



Probabilistic Logic Encoding of Spatial Domains ¹

Author(s):

P. Santos

B. Hummel

V. Fenelon

F. G. Cozman

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

Probabilistic Logic Encoding of Spatial Domains

P. Santos¹, B. Hummel², V. Fenelon^{3,1}, and F. G. Cozman³

¹FEI, São Paulo, Brazil

²University Karlsruhe, Germany

³Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil

Abstract. This paper presents a formalisation of a spatial domain in terms of a qualitative spatial reasoning formalism, encoded in a probabilistic description logic. The QSR formalism chosen is a subset of a cardinal direction calculus and the probabilistic description logic used has the relational structures of the well-known \mathcal{ALC} language, allied with the inference methods of Bayesian Networks. We consider a scenario consisting of a road navigated by an experimental vehicle equipped with three on-board sensors: a digital map, a GPS and a video camera. This paper presents experiments where the proposed formalism is used to answer queries about driving directions, lanes and vehicles.

1 Introduction

Much work in computer vision in the 70's and 80's had as a goal the development of high-level vision, whereby the numerical (or quantitative) processing feeds a symbolic (or qualitative) level of knowledge from where an agent is capable of interpreting the world, and acting in accordance to this interpretation. These early attempts were frustrated by the non-existence (at the time) of efficient algorithms for dealing with uncertainty, of tractable knowledge representation formalisms and also by the rudimentary stage of image-processing algorithms.

Since then, important advancements in Artificial Intelligence (AI) suggest that we may be at the stage of bridging the gap between AI and Computer Vision. The present paper is related to three of these advancements: *Bayesian Networks* [1], which are graphical representations of domain variables that provide efficient probabilistic inference methods; *Description Logics* (DLs) [15], a family of formalisms (sub-sets of first-order logic) that have a positive trade-off between expressivity and complexity; and qualitative spatial reasoning (QSR) [25], that are formalisms for representing and reasoning about space.

The purpose of this paper is to investigate the use of a qualitative spatial reasoning formalism, encoded on a probabilistic description logic, to answer queries about a traffic scenario. The QSR formalism chosen is a sub-set of a cardinal direction calculus [8, 9] and the probabilistic DL used is $CR\mathcal{ALC}$ [23, 28], which has the relational structures of the description logic \mathcal{ALC} [4], allied with the inference methods of Bayesian Networks.

2 Literature Overview

Qualitative spatial reasoning: The aim of QSR is the logical formalisation of knowledge from elementary spatial entities, such as spatial regions, line segments, cardinal directions, and so forth [13, 25]. These formalisms provide the basic machinery for a system to represent and reason about spatial entities on a more abstract level than quantitative methods [25].

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 [2], where image interpretation is reduced to a constraint satisfaction problem on a set of axioms; [3] 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 [14], which was inspired on an earlier proposal for learning event models from visual information [10]. Recently, [16] 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 [11], and was further explored in the interpretation of scenes within a mobile robotics scenario in [21, 31]. In [29], functional and geometric properties of roads and intersections could be inferred using an expressive, deterministic, DL in combination with on-board vehicle sensors.

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. 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 forms of quantification. A key feature of most description logics is that their inference is decidable [15].

In recent years there have been an increasing interest in the combination of probabilistic reasoning and logics (and with description logics in particular) [22, 27, 20]. 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 [5, 6] a number of distinct probabilistic logics were proposed where probabilities were defined over subsets of domain elements. These logics have difficulty in handling probabilistic assertions over individuals, as statistical information over the domain does not imply information about individuals; this is known as the *direct inference* problem [7]. The direct inference problem is solved in [18] by adopting probabilities only on assertions. An alternative way around the direct inference problem is to assign probabilities to subsets of interpretations, as assumed in [17, 26] and is also at the kernel of the probabilistic DL we use here.

3 The credal \mathcal{ALC}

The credal \mathcal{ALC} ($\text{CR}\mathcal{ALC}$) [28] is a probabilistic extension of the \mathcal{ALC} description logic [4]. 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: $\mathcal{I}(C \sqcap D) = \mathcal{I}(C) \cap \mathcal{I}(D)$; $\mathcal{I}(C \sqcup D) = \mathcal{I}(C) \cup \mathcal{I}(D)$; $\mathcal{I}(\neg C) = \mathcal{D} \setminus \mathcal{I}(C)$; $\mathcal{I}(\exists r.C) = \{x \in \mathcal{D} \mid \exists y : (x, y) \in \mathcal{I}(r) \wedge y \in \mathcal{I}(C)\}$; $\mathcal{I}(\forall r.C) = \{x \in \mathcal{D} \mid \forall y : (x, y) \in \mathcal{I}(r) \rightarrow y \in \mathcal{I}(C)\}$. 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 $\mathcal{I}(C) = \mathcal{I}(D)$ (e.g. $C \sqsubseteq (\exists \text{hasSibling.Woman}) \sqcap (\forall \text{buys.}(\text{Fish} \sqcup \text{Fruit}))$ indicates that C contains only individuals who have sisters and buy fruits or fishes).

In the probabilistic version of \mathcal{ALC} (CRALC), 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:

$$P(C) \in [\underline{\alpha}, \bar{\alpha}], \quad (1)$$

$$P(C|D) \in [\underline{\alpha}, \bar{\alpha}], \quad (2)$$

$$P(r) \in [\underline{\beta}, \bar{\beta}]. \quad (3)$$

We write $P(C|D) = \underline{\alpha}$ when $\underline{\alpha} = \bar{\alpha}$, $P(C|D) \geq \underline{\alpha}$ when $\underline{\alpha} < \bar{\alpha} = 1$, 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 CRALC 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 $[\underline{\alpha}, \bar{\alpha}]$. That is,

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

Informally, the semantics can be represented as $\forall x \in \mathcal{D} : P(C(x)) \in [\underline{\alpha}, \bar{\alpha}]$. The semantics of Expressions (2) and (3) is then:

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

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

Given a finite domain, a set of sentences in CRALC specifies probabilities over all instantiated concepts and roles. In general, a set of probabilities is specified by a set of sentences; 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 [28]. So, 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 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 [28].

4 Cardinal Direction Calculus

The cardinal direction calculus (CDC) [8] 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).

We define a CDC inspired on the formulation given in [9], 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 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 we assume that each road defines its local cardinal direction system, whereby the directions “Down-Up” instead of “South-North” goes from the origin of the road towards its end, following the road’s centre line. In other words, the “Down-Up” 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, refer as *right-left*, 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.




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

5 The CRALC encoding of a traffic scenario

This section presents a formalisation in CRALC of a road traffic domain where incomplete sensor data and domain knowledge can be jointly exploited 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 chosen is composed of a road, where each of its lanes has either the direction *going up* or the direction *going 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 the following functional properties of the ego-road using on-board vehicle sensors:

- Which lane corresponds to the ego-lane? 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. [19]).
- Which driving direction does each lane permit, “going up” or “going down”?

Extending these tasks in order to allow queries about turning directions and multiple traffic actors should be straightforward once the above issues are solved.




Fig. 2. (a) original scene and (b) Example for sensor input from on-board camera, digital map, and map-matched GPS from scene (a).

Sensor Input: The sensors input available to solve that task are:

- *Video-based divider marking recognition:* 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.
- *Map-matched GPS position:* retrieves the current road from a digital map and provides the vehicle’s driving direction on that road segment. The algorithm proposed in [19] has been shown to be accurate under batch-processing.
- *Digital navigation map:* provides the classification of the road into either one-way or two-way traffic. 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 [29], whereas the data in Table 1(b) was estimated. A typical sensor input is sketched in Figure 2(b), that shows the information obtained by the sensors on the situation in Figure 2(a).

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

| | (a) | (b) | | (c) | |
|----------|-------------------------------|----------------------------------|----------|-------------------------------------|---------|
| | Video: Divider Recognition | Video: Divider Classification | | Digital map: Road Classification | |
| | | Solid | Dashed | Oneway | Twoway |
| Hit rate | .51 | Solid | .80 .067 | Oneway | .99 .01 |
| FA rate | .23 | Dashed | .20 .933 | Twoway | .01 .99 |

Road Building Regulations: A taxonomy of concepts and roles relevant to the traffic task is set up, in which the concept Lane is defined by the two primitives GoingUp and GoingDown, the concept Divider is defined as the union of DashedDivider and

SolidDivider, and Vehicle is either on a one-way road (OnOneWayRoad) or on a two-way road (OnTwoWayRoad):

$$\text{Lane} \equiv \text{GoingUp} \sqcup \text{GoingDown}$$

$$\text{Divider} \equiv \text{DashedDivider} \sqcup \text{SolidDivider}$$

$$\text{Vehicle} \equiv \text{OnOneWayRoad} \sqcup \text{OnTwoWayRoad}.$$

In Formulae (4)–(7) and (9) we use the abbreviation $\text{disjoint}(t_1, t_2, \dots, t_n)$ to represent the set of statements about pairwise disjoint terms, i.e., $t_i \sqsubseteq \neg t_j \forall i, j \in 1, \dots, n, i \neq j$.

$$\text{disjoint}(\text{Vehicle}, \text{Divider}, \text{Lane}) \quad (4)$$

$$\text{disjoint}(\text{GoingUp}, \text{GoingDown}) \quad (5)$$

$$\text{disjoint}(\text{DashedDivider}, \text{SolidDivider}) \quad (6)$$

$$\text{disjoint}(\text{OnOneWayRoad}, \text{OnTwoWayRoad}). \quad (7)$$

The taxonomy of roles consists of CDC relations only. Out of the nine cardinal directions, only three are relevant to the task at hand right (ri), left (le), since the domain does not have cross-roads, and equal (eq):

$$\text{cdc} \equiv \text{ri} \sqcup \text{le} \sqcup \text{eq} \quad (8)$$

$$\text{disjoint}(\text{ri}, \text{le}, \text{eq}). \quad (9)$$

Next, a set of hard constraints about road building regulations are formulated, making use of the concepts and roles introduced before. The Formulae (10) and (11) 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 “going down” lanes.

$$\text{GoingUp} \sqsubseteq \forall \text{ri}. (\text{GoingUp} \sqcup \neg \text{Lane}) \quad (10)$$

$$\text{GoingDown} \sqsubseteq \forall \text{le}. (\text{GoingDown} \sqcup \neg \text{Lane}). \quad (11)$$

Formulae (12) and (13) refer to the dividers function, which may be distinct in different countries; these axioms holds for right-handed traffic. A dashed divider divides two lanes, a solid divider either marks the road border or it separates roads with opposing driving directions. And, the axiom states that a two-way road has traffic in both directions (Formula (14)).

$$\text{DashedDivider} \sqsubseteq \exists \text{ri}. \text{Lane} \sqcap \exists \text{le}. \text{Lane} \quad (12)$$

$$\text{SolidDivider} \sqsubseteq \neg \exists \text{ri}. \text{Lane} \sqcup \neg \exists \text{le}. \text{Lane} \sqcup (\exists \text{cdc}. \text{GoingUp} \sqcap \exists \text{cdc}. \text{GoingDown}) \quad (13)$$

$$\text{OnTwoWayRoad} \sqsubseteq \exists \text{cdc}. \text{GoingUp} \sqcap \exists \text{cdc}. \text{GoingDown}. \quad (14)$$

Sensor Model: Concepts are added to represent all probabilistic inputs from sensors: *SensedOnOneWayRoad*, *SensedOnTwoWayRoad* and *SensedDivider*, that can be either *SensedDashedDivider* or *SensedSolidDivider*. The confusion tables (Tables 1(a)–1(c)) show joint probabilities of an event and its detection by a sensor, and those conditional probabilities are formulated as a set of axioms:

$$P(\text{DashedDivider} | \text{SensedDashedDivider}) = 0.93$$

$$P(\text{SolidDivider} | \text{SensedDashedDivider}) = 0.07$$

$$P(\text{DashedDivider} | \text{SensedSolidDivider}) = 0.20$$

$$P(\text{SolidDivider} | \text{SensedSolidDivider}) = 0.80.$$

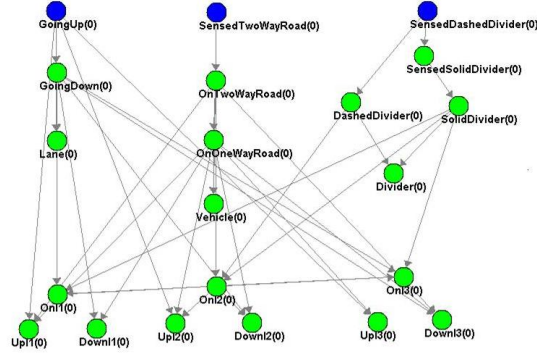


Fig. 3. Bayesian Network representing a traffic domain.

6. Coding and running the scenario

The formalisation presented in the previous section is within the basic definitions of $CRALC$. However, the original role hierarchies are not within the scope of ALC (and, consequently, not within $CRALC$). Therefore, we could not represent directly Formulae 8 and 9. Instead, the spatial information about the domain was implicit in the definitions of each of the lanes and their possible directions. This information was included by grounding the descriptions $\exists eq.Lane$ (“there is a vehicle in the lane”) to new concepts $On1$, $On2$ and $On3$ (“the vehicle is *on* the lane li , for $i \in 1, 2, 3$ ”). An analogous idea was used with respect to the lane directions, where $Upli$ and $Downli$ (for $i \in 1, 2, 3$) were used to represent that the lane li is a going up (resp. down) lane. By merging the roles taken on by the individuals $l1$, $l2$, and $l3$ into concepts $Onli$, $Upli$ and $Downli$, it was possible to represent the Bayesian network for only one individual, the vehicle, and not for $\{l1, l2, l3$ and $\nu\}$. Our solution for representing formulae such as $disjoint(A,B,C)$ was to include probabilistic statements, such as $P(A|B) = 0$ and $P(C|A \vee B) = 0$.

Given the formalisation presented in Section 5 (and the consideration above), the system generated automatically the Bayesian network for only one individual represented in Figure 3, where the nodes in blue are observed variables, i.e. sensors’ states. It is now possible to answer the queries specified in Section 5, which correspond to the following:

1. $\operatorname{argmax}_i P((v : Onli))$, i.e. li is the lane with maximum probability of being the vehicle’s (v) ego-lane .
2. $\forall i : P(li : GoingUp)$, i.e. for each lane li , the probability of being a GoingUp lane.

Consulting the network in Figure 3 for all of the eight possible states of the three sensors, we obtained the answers presented in Tables 2 and 3 for the queries 1 and 2 respectively. In these tables we used the abbreviations $STWR$ for $SensedOnTwoWayRoad$ and SDD for $SensedDashedDivider$. Table 2 shows probable lane on which the vehicle v is driving ($\operatorname{argmax}_{li} P((v, Onli))$), given the evidences, represented on the first three columns. The first line of the table, for instance, represents the state where the sensor obtained GoingDown, vehicle on a *one way road* and a *solid divider*. Given these evidences the node $Onli$ with the highest probability was $On3$. This case is shown in Figure 4(c).

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 | $\text{argmax}_{i_i} P((v : \text{Onli} A))$ |
|---------|------|-------|--|
| GoingUp | STWR | SDD | |
| 0 | 0 | 0 | l3 |
| 0 | 0 | 1 | l1 \vee l2 |
| 0 | 1 | 0 | l2 |
| 0 | 1 | 1 | l3 |
| 1 | 0 | 0 | l1 |
| 1 | 0 | 1 | l2 \vee l3 |
| 1 | 1 | 0 | l1 |
| 1 | 1 | 1 | l2 |




Fig. 4. Examples of three traffic situations, where the vehicle is on a one way road and going down.

Table 3 represents the probabilities of each of the l_i lanes being a GoingDown lane, given the evidences on the first three columns (the probability of GoingUp 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 GoingDown, which is consistent with the evidences GoingDown for the vehicle, and SensedOnOneWayRoad. Similarly for the remainder sensor states represented in the table.

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 | $P(l1:\text{GoingDown} A)$ | $P(l2:\text{GoingDown} A)$ | $P(l3:\text{GoingDown} A)$ |
| 0 | 0 | 0 | 0.99 | 0.99 | 1.00 |
| 0 | 0 | 1 | 0.99 | 0.99 | 1.00 |
| 0 | 1 | 0 | 0.01 | 0.76 | 1.00 |
| 0 | 1 | 1 | 0.01 | 0.95 | 1.00 |
| 1 | 0 | 0 | 0.00 | 0.01 | 0.01 |
| 1 | 0 | 1 | 0.00 | 0.00 | 0.01 |
| 1 | 1 | 0 | 0.00 | 0.61 | 0.99 |
| 1 | 1 | 1 | 0.00 | 0.09 | 0.99 |

7 Conclusion

The representation of QSR systems into description logics is a recent endeavour [12, 30]. The major difficulty of this task is the representation of transitive relations, which are fundamental pieces of spatial knowledge. In particular, [12] presents undecidability results of various \mathcal{ALC} extensions that allow composition-based role inclusion axioms, such as $A \sqsubseteq B \sqcap R_1 \cup \dots \cup R_n$ [30]. Decidability of description logic representations of spatial formalisms were proved in [30] 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.

In this paper we investigated the formalisation of a spatial domain into a probabilistic extension of a basic description logic, $CR\mathcal{ALC}$. In this formalisation we were capable of using the expressivity of a relational formalism (the description logic \mathcal{ALC}), with the treatment of uncertainty provided by Bayesian networks, with which sensor model was encoded. To the best of our knowledge, this paper presented the first principled approach on sensor modelling in a logic language.

This work was successful in showing that the expected queries were consequences of the formalisation of the assumed domain. Given this initial success we conjecture that there is a suitable extension of $CR\mathcal{ALC}$ capable of representing (and reasoning about) spatial domains from any qualitative spatial reasoning system. The development of this formalism is a task of future research

Acknowledgements

This work has been partially supported by Fapesp Project 2008/03995-5 (LogProb). Paulo Santos acknowledges travel 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. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
2. R. Reiter e A. Mackworth. A logical framework for depiction and image interpretation. *Artificial Intelligence*, 41(2):125–155, 1989.
3. T. Matsuyama e V. S. Hwang. *SIGMA: A Knowledge-Based Image Understanding System*. Plenum Press, New York, U.S., 1990.
4. M. Schmidt-Schauss e G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.
5. M. Jaeger. Probabilistic reasoning in terminological logics. In *Principles of Knowledge Representation (KR)*, p. 461–472, 1994.
6. F. Sebastiani. A probabilistic terminological logic for modelling information retrieval. In *International ACM Conference on Research and Development in Information Retrieval*, p. 122–130, Dublin, Ireland, 1994. Springer-Verlag.
7. F. Bacchus, A. Grove, J. Y. Halpern, e D. Koller. From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87:75–143, 1996.
8. A. U. Frank. Qualitative spatial reasoning: Cardinal directions as an example. *International Journal of Geographical Information Science*, 10(3):269–290, 1996.
9. G. Ligozat. Reasoning about cardinal directions. *J. Vis. Lang. Comput.*, 9(1):23–44, 1998.

10. J. Fernyhough, A. G. Cohn, e D. C. Hogg. Constructing qualitative event models automatically from video input. *Image and Vision Computing*, 18:81–103, 2000.
11. D. Randell, M. Witkowski, e M. Shanahan. From images to bodies: Modeling and exploiting spatial occlusion and motion parallax. In *Proc. of IJCAI*, p. 57–63, Seattle, U.S., 2001.
12. M. Wessel. Obstacles on the way to qualitative spatial reasoning with description logics: Some undecidability results. In *Description Logics*, volume 49 of *CEUR Workshop Proceedings*, 2001.
13. A. G. Cohn e S. M. Hazarika. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 46(1-2):1–29, 2001.
14. S. M. Hazarika e A. G. Cohn. Abducing qualitative spatio-temporal histories from partial observations. In *Proc. of KR*, p. 14–25, Toulouse, France, 2002.
15. F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, e P.F. Patel-Schneider. *Description Logic Handbook*. Cambridge University Press, 2002.
16. B. Bennett, A. Cohn, e D. Magee. Enforcing global spatio-temporal consistency to enhance reliability of moving object tracking and classification. *Künstliche Intelligenz*, 2:32–35, 2005.
17. P. C. G. da Costa e K. B. Laskey. Of Klingons and starships: Bayesian logic for the 23rd century. In *Conference on Uncertainty in Artificial Intelligence*, 2005.
18. M. Dürig e T. Studer. Probabilistic ABox reasoning: preliminary results. In *Description Logics*, p. 104–111, 2005.
19. B. Hummel. *Dynamic and Mobile GIS: Investigating Changes in Space and Time*, chapter Map Matching for Vehicle Guidance. CRC Press, 2006.
20. B. Milch e S. Russell. First-order probabilistic languages: into the unknown. In *International Conference on Inductive Logic Programming*, 2007.
21. P. Santos. Reasoning about depth and motion from an observer’s viewpoint. *Spatial Cognition and Computation*, 7(2):133–178, 2007.
22. A. Nikolov, V. Uren, E. Motta, e A. de Roeck. Using the Dempster-Shafer theory of evidence to resolve ABox inconsistencies. In *Workshop on Uncertainty Reasoning for the Semantic Web*, 2007.
23. F. G. Cozman e R. B. Polastro. Loopy propagation in a probabilistic description logic. In *SUM*, p. 120–133, 2008.
24. B. Neumann e R. Möller. On scene interpretation with description logics. *Image Vision Comput.*, 26(1):82–101, 2008.
25. A. G. Cohn e J. Renz. Qualitative spatial representation and reasoning. In *Handbook of Knowledge Representation*, p. 551–596. Elsevier, 2008.
26. C. D’Amato, N. Fanizzi, e T. Lukasiewicz. Tractable reasoning with Bayesian description logics. In *SUM*, p. 146–159, Berlin, Heidelberg, 2008. Springer-Verlag.
27. C. P. de Campos, F. G. Cozman, e 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.
28. F. G. Cozman e R. Polastro. Complexity analysis and variational inference for interpretation-based probabilistic description logics. In *Proc. of the Conference on Uncertainty in Artificial Intelligence*, 2009.
29. B. Hummel. *Description Logic for Intersection Understanding at the Example of Urban Road Intersections*. Phd Thesis, April 2009.
30. M. Cristani e 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*, p. 5–10, 2009.
31. P. Santos, H. Dee, e V. Felon. Qualitative robot localisation using information from cast shadows. In *Proc. of IEEE International Conference on Robotics and Automation*, Kobe, Japan, 2009.