description logics in particular) [25, 11, 23]. This combination is not only motivated by pure theoretical interest, but it is very relevant from an application standpoint in order to equip a reasoning system with relational inferences capable of making also probabilistic assessments.

In [18] a number of distinct probabilistic logics were proposed where probabilities were defined over subsets of domain elements. These logics, however, have difficulties to handle probabilistic assertions over individuals, as statistical information over the domain does not preclude information about individuals (this is known as the *direct inference* problem [2]). This problem is also present in various formalisms, as summarised in [11].

The direct inference problem is solved in [12] by adopting probabilities only on assertions. An alternative way around the direct inference problem is to assign probabilities to subsets of interpretations, rather than subsets of the domain. This solution was assumed in [9, 10] and also in CR$\mathcal{ALC}$.

2.4 CR$\mathcal{ALC}$: credal $\mathcal{ALC}$

This section summarizes CR$\mathcal{ALC}$, a probabilistic extension of the popular $\mathcal{ALC}$ description logic [31]. CR$\mathcal{ALC}$ inherits all the constructs of $\mathcal{ALC}$, summarised as follows. The basic vocabulary of $\mathcal{ALC}$ contains individuals, concepts (sets of individuals) and roles (binary relations of individuals). Given two concepts C and D, they can be combined to form new concepts from *conjunction* ($C \sqcap D$), *disjunction* ($C \sqcup D$), *negation* ($\neg C$), *existential* restriction ($\exists r.C$) and *value restriction* ($\forall r.C$).

A concept *inclusion*, $C \sqsubseteq D$, indicates that the concept D contains the concept C and a *definition*, $C \equiv D$, indicates that the concepts C and D are identical. The set of inclusions and definitions constitute a *terminology*. In general, a terminology is constrained to be acyclic, i.e., no concept can refer to itself in inclusions or definitions.

The semantics of $\mathcal{ALC}$ is defined by a domain $\mathcal{D}$ and an interpretation function $\mathcal{I}$, which maps: each individual to a domain element; each concept to a sub-set of $\mathcal{D}$; and, each role to a binary relation $\mathcal{D} \times \mathcal{D}$, such that the following holds: $\mathcal{I}(C \sqcap D) = \mathcal{I}(C) \cap \mathcal{I}(D)$; $\mathcal{I}(C \sqcup D) = \mathcal{I}(C) \cup \mathcal{I}(D)$; $\mathcal{I}(\neg C) = \mathcal{D} \backslash \mathcal{I}(C)$; $\mathcal{I}(\exists r.C) = \{x \in \mathcal{D} | \exists y : (x,y) \in \mathcal{I}(r) \wedge y \in \mathcal{I}(C)\}$; $\mathcal{I}(\forall r.C) = \{x \in \mathcal{D} | \forall y : (x,y) \in \mathcal{I}(r) \rightarrow y \in \mathcal{I}(C)\}$.

An inclusion $C \sqsubseteq D$ holds if and only if $\mathcal{I}(C) \subseteq \mathcal{I}(D)$, and a definition $C \equiv D$ holds if and only if $\mathcal{I}(D) = \mathcal{I}(D)$. For instance, $\mathsf{C} \sqsubseteq (\exists\ \mathsf{hasSibling.Woman}) \sqcap (\forall \mathsf{buys.(Fish} \sqcup \mathsf{Fruit}))$ indicates that C contains only individuals who have sisters and buy fruits or fishes.

In the probabilistic version of $\mathcal{ALC}$ (CR$\mathcal{ALC}$), on the left hand side of inclusions/definitions only concepts may appear. Given a concept name C, a concept D and a role name r, the following probabilistic assessments are possible:

$$P(C) \in [\underline{\alpha}, \overline{\alpha}], \quad (1)$$
$$P(C|D) \in [\underline{\alpha}, \overline{\alpha}], \quad (2)$$
$$P(r) \in [\underline{\beta}, \overline{\beta}]. \quad (3)$$

We write $P(C|D) = \underline{\alpha}$ when $\underline{\alpha} = \overline{\alpha}$, $P(C|D) \geq \underline{\alpha}$ when $\underline{\alpha} < \overline{\alpha} = 1$, and so on.

In order to guarantee acyclicity, no concept is allowed to use itself in deterministic (or probabilistic) inclusions and definitions.

The semantics of CR$\mathcal{ALC}$ is based on probabilities over interpretations so that the direct inference problem can be avoided. In other words, probabilistic values are assigned to the set of all interpretations. The semantics of Formula (1) is, thus: for any $x \in \mathcal{D}$, the probability that x belongs to the interpretation of C is in $[\underline{\alpha}, \overline{\alpha}]$. That is,

$$\forall x \in \mathcal{D} : P\Big(\Big\{ \mathcal{I} : x \in \mathcal{I}(C) \Big\} \Big) \in [\underline{\alpha}, \overline{\alpha}].$$

Informally, the semantics can be represented as:

$$\forall x \in \mathcal{D} : P(C(x)) \in [\underline{\alpha}, \overline{\alpha}].$$

The semantics of Expressions (2) and (3) is then:

$$\forall\ x \in \mathcal{D}\ :\ P(C(x)|D(x)) \in [\underline{\alpha}, \overline{\alpha}],$$
$$\forall\ (x,y) \in \mathcal{D} \times \mathcal{D}\ :\ P(r(x,y)) \in [\underline{\beta}, \overline{\beta}].$$

Given a finite domain, a set of sentences in CR$\mathcal{ALC}$ specifies probabilities over all instantiated concepts and roles. In general, a set of probabilities is specified by a set of sentences (for example, one may specify $P(C) \in [0.2, 0.3]$, allowing all probability values in an interval). A few assumptions guarantee that a single probability distribution is specified by a set of sentences: unique-names, point-probabilities on assessments, rigidity of names [6]. Under these assumptions, a finite domain and a set of sentences specify a unique Bayesian network over the instantiated concepts and roles. To compute the probability of a particular instantiated concept or role, one can generate this Bayesian network and then perform probabilistic inference in the network. Because the domains we deal with in this paper are relatively small, we follow this propositionalisation strategy in our examples. For large domains it may be impractical to explicitly generate a Bayesian network. In this case, approximate algorithms can be used and, in particular, algorithms based on variational methods have been developed with success [6].

3 CR$\mathcal{ALC}$ encoding of a traffic scenario

This section presents a formalisation in CR$\mathcal{ALC}$ of a road traffic domain. Incomplete sensor data and domain knowledge in the form of road building regulations are used to solve functional lane recognition tasks. Let *ego-road* and *ego-lane* denote, respectively, the road and the particular lane on which a vehicle is driving.

The scenario we represent consists of a road where each lane goes either *up* or *down*. Dividing every pair of adjacent lanes is either a *dashed divider* or a *solid divider*. The scenario also contains an experimental vehicle equipped with three on-board sensors: a digital map, a GPS and a video camera. The task of the formalism is to estimate two functional properties of the ego-road using on-board vehicle sensors. First, which lane corresponds to the ego-lane? The answer is one element of the set: $\{1, \ldots, n\}$, where n is the number of lanes in the road. This task is derived from the fact that current differential GPS receivers are able to reliably determine a vehicle's ego-road, but not its ego-lane (e.g. [16]). Second, which driving direction does each lane permit? The answer is either up or down the road, relative to the ego-road's coordinate system.

The functional properties of lanes that are adjacent to the ego-lane are poorly picked out by state-of-the-art vehicle sensors. One reason for that is the narrow field of view of cameras pointing in the driving direction, which causes blind spots over a large portion of the adjacent lanes. Besides, other vehicles frequently occlude relevant image cues, such as divider markings, arrow markings, and traffic signs. Finally, some properties are not explicitly given in the form of symbols but need to be derived from the context by the human driver (e.g. right-handed traffic rules, as assumed in this work).

These observations are reflected by the sensor input available to solve this task.

First, a video-based divider marking recognition is available. Such a sensor recognises lane divider markings on the right of the vehicle and classifies them into either dashed or solid divider lines. Hit and false alarm rate of the recognition task, and the confusion table of the classification task, are given in Tables 1(a) and 1(b), respectively.

Second, a vehicle has a map-matched GPS position that retrieves the current road from a digital map and provides the vehicle's driving direction on that road segment, discretised into "going up" or "going down" relative to the road's coordinate system [16]

Third, a digital navigation map is available, providing classification of the road into either one-way or two-way traffic and an estimate for the number of lanes. Table 1(c) is a confusion table for this classification task.

It is worth pointing out that tables 1(a) and 1(c) are based on comparing the algorithm's outcomes with ground truth [17], whereas the data in Table 1(b) was estimated.

Table 1. Sensor model. In the confusion tables (b) and (c), columns denote ground truth and rows denote estimates.

(a) Video: Divider Recognition

Hit rate	.51
FA rate	.23

(b) Video: Divider Classification

	Solid	Dashed
Solid	.80	.067
Dashed	.20	.933

(c) Digital map: Road Classification

	Oneway	Twoway
Oneway	.99	.01
Twoway	.01	.99

A taxonomy of concepts and roles relevant to the traffic task is now presented (mostly from our previous publication [30]). The concept Lane is defined using two primitives, Up and Down; the concept Divider is defined as the union of DashedDivider and SolidDivider, and Vehicle is either going up (GoingUp) or going down (GoingUp):

$$\mathsf{Lane} \equiv \mathsf{Up} \sqcup \mathsf{Down} \quad (4)$$

$$\mathsf{Divider} \equiv \mathsf{DashedDivider} \sqcup \mathsf{SolidDivider} \quad (5)$$

$$\mathsf{Vehicle} \equiv \mathsf{GoingUp} \sqcup \mathsf{GoingDown} \quad (6)$$

In Formulae (7)–(11) and (13) below we use the abbreviation $\mathsf{disjoint}(\mathsf{t}_1, \mathsf{t}_2, \ldots, \mathsf{t}_n)$ to represent the set of statements about pair-wise disjoint terms, i.e., $\mathsf{t}_i \sqsubseteq \neg \mathsf{t}_j \; \forall i, j \in 1, ..., n, i \neq j$.

$$\mathsf{disjoint}(\mathsf{Vehicle, Divider, Lane}) \quad (7)$$

$$\mathsf{disjoint}(\mathsf{Up,Down}) \quad (8)$$

$$\mathsf{disjoint}(\mathsf{DashedDivider, SolidDivider}) \quad (9)$$

$$\mathsf{disjoint}(\mathsf{GoingUp, GoingDown}) \quad (10)$$

$$\mathsf{disjoint}(\mathsf{OnOneWayRoad, OnTwoWayRoad}) \quad (11)$$

The taxonomy of roles consists of CDC relations only. Out of the nine cardinal directions, only three are relevant to the task at hand east (e), west (w) and equal (eq), since the domain does not have cross-roads.

In this paper the cardinal direction "west" is implicit, as it is not directly defined but it is used in some restrictions such as DashedDivider. (this is more efficient than the representation strategy used in our previous publication [30]). Another change from our previous work is that we use a global point of view (bird-eye) with fixed coordinates (north, south, east, west). This simplifies inference through Bayesian networks, as discussed later.

The relation eq represents the fact that a vehicle being located on a particular lane. We have:

$$\mathsf{cdc} \equiv \mathsf{e} \sqcup \mathsf{w} \sqcup \mathsf{eq} \quad (12)$$

$$\mathsf{disjoint}(\mathsf{e, w, eq}) \quad (13)$$

A set of hard constraints is now defined on road building regulations. The first two constraints (Formulae (14) and (15)) formalise the semantics of right-handed traffic: to the right of a lane allowing for traffic *going up* the road (with respect to the road's egocentric coordinate system) there must only be lanes allowing for "going up" traffic, and to the left of traffic *going down* the road there must only be "down" lanes. When a vehicle is "going up" in a lane with direction *up*, to its east there is a "solid divider" or to its west there is a "dashed divider"and there is also a lane to its east that is *up*. Similarly, when a vehicle is "going down" in a lane with direction *down*, to its west there is a "solid divider" or to its west there is a "dashed divider" and there is also a lane to its west that is *down*.

$$\mathsf{GoingUp} \sqsubseteq \exists \mathsf{e}.(\mathsf{SolidDivider} \sqcap \neg\mathsf{Lane}) \sqcup \exists \mathsf{e}.(\mathsf{DashedDivider} \sqcap \mathsf{Up}) \quad (14)$$

$$\mathsf{GoingDown} \sqsubseteq \forall \mathsf{w}.(\mathsf{SolidDivider} \sqcap \neg\mathsf{Lane}) \sqcup \exists \mathsf{w}.(\mathsf{DashedDivider} \sqcap \mathsf{Down}) \quad (15)$$

Formulae (16) and (17) refer to the dividers function, which may be distinct in different countries. A dashed divider divides two lanes with same driving direction, whereas a solid divider either marks the road border or it separates roads with opposing driving directions:

$$\mathsf{DashedDivider} \sqsubseteq (\exists \mathsf{e}.\mathsf{Up} \sqcap \exists \mathsf{w}.\mathsf{Up}) \sqcup (\exists \mathsf{e}.\mathsf{Down} \sqcap \exists \mathsf{w}.\mathsf{Down}) \quad (16)$$

$$\mathsf{SolidDivider} \sqsubseteq (\neg\exists \mathsf{e}.\mathsf{Lane} \sqcap \exists \mathsf{w}.\mathsf{Lane}) \sqcup (\neg\exists \mathsf{w}.\mathsf{Lane} \sqcap \exists \mathsf{e}.\mathsf{Lane}) \sqcup (\exists \mathsf{e}.\mathsf{Up} \sqcap \exists \mathsf{w}.\mathsf{Down}) \sqcup (\exists \mathsf{e}.\mathsf{Down} \sqcap \exists \mathsf{w}.\mathsf{Down}) \quad (17)$$

Finally, the following axiom states that a two-way road has traffic in both directions (Formula (18)).

$$\mathsf{OnTwoWayRoad} \sqsubseteq \exists \mathsf{cdc}.\mathsf{Up} \sqcap \exists \mathsf{cdc}.\mathsf{Down} \quad (18)$$

The probabilities given in Tables 1a-1c can be justified as follows. First, new concepts (prefixed by Sensed) are added as subclasses of Sensor for all probabilistic inputs:

$$\mathsf{Sensor} \equiv \mathsf{SensedOnOneWayRoad} \sqcup \mathsf{SensedOnTwoWayRoad} \sqcup \mathsf{SensedDivider} \quad (19)$$

$$\mathsf{SensedDivider} \equiv \mathsf{SensedDashedDivider} \sqcup \mathsf{SensedSolidDivider} \quad (20)$$

The confusion tables (Tables 1(a)–1(c)) show joint probabilities of an event and its detection by a sensor. These probabilities can be represented as conditional probabilities using the definition of conditional events in terms of their intersections, such as: $P(\text{Solid}_{\text{GT}} \mid \text{Dashed}_{\text{S}}) = P(\text{Solid}_{\text{GT}} \sqcap \text{Dashed}_{\text{S}}) / P(\text{Dashed}_{\text{S}}) = 0.05/0.75 = 0.07$, where the subscript GT denotes "ground truth" and S stands for "sensed". The sensor model from Table 1b (for instance) can now be elegantly formulated as a set of axioms as follows:

$$P(\mathsf{DashedDivider}|\mathsf{SensedDashedDivider}) = 0.93 \quad (21)$$

$$P(\mathsf{SolidDivider}|\mathsf{SensedDashedDivider}) = 0.07 \quad (22)$$

$$P(\mathsf{DashedDivider}|\mathsf{SensedSolidDivider}) = 0.20 \quad (23)$$

$$P(\mathsf{SolidDivider}|\mathsf{SensedSolidDivider}) = 0.80 \quad (24)$$

An analogue set of axioms is used for confusion Table 1. For recognition tasks the sensor model can be translated as follows:

$$P(\mathsf{Divider}|\mathsf{SensedDivider}) = 1 - \text{false alarm rate} = 0.77 \quad (25)$$

$$P(\mathsf{SensedDivider}|\mathsf{Divider}) = \text{hit rate} = 0.51 \quad (26)$$

With the previous axioms, the $\mathcal{T}$Box is fully specified.

DL is used in this work as a specification language from which a Bayesian description is derived. In the present context, a DL description is used to encode high-level knowledge, such as the permitted driving directions. The use of a DL-based probabilistic logics gives us guarantees concerning expressivity and complexity that are not available when one resorts to full first-order probabilistic logic.

4 Coding and running the scenario

Most axioms presented in the previous section are within the basic definitions of CR$\mathcal{ALC}$. However, the original role hierarchies are not even within the scope of $\mathcal{ALC}$ (and, consequently, not within CR$\mathcal{ALC}$). In our previous publication [30] we circumvented these axioms to construct a TBox with the confines of CR$\mathcal{ALC}$. In this paper we wish to follow a different strategy.

As any TBox in CR$\mathcal{ALC}$ is turned into a Bayesian network upon inference, we have handled hole hierarchies and disjoint concepts within the transformation TBox→Bayesian network. Concepts are translated into nodes of the network in such a way that concepts that are to the left are translated into parents of concepts that are to the right. To handle disjoint concepts, we create or-exclusive nodes that are always set to true in the network. We can then determine whether a lane is *up* or *down* using Formulae (16), (17) and (18). Even though concepts in these formulae are defined using roles, we compute their probabilities by conditioning on concepts GoingUp, GoingDow, OnTwoWayRoad, OnTwoWayRoad. The same procedure is used to determine the kind of "divider" that is related to each lane; that is, which divider is to the east of a lane. In this work we focus on the cardinal direction calculus to determine in which lane the vehicle is located. We can determine $P(eq(vehicle, lane))$ using data in the $\mathcal{A}$Box; for instance, if the sensor indicates a divider to the east of the vehicle, and we can then infer the kind of divider for each lane. Hence we can compute the probability of each lane being an ego-lane.

Given the formalisation presented in Section 3 (and the consideration above), the system generated the Bayesian network represented in Figure[5] 2 , where the nodes in red are observed variables, i.e. sensors' states. A detail of this network is shown in Figure 3, that represents a Bayesian Network for one individual (out of the 5 interconnected nets shown in Figure 2). Besides the information in the $\mathcal{A}$Box, there is evidence in the nodes that represent disjoint concepts (that is, the nodes that encode or-exclusive relations) and nodes that indicate whether a network represents a lane, a vehicle, or a divider.

It is now possible to answer the queries specified in Section 3, which correspond to the following:

1. $\text{argmax}_{\mathsf{li}} P((v : \mathsf{li} : eq))$, i.e. li is the lane with maximum probability of being the vehicle's (v) ego-lane .
2. $\forall i : P(\mathsf{li} : \mathsf{Up})$, i.e. for each lane li, the probability of being a Up lane.

Consulting the network in Figure 2 for all of the eight possible states of the three sensors, we obtain the answers presented in Tables 2 and 3 for the queries 1 and 2 respectively (we employ the abbreviations $STWR$ for SensedOnTwoWayRoad and SDD for SensedDashedDivider).

Table 2 shows the most probable lane on which the vehicle v is driving ($\text{argmax}_{li} P((v, \mathsf{li} : eq))$), given the evidences, represented on the first three columns. The first line of the table, for instance, represents the state where the GPS obtained GoingDown, the map sensed that the vehicle was on a *one way road* and the vision system sensed a *solid divider*. Given these evidences the node li with the highest probability (on the network of Figure 2) was l1. This case is shown in Figure 4(c). On the second line of Table 2, however, the GPS and the map sensor remained in the state just described, but the vision sensed a dashed divider (instead of a solid one). In this case, there were two hypotheses with equal probabilities[6]: l2 and l3 (as repre-

[5] For colour image, please refer to the electronic version of this paper.
[6] They differ on the third decimal digit.

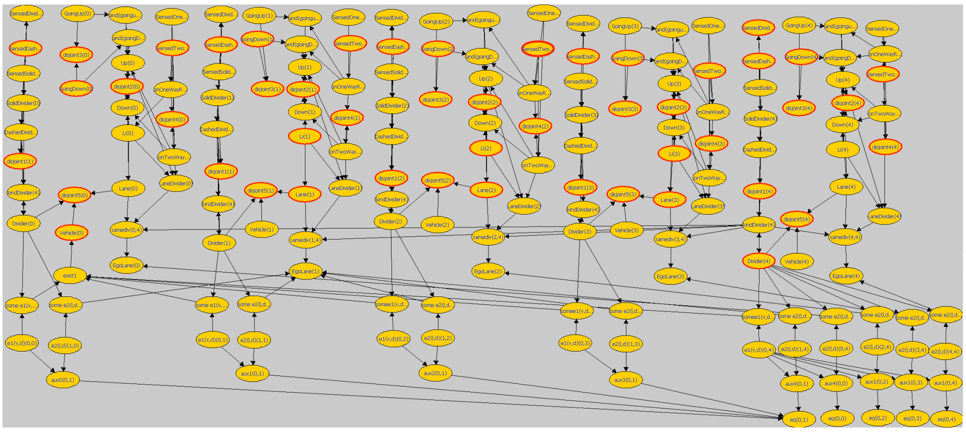

Figure 2. Bayesian Network representing a traffic domain.

sented in Figures 4(a) and 4(b) respectively). The remainder cases on Table 2 are analogous.

Table 3 represents the probabilities for each of the li lanes to be a Down lane, given the evidence on the first three columns. The probability of Up is the complement of the values stated in the table. Take for instance the first line, the highest probability for l1, l2 and l3 is Down, which is consistent with the evidence GoingDown (for the vehicle) and SensedOnOneWayRoad. Similarly for the remainder sensor states represented in the table.

Table 2. Answer to query 1: the probability on the *ego-lane* given the evidence A (expressed on the first three columns).

GPS	map	video	$\text{argmax}_{li} P((v, \text{li} : eq\|A))$
GoingUP	STWR	SDD	
0	0	0	l1
0	0	1	l2 ∨ l3
0	1	0	l2 ∨ l3
0	1	1	l3
1	0	0	l1
1	0	1	l2 ∨ l3
1	1	0	l1
1	1	1	l2

Table 3. Answer to query 2: the probability for the *lane's driving direction* given the evidence A (expressed on the first three columns).

GPS	map	video	l1	l2	l3
GoingUP	STWR	SDD	P(l1:Down|A)	P(l2:Down|A)	P(l3:Down|A)
0	0	0	0.98	0.99	1.00
0	0	1	0.98	0.99	1.00
0	1	0	0.50	0.5	1.00
0	1	1	0.05	0.5	1.00
1	0	0	0.00	0.00	0.00
1	0	1	0.00	0.00	0.00
1	1	0	0.00	0.50	0.99
1	1	1	0.00	0.50	0.99

In this work the queries presented in Table 2 and 3 were run offline. However, for small size scenarios they could be executed in real time.

A network is generated for each individual; in total we have five individuals: 3 lanes, 1 vehicle and 1 divider. The generation of the resulting network is a non-trivial task that is obviously simplified by the use of a probabilistic description language.

(a) (b) (c)

Figure 4. Three traffic scenarios.

5 Concluding remarks

In this paper we have extended our previous efforts on encoding spatial domains with a probabilistic description logic. We still employed CR$\mathcal{ALC}$ as the basic description language but added features that affect the translation of terminologies into Bayesian networks; namely, we added the ability to handle the role hierarchies and the disjoint concepts that appear in spatial domains. The development of more general inference algorithms is the object of our future work.

Overall, the representation of qualitative spatial reasoning with description logics is a recent endeavour [8]. The major difficulty of this task, which we still face in our work, is the representation of transitive relations. Decidability of description logic representations of spatial formalisms were investigated in [20, 8] for a combination of $\mathcal{ALC}$ with a decidable constraint system (called $\mathcal{ALC}(C)$, where C is the constraint system). The investigation of probabilistic extensions of $\mathcal{ALC}(C)$, and whether decidability is maintained, is an interesting issue for future research.

Acknowledgements

This work has been partially supported by Fapesp Project 2008/03995-5 (LogProb). Paulo Santos acknowledges support from CAPES and CNPq, Brazil. Valquiria Fenelon is a graduate student sponsored by CAPES, Brazil. Fabio Cozman acknowledges support from CNPq, Brazil.

REFERENCES

[1] F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, and P.F. Patel-Schneider, *Description Logic Handbook*, Cambridge University Press, 2002.

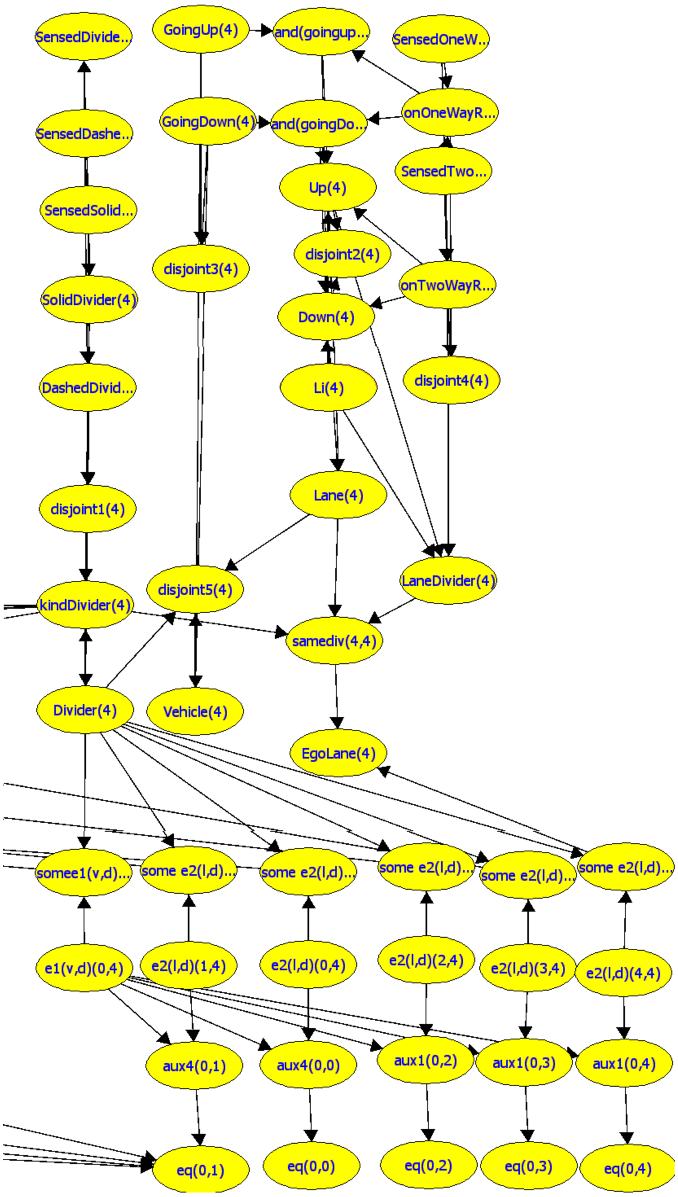

Figure 3. Detail of the Bayesian Network shown in Figure 2.

[2] F. Bacchus, A. Grove, J. Y. Halpern, and D. Koller, 'From statistical knowledge bases to degrees of belief', *Artificial Intelligence*, **87**, 75–143, (1996).

[3] B. Bennett, A. Cohn, and D. Magee, 'Enforcing global spatio-temporal consistency to enhance reliability of moving object tracking and classification', *Künstliche Intelligenz*, **2**, 32–35, (2005).

[4] A. G. Cohn and S. M. Hazarika, 'Qualitative spatial representation and reasoning: An overview', *Fundamenta Informaticae*, **46**(1-2), 1–29, (2001).

[5] A. G. Cohn and J. Renz, 'Qualitative spatial representation and reasoning', in *Handbook of Knowledge Representation*, 551–596, Elsevier, (2008).

[6] F. G. Cozman and R. Polastro, 'Complexity analysis and variational inference for interpretation-based probabilistic description logics', in *Proc. of the Conference on Uncertainty in Artificial Intelligence*, (2009).

[7] F. G. Cozman and R. B. Polastro, 'Loopy propagation in a probabilistic description logic', in *SUM*, pp. 120–133, (2008).

[8] M. Cristani and N. Gabrielli, 'Practical issues of description logics for spatial reasoning', in *Proc. of the AAAI Spring Symposium Benchmarking of Qualitative Spatial and Temporal Reasoning Systems*, pp. 5–10, (2009).

[9] P. C. G. da Costa and K. B. Laskey, 'Of Klingons and starships: Bayesian logic for the 23rd century', in *Conference on Uncertainty in Artificial Intelligence*, (2005).

[10] C. D'Amato, N. Fanizzi, and T. Lukasiewicz, 'Tractable reasoning with Bayesian description logics', in *SUM*, pp. 146–159, Berlin, Heidelberg, (2008). Springer-Verlag.

[11] C. P. de Campos, F. G. Cozman, and J. E. O. Luna, 'Assembling a consistent set of sentences in relational probabilistic logic with stochastic independence', *Journal of Applied Logic: Combining Probability and Logic*, **7**, 137–154, (2009).

[12] M. Dürig and T. Studer, 'Probabilistic ABox reasoning: preliminary results', in *Description Logics*, pp. 104–111, (2005).

[13] J. Fernyhough, A. G. Cohn, and D. C. Hogg, 'Constructing qualitative event models automatically from video input', *Image and Vision Computing*, **18**, 81–103, (2000).

[14] A. U. Frank, 'Qualitative spatial reasoning: Cardinal directions as an example', *International Journal of Geographical Information Science*, **10**(3), 269–290, (1996).

[15] S. M. Hazarika and A. G. Cohn, 'Abducing qualitative spatio-temporal histories from partial observations', in *Proc. of KR*, pp. 14–25, Toulouse, France, (2002).

[16] B. Hummel, *Dynamic and Mobile GIS: Investigating Changes in Space and Time*, chapter Map Matching for Vehicle Guidance, CRC Press, 2006.

[17] B. Hummel, *Description Logic for Intersection Understanding at the Example of Urban Road Intersections*, Phd Thesis, April 2009.

[18] M. Jaeger, 'Probabilistic reasoning in terminological logics', in *Principles of Knowledge Representation (KR)*, pp. 461–472, (1994).

[19] G. Ligozat, 'Reasoning about cardinal directions', *J. Vis. Lang. Comput.*, **9**(1), 23–44, (1998).

[20] C. Lutz and M. Milicic, 'A tableau algorithm for DLs with concrete domains and GCIs', *Journal of Automated Reasoning*, **38**(1–3), 227–259, (2007).

[21] T. Matsuyama and V. S. Hwang, *SIGMA: A Knowledge-Based Image Understanding System*, Plenum Press, New York, U.S., 1990.

[22] 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).

[23] B. Milch and S. Russell, 'First-order probabilistic languages: into the unknown', in *International Conference on Inductive Logic Programming*, (2007).

[24] B. Neumann and R. Möller, 'On scene interpretation with description logics', *Image Vision Comput.*, **26**(1), 82–101, (2008).

[25] A. Nikolov, V. Uren, E. Motta, and A. de Roeck, 'Using the Dempster-Shafer theory of evidence to resolve ABox inconsistencies', in *Workshop on Uncertainty Reasoning for the Semantic Web*, (2007).

[26] J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.

[27] D. Randell, M. Witkowski, and M. Shanahan, 'From images to bodies: Modeling and exploiting spatial occlusion and motion parallax', in *Proc. of IJCAI*, pp. 57–63, Seattle, U.S., (2001).

[28] R. Reiter and A. Mackworth, 'A logical framework for depiction and image interpretation', *Artificial Intelligence*, **41**(2), 125–155, (1989).

[29] P. Santos, H. Dee, and V. Fenelon, 'Qualitative robot localisation using information from cast shadows', in *Proc. of IEEE International Conference on Robotics and Automation*, Kobe, Japan, (2009).

[30] Paulo E. Santos, Britta Hummel, Valquiria Fenelon Pereira, and Fabio G. Cozman, 'Probabilistic logic encoding of spatial domains', in *First International Workshop on Uncertainty in Description Logics (UniDL)*, (2010).

[31] M. Schmidt-Schauss and G. Smolka, 'Attributive concept descriptions with complements', *Artificial Intelligence*, **48**, 1–26, (1991).