Towards a Situation Awareness Framework Based on Primitive Relations

Norbert Baumgartner [1#], Werner Retschitzegger [2#]
[1] *team Communication Technology Management GmbH*
Goethegasse 3/3, 1010 Vienna, Austria, norbert.baumgarter@te-am.net
[2] *Department of Information Systems (IFS), Johannes Kepler University Linz*
Altenbergerstrasse 69, 4040 Linz, Austria, werner@ifs.uni-linz.ac.at

Abstract

Situation awareness, which aims at determining the meaning of information about perceived objects, is the basis for making decisions in heterogeneous, highly dynamic environments. Recently, ontology-based approaches to situation awareness have been proposed. However, these approaches partly reinvent the wheel, since a common approach to ontology-based situation awareness is missing. Our work focuses on the integral task of such an approach, namely the derivation of relevant relations among the objects of interest. In this paper, we introduce the notion of primitive relations which are inherent in achieving situation awareness. We argue that explicitly deriving these primitive relations could be the foundation of a framework for ontology-based situation awareness which is supposed to reduce the efforts involved in developing situation-aware systems in arbitrary application domains.

1. Introduction

With advances in sensor technologies, the amount of information which has to be incorporated into decision making in heterogeneous, highly dynamic environments steadily increases. The resulting information overload complicates coming to the "best" decision. Situation awareness (SAW) provides the basis for increasing the quality of decisions by determining the meaning of information about the perceived objects. Originating from applications of cognitive sciences to the aviation and military domain, SAW has been defined by Endsley [1] as "the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future".

The process of computationally achieving SAW has been further defined within the International Society for Information Fusion (ISIF) [2]. Especially the Joint Directors of Laboratories' (JDL) Data Fusion Model [3] is a, in this community, well-agreed specification of the processes involved in achieving SAW. The JDL Data Fusion Model divides these processes into four consecutive and at the same time interdependent tasks, denoted as *levels*. Whereas the lower levels 0 and 1 deal with signal respectively object assessment, levels 2 and 3, also commonly referred to as higher-level fusion, address situation and impact assessment. Particularly, level 2, situation assessment achieves SAW by estimating relationships among entities [3]. The importance of finding these relevant relations is emphasized in various related work (e.g. [3], [4]). In particular, Kokar [5] suggests the identification of the types of relations that should be derived as one of the key challenges within the field of SAW. Based on the derivation of these relations, i.e. the determination whether or not a relation holds among two objects, one can aggregate objects to situations.

Recently, the usage of formal *ontologies* [6] for SAW has been motivated (e.g. [3]) what resulted in rather domain-specific approaches to the ontology-based assessment of situations (e.g. [7], [8]). Although using ontologies is beneficial, the lack of a common interpretation of the concepts and tasks involved in achieving SAW causes much efforts, since system designers have to partly reinvent the wheel when developing ontology-based SAW systems.

In the scope of this paper, we focus on an integral task when assessing situations, namely the derivation of relations among objects, from an ontological perspective. Based on a classification of types of relations that contribute to situations, we introduce the essential category of *primitive relations* as well as some of their representatives. Being largely domain-independent, we argue that explicitly deriving these primitive relations could be the foundation of a framework for ontology-based SAW. Amongst others, the main advantage of such a framework would be reduction of efforts involved in developing systems achieving SAW.

Our work is elaborated in cooperation with a prominent Austrian highways agency, since the field of road traffic telematics has all characteristics of a typical SAW application domain (e.g. heterogeneous information, highly dynamic objects, mission- and time-critical decisions). Thus, the adequacy of the proposed approach is illustrated by examples from this application area. Incidently, the concepts of road traffic are easy to understand, as one meets such situations in everyday life.

The remainder of this paper is structured as follows: In Section 2, the mentioned primitive relations as well as some concrete members are introduced based on a classification of relation types. Subsequently, in Section 3, the adequacy of our approach to ontology-based SAW is fortified by the application to the domain of road traffic telematics. In Section 4, related work in the field of ontology-based SAW is discussed. Conclusions and further prospects are presented in Section 5.

2. Towards Primitive Relations

In the following, a classification of relation types, which leads to the category of primitive relations, is introduced. In addition, exemplary representatives of primitive relations are presented.

A. Classification of Relation Types

The classification mentioned above is based on two characteristics of relation types which are described as well as motivated in the following.

- *Domain-dependence*—this characteristic is based on the philosophical distinction between formal and material types of relations [9]. Formal relations have a very low domain-dependence, whereas material relations incorporate domain-dependent knowledge. Apart from the fact that specifying the derivation of highly domain-dependent relation types can not be done without extensive knowledge of the domain, another property is interesting. Relations with a high dependence on the domain implicitly use less domain-dependent relations in the course of their derivation, i.e. they could be decomposed into these less domain-dependent relations. For example, a domain-dependent relation `obstructs` determines whether an object on the road network (e.g. lost cargo) obstructs another object (e.g. a car). This relation could, amongst others, implicitly be derived using a rather domain-independent, spatial relation `before` in order to determine whether an object obstructs another object.
- *Focus*—we introduce this characteristic in order to further define the *relevant* types of relations. Remembering that achieving SAW depends on deriving relations among objects, we argue that relation types with a usually large number of instances are not as valuable as their counterparts with few instances—dealing just with relation types with lots of instances would cause difficulties during the aggregation of objects to *meaningful* situations. We call this property, i.e. the number of relations that *usually* hold among arbitrary objects, the *focus* of a relation type and apply it on a scale from low to high.

Figure 1 depicts a classification of relation types regarding the two characteristics domain-dependence (x-axis) and focus (y-axis). The resulting two-dimensional space is separated into four quadrants. This separation is motivated by the growth-share matrix [10] from the field of marketing and business development. According to the two dimensions, each quadrant can be associated with certain characteristics that are valid for the contained types of relations.

The upper right quadrant contains *situational relations*, that is, types of relations with a high domain-dependence and a high focus. The designation *situational* has been chosen, since adhering to the usual top-down approach, these relation types determine situations from a system designer's perspective. Motivated by the work of Gangemi et. al. [9], its lower left counterpart are *primitive relations* which have a low domain-dependence and a low focus. The other two quadrants fulfill the same role as their counterparts stars and dogs in the growth-share matrix. The very useful *leading relations* (stars) have a low domain-dependence, i.e. they may be applied in a generic way. Nevertheless, leading relations are *usually* highly focused, i.e. they describe very specific circumstances. An example for a leading relation is the relation type `isDuplicateOf`. *Nominal relations* (nominal in the sense of nominal members) are highly domain-dependent and have a low focus; remembering that relations are used for identifying situations, nominal relations are—separately regarded—mainly irrelevant for achieving SAW unless implicitly used by a situational relation. For example, the relation type `isOnDifferentLane` which holds among all percept vehicles that are on different lanes is certainly domain-dependent and has a very low focus, i.e. it is a nominal relation. However, it could be implicitly used by situational relations (e.g. `overtakes`). Providing the largest reuse potential, just primitive and situational relations are hence further on considered. Leading relations, although certainly valuable, are beyond the scope of this paper as we regard them to be very rare and difficult to derive.

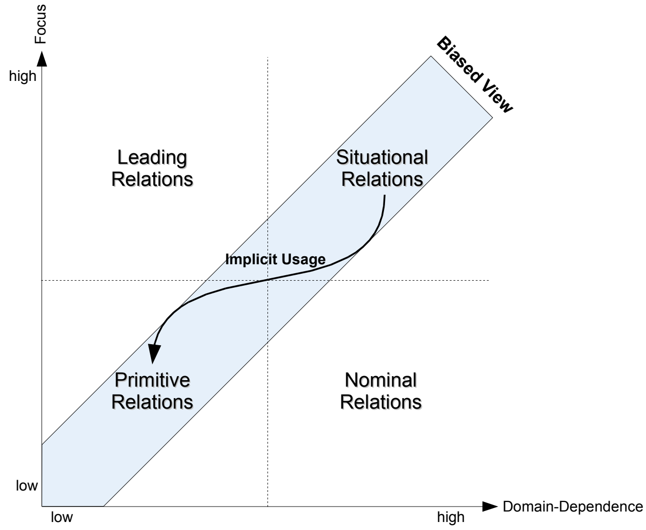

Fig. 1: Classification of Relations

B. Primitive Relations

The classification of relation types in figure 1 indicates two relationships between situational and primitive relations. First, the arrow that connects the two quadrants depicts the implicit usage of primitive relations when deriving situational relations. The non-linear course of this arrow indicates that this relationship between situational and primitive relations is not always strictly hierarchic regarding their domain-dependence. That is, there exists no clear cut hierarchy regarding the mutual usage of relations during their derivation. The second relationship between situational relations and primitive relations follows from the first one. Specific situational relations have a high focus, they pinpoint the situations they support by making use of just the primitive relations they need. This decomposition of situational relations into primitive relations is analog to backward chaining in logic programming [11]. However, this *biased* use of primitive relations, i.e. the restriction to specific situational relations may be too deterministic in most application domains of SAW. It is proposed that the bottom-up, explicit usage of primitive relations when achieving SAW should overcome this limitation.

The practical advantages of incorporating primitive relations into a framework for SAW are threefold. First, one may develop to some degree domain-independent as well as optimized relation derivation algorithms which can be reused in a specific domain. Second, situational relations can be derived by explicitly using existing primitive relations. That is, one may abstract from the details of space, time, etc. and concentrate on the specifics of the to-be-derived situational relations. Finally, the strictly top-down approach, which leads to the deterministic view of relations that contribute to situations, may be levered. That is, by rating relation types according to the degree of their contribution to a situation, also exceptional cases could be dealt with.

In the following, the notion of primitive relations is further illustrated using exemplary members. Therefore, we introduce the notion of *families* of relation types, i.e. relation types which associate objects by examining the same object properties. In fact, just families of primitive relations are described in the following. These families are, apart from being influenced by SAW-specific work, largely based on ontological discussions about formal relations ([9], [12]).

Recapitulating the definition of SAW by Endsley [1], the characteristics that objects are examined "...within a volume of time and space..." leads to the first two kinds of primitive relations: *spatial* and *temporal* relations. A prominent example for a family of *spatial* relations, that are appropriate for regions as primitives in space, is the region connection calculus with eight relations (RCC-8) by Cohn [13]. The relations of RCC-8 are `DC` (disconnected), `EC` (externally connected), `PO` (partly overlapping), `EQ` (equal), `TPP` (tangential proper part) with its inverse `TPPi`, and `NTPP` (non-tangential proper part) with its inverse `NTPPi`. Figure 2 [13] depicts these disjoint relations over the regions a and b as well as transitions between them. However, although RCC-8 is appropriate for the spatial primitive "region", dependent on the application domain, additional relation types for different primitives (e.g. point) should be considered (cf. [13] for an overview of approaches).

As mentioned above, *temporal* relations are also inherent to SAW. As with spatial relations, there are different primitives when dealing with temporal relations. In short, one may distinguish between theories that are based on time points or time intervals [14]. An example for a prominent family of primitive relations over time intervals is Allen's [15] time intervals algebra (e.g. `before`, `after`, `during`).

Comparing the characteristics of spatial and temporal relations, spatial relations can be regarded to be more domain-dependent—although object properties, which define the location respectively extent of an object, are specified similarly across domains, the underlying topologies differ, whereas the topology of time is well-known and therefore more independent from the domain. Furthermore, the focus of the members of both families is very low, since all spatial relations among all objects without regarding the temporal dimension (as well as the other way round) have to be considered.

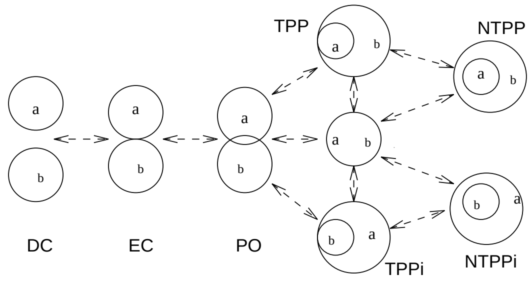

Fig. 2: RCC-8 relations [13]

Going on in the search for primitive relations, one may consider current approaches to SAW. In this respect, *causality* is evident (e.g. bayesian belief networks [4] model dependencies among objects). Hence, the determination whether or not objects influence each other is fundamental to achieving SAW. Unfortunately, causal relations (e.g. `causes` and its inverse `causedBy`) are much more domain-dependent than their spatio-temporal counterparts. For determining whether a relation `causes` holds among two objects, at least some domain-dependent knowledge is necessary. Furthermore, causal relations are also more focused than spatial or temporal relations. However, we argue that causal relations are partly decomposed to spatio-temporal relations (e.g. in order to `cause` a traffic jam, an accident has to occur `before` or `at the same time` as the traffic jam, cf. [16]). Being aware of this circumstance, causal relations are regarded to be primitive relations. There are still further candidates for primitive relations. First, the family of primitive relations dealing with the *composition* of objects, e.g. the relations `isPartOf` and its inverse `consistsOf`, are inherent to any domain. From a philosophical point of view, the composition or parthood of objects is related with the field of mereology, which itself is akin to topology. Examining the introduced spatio-temporal primitive relations, it is evident that they are based on topologies of time respectively space. For example, the relation `isPartOf` is related to the RCC-8 relations `TPP` and `NTPP`. Nevertheless, the composition of objects can be regarded to be more universal than spatial or temporal subsumption, since time and space are not always relevant (e.g. there are entities that persist beyond space and time). Thus, relation types dealing with composition are also regarded to be primitive relations.

Finally, a further interesting family of primitive relations is based on *type inheritance*. That is, objects may be instances of the same object type respectively are in the same branch of the type hierarchy. For example, the relation types `hasSameType`, `isSpecializationOf` and its inverse isGeneralizationOf are proposed for appropriate members of this family.

Figure 3 depicts the suggested classification of the presented families of primitive relations. It reflects both, the focus and the domain-dependence indicated by the positions of the primitive relations and the interdependencies as discussed above, shown by arrows. Note that these arrows are just rather typical examples which should illustrate the concept. Furthermore, placeholders for relation types as well as exemplary situational relations (e.g. an accident decelerates a vehicle by obstructing it) are shown. The bottom-up, *explicit* usage of primitive relations widens the scope of SAW, since not just the implicitly used primitive relations contribute to situations.

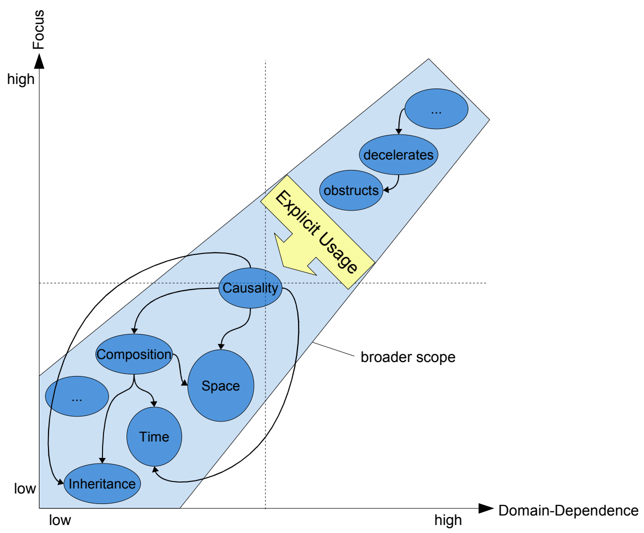

Fig. 3: Relationships and classification of primitive relations

3. Proof of Concept

In this section, our approach regarding the explicit use of primitive relations is illustrated using an example from the field of road traffic telematics. Thereby, the proposed advantages are fortified through real-world use cases.

The modeled concepts of the domain focus on typical traffic information for motorists. Theses are `vehicle` as well as `incident` with its derived types `accident` and `traffic jam`. There are several relation types which should be derived in the course of determining the situations of interest. The first one is `causes` which is a directed relation type among accidents and traffic jams that occur at adjacent spatio-temporal locations. Furthermore, an accident may `involve` vehicles, whereas a traffic jam may `obstruct` the following vehicles. Certainly, the relation types `involves` and `obstructs` may be classified as situational relations, whereas `causes`, because of its domain-independent meaning, is in the gray area between situational and primitive relations.

An exemplary scenario is depicted in Figure 4. An accident, involving a vehicle, causes a traffic jam that obstructs a currently approaching vehicle.

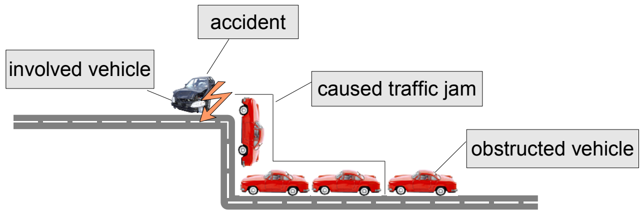

Fig. 4: Exemplary scenario

In the following, the advantages stated in the previous section are fortified according to this example. First, one can imagine optimized relation derivation algorithms that are also applicable to the field of road traffic telematics. Usually, each object is projected onto the road network. The underlying spatial topology of road networks are graphs, thus, common graph algorithms (e.g. wayfinding) are adequate as default implementations for deriving spatial relations.

Second, the explicit derivation of primitive relations allows a designer of a SAW system to abstract from details of space, time, etc. For example, examining the relation type `obstructs`, may—apart from the specification of the object types `incident` and `vehicle`—be decomposed into spatio-temporal primitive relations. Figure 5 shows the explicitly used primitive relations `PO` (a region forming the boundary of the traffic jam partly overlaps another region marking off the approaching vehicle) and `contains` (the time interval the traffic jam lasts contains the time interval the vehicle approaches). That is, it is not necessary to go into the details of spatio-temporal representation. Actually, the underlying topology could be changed without affecting the system designers composition.

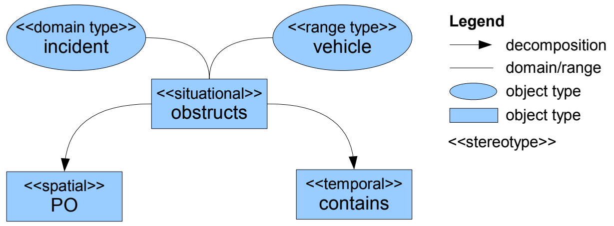

Fig. 5: Specification of `obstructs`

Finally, this explicit specification induces a too deterministic view of situations. That is, by aggregating only objects that obstruct, involve, or cause each other to situations would result in insufficient results. For instance, the approaching vehicle in Figure 4 does not overlap the traffic jam, thus, the relation `obstructs` would not be derived and the objects would not be be regarded to be in a situational context. By deriving all primitive relations in a bottom-up manner, this deterministic approach could be levered. For example, rating the relevance of primitive relations, a system designer could explicitly determine the primitive relations that contribute to situations respectively are derived. Being noticed to be relevant in our case, the relation type `EC` (externally connected) could be additionally derived (without being used by a situational relation). Consequently, the vehicle and the traffic jam could be aggregated to a situation. In the end, system development could be reduced to a mere parametrization of the existing default implementations of primitive relations (e.g. rating of relevant relation types).

4. Related Work

In this section, related work is discussed and the contribution of this paper is elaborated. To the best of our knowledge, the only framework for ontology-based situation assessment has been elaborated by Matheus et. al. [17]. They developed SAWA—an assistant for higher-level fusion and SAW—which uses an upper ontology for situation assessment. With regards to relation types, Matheus et. al. motivate the category of *standing relations* which are high-level relation types that represent the goal of the situation assessment process. In fact, standing relations can be regarded as highly focused as well as domain-dependent situational relations that determine the set of derived relations. In contrast, we explicitly derive primitive relations in a bottom-up fashion; thereby, the biased view arising from the top-down approach of standing relations is prevented. Furthermore, the proposed incorporation of primitive relations reduces the efforts of a system designer when specifying situational relations.

5. Conclusions

In this paper, we proposed a classification of relations that contribute to SAW. We argued that one of the resulting categories, primitive relations, may provide the fundament for developing a framework for ontology-based SAW. The adequacy of explicitly used primitive relations and their introduced exemplary members have been fortified by their application to the domain of road traffic telematics. Thereby, our claim that such a framework reduces the efforts involved in developing concrete systems for achieving SAW has been illustrated.

The suggested primitive relations result from a mainly ontological view of the problem. However, the number of relations to be derived is likely to grow exponentially what endangers the algorithmic feasibility, i.e. the relevant primitive relations have to be chosen carefully. In future work, we aim at elaborating characteristics of appropriate, computationally feasible primitive relations. Furthermore, our work focuses on developing an upper ontology for SAW, which—amongst other concepts—employs primitive relations. In the long term, we are going to apply our findings to the domain of road traffic telematics, in order to support traffic operators achieving situation awareness.

References

[1] M.R. Endsley. *Situation Awareness Analysis and Measurement*, chapter Theoretical Underpinnings of Situation Awareness: A Critical Review, pages 3–33. Lawrence Erlbaum Associates, New Jersey, USA, 2000.

[2] International Society of Information Fusion. `http://www.isif.org`.

[3] James Llinas, Christopher Bowman, Galina Rogova, and Alan Steinberg. Revisiting the JDL data fusion model II. In *Proceedings of the Seventh International Conference on Information Fusion, Stockholm, Sweden*, pages 1218–1230, June 2004.

[4] Vladimir Gorodetsky, Oleg Karasaev, and Vladimir Samoilov. On-Line Update of Situation Assessment: A Generic Approach. *International Journal of Knowledge-Based and Intelligent Engineering Systems*, 9(4):361–365, November 2005.

[5] M. M. Kokar. Situation Awareness: Issues and Challenges. In *Proceedings of the Seventh International Conference on Information Fusion, Stockholm, Sweden*, pages 533–534, June 2004.

[6] Thomas R. Gruber. A Translation Approach to Portable Ontology Specification. *Knowledge Acquisition*, 5(2):199–220, June 1993.

[7] Anne-Claire Boury-Brisset. Ontology-based approach for information fusion. In *Proceedings of the Sixth Conference on Information Fusion, Cairns, Queensland, Australia*, July 2003.

[8] Paul R. Smart, Nigel R. Shadbolt, Leslie A. Carr, and Monica C. Schraefel. Knowledge-based information fusion for improved situational awareness. In *Proceedings of the 8th International Conference on Information Fusion 2005, Philadelphia, USA*, July 2005.

[9] A. Gangemi, N. Guarino, C. Masolo, and A. Oltramari. Understanding top-level ontological distinctions. In *In Proceedings of the Seventeenth International Joint Conference on Articficial Intelligence (IJCAI-01) Workshop on Ontologies and Information Sharing, Seattle, Washington, USA*, 2001.

[10] The Boston Consulting Group. The Growth-Share Matrix. `http://www.bcg.com/this_is_bcg/mission/growth_share_matrix.html`, 1970.

[11] William F. Clocksin and Christopher S. Mellish. *Programming in Prolog*. Number ISBN 3-540-00678-8. Springer-Verlag, 5th edition, August 2003.

[12] Barry Smith and Pierre Grenon. The cornucopia of formal-ontological relations. *Dialectica*, 58:3:279–296, 2004.

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

[14] Ubbo Visser. *Intelligent Information Integration for the Semantic Web*. Springer-Verlag, Berlin, Germany, 2004.

[15] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.

[16] John F. Sowa. *Knowledge Representation — Logical, Philosophical, and Computational Foundations*. Brooks/Cole, Pacific Grove, California, USA, 2000.

[17] C. Matheus, M. Kokar, K. Baclawski, J. Letkowski, C. Call, M. Hinman, J. Salerno, and D. Boulware. SAWA: An assistant for higher-level fusion and situation awareness. In Belur V. Dasarathy, editor, *Proceedings of SPIE Conference on Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2005, Orlando, Florida, USA*, volume 5813, pages 75–85, March 2005.