Introduction to the special issue

Inderjeet Mani — Philippe Muller

1. Introduction

References to temporal and spatial locations are ubiquitous in natural language communication. Speakers position events in time in terms of relationships in a calendar, and they describe the spatial positions of entities in terms of relations with other known entities, or within an accepted frame of reference (contextual, or geographical). In both cases, speakers can take their positions in time and space into account. Temporal descriptions in natural language tell us when events happened, and in what order, how long they last, and their distance in time from the speech time. Spatial descriptions tell us where entities are located, how far away they are, their spatial configurations with respect to each other, and their size and shape. This wealth of spatial and temporal information in natural language has stimulated a great deal of interest in leveraging such information in practical applications. For example, in information retrieval, there has been considerable emphasis on geo-location (Cheng *et al.*, 2010; Roller *et al.*, 2012), as well as using temporal information for improved search and browsing (Alonso *et al.*, 2011). In natural language question-answering, users can receive answers to questions such as when an event occurred, or which events occurred prior to a particular event (Saquete *et al.*, 2009). Systems also can rank answers by considering places mentioned in them that are associated with places in the question, as in IBM's DeepQA [1]. In information extraction, systems can extract and plot the geographical distributions of events such as natural disasters, weather, and disease outbreaks from news articles and display them on interactive maps, e.g., Health Map from ProMed [2]. Such research relies in part on accurate disambiguation and geo-location of place-names (e.g., Garbin and Mani, 2005; Mani *et al.*, 2010; Roller *et al.*, 2012). In medical information extraction from clinical narrative reports, Hripcsak *et al.* (2009) have mined duration assertions (like "three weeks ago") and determined their uncertainty by correlating them with facts from a structured clinical database. In summarization within an information retrieval setting, Allan *et al.* (2001) have explored temporal summaries of news topics, and Aker *et al.*

---

[1] `http://www-03.ibm.com/innovation/us/watson/building-watson/how-watson-works.html`.

[2] `http://healthmap.org/promed/`

(2012) have examined summarization of documents relevant to geo-tagged images for improved indexing in image retrieval. In generating weather forecasts, spatial and temporal information in structured data has been mapped to natural language descriptions (Turner *et al.*, 2010). For navigation, Vogel and Jurafsky (2010) have trained a system to follow navigational spatial directions found in an annotated corpus, the HCRC Map Task Corpus (Anderson *et al.*, 1991). In computer graphics and scene generation, Coyne and Sproat (2001) have taken spatial configurations in text and rendered them in graphical displays; Åkerberg *et al.* (2003) take natural language accident reports and process temporal and spatial information in them to generate animations. The spate of recent work in annotation and evaluation in this area has helped to move from essentially theoretical models in the 1990s to arrive at research that has a more empirically justified foundation. The objective of this special issue is to present new developments in the processing of temporal and spatial information in language, from theoretical, practical and methodological points of view.

2. Linguistic Challenges

To address the needs of applications such as the ones above, the sheer variety of spatial and temporal expressions in natural language is something that has to be confronted. Dates and times, explicit temporal relations (e.g., “before”, “simultaneous”, etc.), and tense and aspect information have to be analyzed in the face of considerable variation in the manner in which they are expressed in different languages. For example, some languages lack tense altogether and rely instead on modals to distinguish between “real” events and others, while others have weak tense systems but rich aspect marking (Comrie, 1985; Mani *et al.*, 2005). Events in narrative can be described in their order of occurrence, but more often, the text order is different from the underlying order. Rules for constructing partial orderings of events given a natural language discourse have been the subject of much research. In Discourse Representation Theory (DRT) (Kamp and Reyle, 1993), rules that take into count tense and aspect have been developed within a compositional semantics that include, as defaults, the narrative ordering rule (past tense events succeed each other) and the stative rule (statives overlap with the event in the adjacent clause). In turn, spatial information is expressed by a variety of linguistic devices that vary greatly across languages, and can express both topological relations (“inside”, “at”, “along”, etc.) and orientation relations (“to the left”, “above”, “east”, etc.), size, shape, etc. These orientation relations involve multiple coordinate systems (or frames of reference) that are distributed unevenly across languages (Levinson, 2003). In the absolute frame of reference, the relation between two spatial entities (figure and ground) involves a coordinate system anchored to fixed bearings that is centered on the ground entity (e.g., “Iran is west of India”). In the intrinsic frame, the coordinate system is centered on the ground object, viz., “the child is in front of the TV”. In the relative frame, e.g., “the lamp is in front of the TV”, the relation between figure and ground is seen from the point of view of a coordinate system centered on the viewer, though a second system centered on the ground entity may also be present. There is thus a considerable challenge in distinguishing, from natural language descriptions, the frames involved, given the resources available to express them in the particular language, and a further challenge in representing their semantics. Language does more than simply describe static spatial configurations and temporal orderings of events. Descriptions of moving objects are common in language, and this involves, from a processing standpoint, identifying the event, the moving object, the path traversed through space (including the medium), the temporal characteristics, and the manner of movement. In turn, extracting such information requires taking into account the two distinct ways in which languages express concepts of motion: verb-framing and satellite-framing (Talmy, 1985; Talmy, 2000). In verb-framing, found in French and other Romance languages, as well as Semitic languages, Turkish, etc., the verb has a morpheme that conflates the path of motion, whereas the manner is optionally expressed by adjuncts. In satellite-framing, found in English and other Germanic languages as well as Slavic languages, the main verb conflates the manner or cause of motion, whereas the path is expressed by "satellites" involving particles and adjuncts. Natural language processing systems must be able to unpack language-specific path and manner information within a compositional semantics, in order to construct a proper characterization of the motion. Asher and Sablayrolles (1995) and Muller (1998) offer systematic classifications of motion verbs in terms of a lexical semantics that hinges on particular spatiotemporal primitives. In addition to representing specific types of temporal and spatial relationships, systems must address the context-dependence, vagueness and ambiguity inherent in the use of natural language. Most temporal and spatial references depend on the prior discourse context. References to seasons or expressions of proximity ("near", "almost") involve fuzzy boundaries and implicit scales. Some linguistic information may be implicit or absent (for example, many events described in a narrative may not have times explicitly associated with them), and often the information is present but ambiguous; for example, there is a need to disambiguate between different referents for a place name e.g., "Victoria"), to distinguish spatial versus temporal senses of prepositions like "before", to determine whether "on the left" means the relative frame (the viewer's left) or the intrinsic frame (the object's left, assuming it has one). The interdependence between temporal interpretation and discourse structure has long been noted (Hitzeman *et al.*, 1995; Asher and Lascarides, 1993).

3. Temporal and Spatial Ontology

3.1. *Introduction*

The representations that natural language maps to must be precise enough to support the inferential needs of the application. The temporal and spatial information from language must be represented in terms of entities such as events, times, temporal relations, as well as places, paths, and spatial relations. It is worth bearing in mind, however, that many of the distinctions made in language, such as between figure and ground, are more or less ignored in formalizations, which treats relations between pairs of entities that are interchangeable. Similarly, parthood is often intransitive in natural language, topological relations (like "on") involve support relations, and orientation relations like "above" or "in front of" can depend on the typical use of the ground object (Herskovits, 1986). These functional notions are largely ignored in formal ontologies. Given that ontologies for space and time have been the focus of much research in logic and AI, a survey of these is not possible in this short introduction; see the References for more details. Instead, we provide an overview of some of the key issues in representation, commenting on the adequacy and expressiveness of the representation from a natural language standpoint.

3.2. *Time*

In the case of time, one may choose to treat intervals as primitive, as in the interval calculus of Allen (1983). There, 13 qualitative relations between intervals are defined in terms of a single relation of contiguity, called "meets", between one time interval and another. The interval calculus has been formalized in terms of first order logic (Van Benthem, 1983; Allen and Hayes, 1985); the latter axiomatization specifies that the meets (i.e., meetings) of intervals are unique, that pairs of intervals meet in linear order, that the union of every pair of intervals exist, and that time is infinite while intervals are finite. Other models, which treat time as finite, or bounded, in the past, future, or both, are of course clearly possible. The interval calculus has been extremely popular in AI and especially so for NLP. Whether these 13 relations are sufficient to address the varieties of temporal relations expressed in natural languages by expressions such as "after", "during", "while", etc. is an open question, but there is so far every reason to believe that no language makes a distinction that is not expressible using the underlying set of 213 different disjunctions of possible temporal relations. It is more likely that the calculus is too fine-grained; for example, "meet" may not typically be part of commonsense reasoning from language (Mani *et al.*, 2003). Alternatively, one might start with intervals as sets of points, i.e., instants, so that, say, an interval is any convex set of instants. In that case, one has to face the criticism that points are not central to commonsense inference, especially from natural language. From a technical standpoint, it also means contending with some representational messiness, in terms of deciding whether an instant is the end-point of one interval or the start-point of another. A more ontologically agnostic approach is found in Galton (1990), where instants both limit intervals and are (properly) within intervals, while refraining from committing to whether intervals are constituted of instants. In addition to the instant/interval representation and temporal boundedness, other key issues for time have to do with whether it is discrete or continuous, and whether it branches, in the past or future. For these latter issues, natural language does not offer any clear prescriptions; it is worth noting that branching pasts as well as futures are quite common in hypothetical reasoning in natural language, as in the example of "regret" and "hope". Logics for branching time include the system of McDermott (1982), and Computational Tree Logic (Huth and Ryan, 2004). Last but not least, representations of time have to be able to deal with tense distinctions in natural language, distinguishing, say, between past perfect and simple past tense. Two approaches have been taken to representing tense: (i) as a set of constraints involving relationships among speech times, event times, and reference times (Reichenbach, 1947), where the times in question are instants; (ii) as an operator in a logic that shifts the truth conditions of propositions (for example, [Past X] is true at a time instant t if there is a time instant t' prior to t such that X is true at time t') (Prior, 1967). While Reichenbach's scheme has been applied to numerous languages, Prior's tense logic is more computational in nature. However, the latter allows free iteration of tense operators that results in overgeneration, expressing tenses that cannot exist in any language, while also failing to distinguish between present perfect and simple past tense. For more on logics of time, see Van Benthem (1983).

3.3. *Space*

The choices among primitive representations for space parallel the distinctions made for time, but there are specific differences from the temporal domain, the most salient of which is the increase in dimensionality. If regions are considered to be sets of points, two regions can be said to be connected if the intersection of their closures is non-null. The possible relations between the interior, exterior, and boundary of one region and those of the other region (where the regions in question are 2D, one-piece, without holes), give rise to the 9 relations of the 9-Intersection Calculus (9IC) (Egenhofer, 1991). Alternatively, if regions are taken as primitive, a primitive notion of connectedness can in turn be used to define relations of parthood, equality, and overlap between regions, giving rise to the 8 underlying relations of the Region Connection Calculus (RCC-8) (Randell *et al.*, 1992). These relations can be mapped to topological relations in natural language, expressed by senses of prepositions such as "in", "on", etc, if functional factors mentioned above are left aside. Another spatial distinction (a spatial analogue, in a sense, of tense) has to do with orientation, as expressed by "to the left", "above", etc. Here, spatial calculi have considered primitive entities such as regions and points, as well as lines. For example, in the Cardinal Direction Calculus (Skiadopoulos and Koubarakis, 2005), which can be used to model the absolute frame of reference, there are 511 relations between regions laid out contiguously in a grid. For the intrinsic frame, the Dipole Calculus (Moratz *et al.*, 2011) considers 24 underlying relations between points and oriented lines. For modeling the relative frame, the Double Cross Calculus (Freksa, 1992) considers 17 relations between figure, ground and viewer, each considered as points. As indicated by the varying numbers of relations above, the formal representation is sometimes too coarse-grained to distinguish natural language meanings, while in some cases the formal representation makes distinctions that are not needed for representing natural language. For background on spatial logics, see Aiello *et al.* (2007), and more generally on qualitative representations for space, see Cohn and Renz (2007).

3.4. *Events*

Events can be treated as atomic, or can form clusters, or be expressed in terms of hierarchies as part of plans. For example, a terrorist event may have different subclasses, such as a bombing, a plane crash, etc., and an event of manufacturing an airplane can involve thousands of different sub-events. At an atomic level, events involve changes of state, and this dynamic property differentiates them from states per se. Further, while states (such as "knowing") are individuated in that every subinterval of the state has the same properties, events do not have such a strong subinterval property. In AI reasoning formalisms such as the event calculus (Kowalski and Sergot, 1986), events are temporally anchored, in that they happen at a particular time instant, and events also initiate and terminate the properties that hold for an individual in a given state; thus "give" is modeled as a transfer of possession. In planning formalisms, in addition to taking the temporal durations of events into account, events have specific pre-conditions (tests) and post-conditions (updates) that can result from the event's execution. The interval calculus as elaborated in Allen (1984) distinguishes between properties that hold at a time interval, and events that occur during the interval (specifically, the maximal interval during which the event occurs). A nice property of the interval calculus is that in addition to event anchoring, event ordering naturally falls out of the temporal relations between intervals during which the events occur. This makes it again especially relevant to modeling the ordering of events in narrative. However, as pointed out by Galton (1990), the interval calculus, unlike his own hybrid instant-interval ontology, fails to adequately model continuous motion; thus, it fails to distinguish between "John is in San Francisco" and "John is at rest in San Francisco".

3.5. *Paths and Motion*

Paths are crucial for the representation of the spatiotemporal trajectories of motion events as well as references to purely spatial routes (e.g., "the road from Chiang Mai to Sukkothai"). Spatiotemporal paths can be viewed as temporal orderings of spatial configurations; thus each configuration can be modeled in a spatial calculus while the ordering can be modeled in a temporal calculus, as in Galton (2000) which uses RCC-8 for the spatial component. The approach of Dynamic Interval Temporal Logic (Pustejovsky and Moszkowicz, 2011) is similar to Galton's, but views interpreting of the meaning of a motion expression in natural language as creating an executable program. Movement is represented in terms of a basic change-of-location predicate that is used to further define directed movement and path predicates (like "enter"). To address relative motion, some authors have proposed spatio-temporal theories providing either purely topological calculi (Muller, 1998; Bennett *et al.*, 2000) or with the integration of orientation change, such as in the Qualitative Trajectory Calculus of Van de Weghe (2004), where the positions of two objects at different moments in time are compared, using both RCC-8 and the Double-Cross Calculus. By representing different values of the RCC-8 relation of DC, i.e., disconnection among regions, it is possible to model relative motions described by "approach", "follow", etc. For more on qualitative calculi for representing motion and their expressiveness in capturing natural language distinctions, see Mani and Pustejovsky (2012).

4. Issues

4.1. *Granularity*

Ontologies for time and space should ideally be augmented with various mechanisms to support reasoning at different levels of granularity. For example, it is often necessary to carry out temporal arithmetic that compares times that have different granularities (such as weeks and months), or to deal with combinations of different types of representations, e.g., topological and orientation calculi (for prepositions like “over”, which involve both aspects), or mixing of bounding boxes or other regions with points (Papadias *et al.*, 1995). Finally, spatial ontologies require integration with gazetteers and other database resources used in GIS systems. Most of the qualitative frameworks mentioned previously have variants in which the information expressed is underspecified, by conflating predicates with neighbouring semantics (e.g. the relations “starts”, “during” and “finishes” in Allen relation could be conflated in a single “is included” relation).

4.2. *Annotation*

Annotation schemes that represent linguistic information need to be expressive enough to capture the information needed for the task or application. At the same time, they need to be sufficiently lightweight to allow for annotation with high reliability. Ideally, they should be extensible to multiple languages at a relatively low cost. Annotation schemes like TimeML (Pustejovsky *et al.*, 2005), which represents events and times as intervals, with 13 temporal relations between them, the TIMEX2 specification for time expressions (Ferro *et al.*, 2005), and SpatialML (Mani *et al.*, 2010) have been proven to be extensible across languages. Extensions to TimeML and SpatialML to handle vagueness and imprecision have also been explored (Snoussi *et al.*, 2012). While TIMEX2 and SpatialML have been reliably annotated, TimeML in particular is hard to annotate reliably, with agreement among annotators on temporal relations being around 50% (Mani *et al.*, 2006). To address this, more constrained annotation procedures have been considered, such as guiding the event pairs to be considered (Bethard *et al.*, 2012), allowing for more coarse-grained annotation (Bramsen *et al.*, 2006) or focusing on the most reliable relations in existing data (Chambers and Jurafsky, 2008).

4.3. *Data*

The annotation of corpora with time expression tags (either TIMEX2, or TIMEX3 as used in TimeML) has been carried out on a variety of languages, including Arabic, Chinese, English, French, Hindi, Italian, Korean, Persian, Portuguese, Spanish, Swedish, and Thai. The scheme has been applied to news, scheduling dialogs, email, parliamentary meetings, medical narratives, accident reports, fiction, etc. "Timebanks" based on TimeML have been created originally for English (Pustejovsky *et al.*, 2003), but also for Catalan (Saurí and Badia, 2012), French (Bittar *et al.*, 2011), German (Spreyer and Frank, 2008), Italian (Caselli *et al.*, 2011), Korean, Portuguese (Costa and Branco, 2012), Romanian (Forascu and Tufi, 2012), and Spanish (Nieto *et al.*, 2011), among others. Bethard *et al.* (2012) have released their restricted annotation on a set of short stories for children. Annotated corpora for spatial information is still scarcer, but SpatialML corpora have been created for Mandarin Chinese and English (Mani *et al.*, 2010).

4.4. *Evaluation*

While there has been considerable progress in terms of evaluation campaigns at least for temporal tasks, e.g. TempEval (Verhagen *et al.*, 2010), there is no agreed upon methodology for many subtasks in the field. In particular, for event ordering, authors aim at various levels of complexity, either considering all distinct pairs of events (Mani *et al.*, 2006), pairs of successively described events (Mani *et al.*, 2003), same-sentence events (Lapata and Lascarides, 2006) (now also tasks in the Tempeval campaigns). Larger groups of events (and event hierarchies) have generally been avoided, with some notable exceptions (Bramsen *et al.*, 2006; Chambers and Jurafsky, 2008; Denis and Muller, 2011), as the task is much more complex, and because it is harder to fix a reference annotation (in particular, different temporal graphs can implicitly agree because of inference issues). For spatial tasks, there have been only a few evaluation campaigns, including GeoCLEF (Mandl *et al.*, 2008), LogCLEF (Mandl *et al.*, 2009), and SemEval (Kordjamshidi *et al.*, 2012).

4.5. *Inference*

Temporal and spatial relations have semantic properties that humans take into consideration when they annotate. For instance, order and inclusion are transitive, and a temporal or a spatial situation can be fully described concisely with only a few predicates. These properties have, historically, motivated the use of well-studied inference frameworks such as constraint systems based on interval algebras (e.g., the interval calculus, RCC-8, etc.). However, they complicate evaluation, as they require a reference scheme that includes both annotation and inferred relations, and in addition, such inferred relations may bias classical evaluation measures. Some authors have explicitly addressed this problem by proposing tailored evaluation procedures (Tannier and Muller, 2011; UzZaman and Allen, 2011). Finally, the issue of efficiency of inference related to temporal and spatial relations is crucial; for more on this see Van Beek (1992) and Kontchakov *et al.* (2008).

5. Papers

This special issue's papers focus more on spatial problems than temporal ones, perhaps reflecting the more recent interest of NLP researchers on this important field, which is growing in interest motivated in part by the widespread use of GIS and location-aware services. The temporally oriented papers focus on isolating the most important pieces of information, either dates, events or event relations, contrasting with the more common approaches which tend to be exhaustive. Kessler *et al.* aim at building chronologies, summaries of important events, and they detect dates from news corpus, using syntactic patterns and handmade interpretation rules. They also use frequency and informativity measures to estimate their importance, or saliency, as events related to such dates are indexed, and by machine learning trained on manually written chronologies. They evaluate the interest of this approach on French and English data, both by comparing the resulting chronologies to manually built ones and also by human judgments of the selected dates. Good inter-annotator agreement on the latter seems to validate the usefulness of the notion of salient dates. Marsic focuses on the difficult task of selecting relevant relations between events, a heavily biased problem that has resisted recent efforts, since most event pairs bear no explicit relation in a document. There is also low inter-annotator agreement on this aspect, and Marsic aims at improving the results at least on sentence-internal event pairs, since they are the most numerous. She lists a number of syntactic rules that can be used to determine the relation between two events in that context, implements them to obtain high accuracies on the TimeBank corpus, and also advocates for similarly more detailed guidelines for annotators.

The spatially oriented papers cover diverse issues, from extraction of spatial information (places descriptions) to geolocalisation for navigation or semantic interpretation. Pustejovsky *et al.*'s paper builds on the TimeML and SpatialML specifications, integrating them into a new framework, ISO-Space 1.4. It represents qualitative spatial relations and distances between locations, as well as paths of objects in motion. The annotation of data using the specification, when carried out, can allow for richer integration of language annotation and spatial reasoning than has hitherto been possible. Moot's paper analyzes motion and path expressions in 19th Century French travelogues. His system uses a compositional semantics framework that subscribes to the temporal ordering framework of DRT, with a spatial semantics that allows for RCC-8 relations among atomic regions (paths here are construed as regions) as well as more complex Boolean combinations of regions, as in Kontchakov *et al.* (2008). His lexical semantics for motion verbs leads to a revision of the classification of Asher and Sablayrolles (1995). Gaio *et al.* focused on the extraction of spatial descriptions that can be mapped to locations in a geographical ontology, for better document indexing.

Lexical and syntactic pattern matching is used to feed a semantic process where spatial relations are interpreted to place locations with respect to known landmarks. Motion descriptions also help to isolate the relevant pieces of information in their corpus of travel journals, the same used by Moot. Blaylock *et al.* describe a dialogue system for navigation which maps street-level natural language path descriptions to synchronized GPS tracks. It is based on their PURSUIT corpus of annotated geospatial path descriptions, acquired from a task where cars with microphones and GPS follow one another. Their approach integrates natural language understanding with path finding methods such as graph search and particle filtering.

Acknowledgments

We thank the journal editors Éric de la Clergerie, Béatrice Daille, Yves Lepage, Jean-Luc Minel and François Yvon, as well as the following reviewers, for their considerable efforts in making this special issue possible: Salah Ait-Mokhtar, Nicholas Asher, Jason Baldridge, John Bateman, Delphine Battistelli, Nate Blaylock, Harry Bunt, Pascal Denis, Leon Derczynski, Christine Doran, Patrice Enjalbert, Michel Gagnon, Mauro Gaio, Robert Gaizauskas, Laurent Gosselin, Caroline Hagège, Guy Lapalme, Gérard Ligozat, Denis Maurel, Richard Moot, Patrick Paroubek, Laurent Prévot, James Pustejovsky, Patrick Saint-Dizier, Xavier Tannier, Isabelle Tellier, Marc Verhagen, and Annie Zaenen.

6. References

Aiello M., Pratt-Hartmann I., Van Benthem J. (eds), *Handbook of Spatial Logics*, Springer, 2007.

Aker A., Fan X., Sanderson M., Gaizauskas R. J., "Investigating Summarization Techniques for Geo-Tagged Image Indexing", *Proceedings of the International Conference on European Conference on Information Retrieval (ECIR)*, p. 472-475, 2012.

Åkerberg O., Svensson H., Schulz B., Nugues P., "CarSim: an Automatic 3D Text-to-Scene Conversion System Applied to Road Accident Reports", *Proceedings of the 10th Conference on European Chapter of the Association for Computational Linguistics*, Budapest, Hungary, p. 191-194, 2003.

Allan J., Gupta R., Khandelwal V., "Temporal Summaries of News Topics", *SIGIR*, p. 10-18, 2001.

Allen J., "Maintaining Knowledge About Temporal Intervals", *Communications of the ACM*, vol. 26, no. 11, p. 832-843, 1983.

Allen J., "Towards a General Theory of Action and Time", *Artificial Intelligence*, vol. 23, p. 123-154, 1984.

Allen J., Hayes P., "A Commonsense Theory of Time", *Proceedings of IJCAI'85*, p. 528-531, 1985.

Alonso O., Gertz M., Baeza-Yates R. A., “Enhancing Document Snippets Using Temporal Information”, *Proceedings of String Processing and Information Retrieval (SPIRE 2011)*, Pisa, Italia, p. 26-31, 2011.

Anderson A., Bader M., Bard E., Boyle E., Doherty G., Garrod S., Isard S., Kowtko J., McAllister J., Miller J. *et al.*, “The HCRC Map Task Corpus”, *Language and Speech*, vol. 34, no. 4, p. 351-366, 1991.

Asher N., Lascarides A., “Temporal Interpretation, Discourse Relations, and Commonsense Entailment”, *Linguistics and Philosophy*, vol. 16, p. 437-493, 1993.

Asher N., Sablayrolles P., “A Typology and Discourse Semantics for Motion Verbs and Spatial PPs in French”, *Journal of Semantics*, vol. 12, no. 1, p. 163-209, 1995.

Bennett B., Cohn A. G., Torrini P., Hazarika S. M., “Describing Rigid Body Motions in a Qualitative Theory of Spatial Regions”, *AAAI/IAAI*, p. 503-509, 2000.

Bethard S., Kolomiyets O., Moens M.-F., “Annotating Story Timelines as Temporal Dependency Structures”, *in* N. Calzolari, K. Choukri, T. Declerck, M. U. Dogan, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis (eds), *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC’12)*, European Language Resources Association (ELRA), Istanbul, Turkey, May, 2012.

Bittar A., Amsili P., Denis P., Danlos L., “French TimeBank: an ISO-TimeML Annotated Reference Corpus”, *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Portland, Oregon, USA, p. 130-134, June, 2011.

Bramsen P., Deshpande P., Lee Y. K., Barzilay R., “Inducing Temporal Graphs”, *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, Sydney, Australia, p. 189-198, July, 2006.

Caselli T., Bartalesi Lenzi V., Sprugnoli R., Pianta E., Prodanof I., “Annotating Events, Temporal Expressions and Relations in Italian: the It-TimeML Experience for the Ita-TimeBank”, *Proceedings of the 5th Linguistic Annotation Workshop*, Association for Computational Linguistics, Portland, Oregon, USA, p. 143-151, June, 2011.

Chambers N., Jurafsky D., “Jointly Combining Implicit Constraints Improves Temporal Ordering”, *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Honolulu, Hawaii, p. 698-706, October, 2008.

Cheng Z., Caverlee J., Lee K., “You Are Where You Tweet: a Content-Based Approach to Geo-locating Twitter Users”, *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, ACM, p. 759-768, 2010.

Cohn A., Renz J., “Qualitative Spatial Reasoning”, *in* V. F. Van Harmelen, V. Lifschitz, B. Porter (eds), *Handbook of Knowledge Representation*, Elsevier Science, 2007.

Comrie B., *Tense*, Cambridge University Press, Cambridge, 1985.

Costa F., Branco A., “TimeBankPT: a TimeML Annotated Corpus of Portuguese”, *in* N. Calzolari, K. Choukri, T. Declerck, M. U. Dogan, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis (eds), *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC’12)*, European Language Resources Association (ELRA), Istanbul, Turkey, May, 2012.

Coyne B., Sproat R., “WordsEye: an Automatic Text-to-Scene Conversion System”, *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, p. 487-496, 2001.

Denis P., Muller P., “Predicting Globally-Coherent Temporal Structures From Texts Via Endpoint Inference and Graph Decomposition”, *Proceedings of IJCAI 2011*, p. 1788-1793, 2011.

Egenhofer M., “Reasoning About Binary Topological Relations”, *in* O. Gunther, H. Schek (eds), *Advances in Spatial Databases - SSD’91 Proceedings*, n° . 525 in *Lecture Notes in Computer Science*, Springer Verlag, Berlin, p. 143-160, 1991.

Ferro L., Gerber L., Mani I., Sundheim B., Wilson G., “TIDES 2005 Standard for the Annotation of Temporal Expressions”, *The MITRE Corporation*, 2005.

Forascu C., Tufi D., “Romanian TimeBank: an Annotated Parallel Corpus for Temporal Information”, *in* N. Calzolari, K. Choukri, T. Declerck, M. U. Dogan, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis (eds), *Proceedings of the 8th International Conference on Language Resources and Evaluation (LREC’12)*, European Language Resources Association (ELRA), Istanbul, Turkey, May, 2012.

Freksa C., “Using Orientation Information for Qualitative Spatial Reasoning”, *in* A. Frank, I. Campari, U. Formentini (eds), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, Proceedings of the International Conference GIS – From Space to Territory*, LNCS, Springer, Berlin, p. 162-178, 1992.

Galton A., “A Critical Examination of Allen’s Theory of Action and Time”, *Artificial Intelligence*, vol. 42, p. 159-188, 1990.

Galton A., *Qualitative Spatial Change*, Oxford University Press, Oxford, 2000.

Garbin E., Mani I., “Disambiguating Toponyms in News”, *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Vancouver, British Columbia, Canada, p. 363-370, October, 2005.

Herskovits A., *Language and Spatial Cognition. An Interdisciplinary Study of the Prepositions in English*, Cambridge University Press, Cambridge, London, 1986.

Hitzeman J., Moens M., Grover C., “Algorithms for Analysing the Temporal Structure of Discourse”, *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*, ACL, p. 253-260, 1995.

Hripcsak G., Elhadad N., Chen Y., Zhou L., Morrison F., “Using Empiric Semantic Correlation to Interpret Temporal Assertions in Clinical Texts”, *Journal of the American Medical Informatics Association*, vol. 16, no. 2, p. 220-227, 2009.

Huth M., Ryan M., *Logic in Computer Science: Modelling and Reasoning About Systems*, vol. 2, Cambridge University Press, Cambridge, UK, 2004.

Kamp H., Reyle U., *From Discourse to Logic*, Kluwer Academic Publishers, 1993.

Kontchakov R., Pratt-Hartmann I., Wolter F., Zakharyaschev M., “On the Computational Complexity of Spatial Logics With Connectedness Constraints”, *Logic for Programming, Artificial Intelligence, and Reasoning*, Springer, p. 574-589, 2008.

Kordjamshidi P., Bethard S., Moens M., “SemEval-2012 Task 3: Spatial Role Labeling”, *Proceedings of the 6th International Workshop on Semantic Evaluation (SemEval)*, 2012.

Kowalski R., Sergot M., “A Logic-Based Calculus of Events”, *New Generation Computing*, vol. 4, p. 67-95, 1986.

Lapata M., Lascarides A., “Learning Sentence-Internal Temporal Relations”, *Journal of Artificial Intelligence Research*, vol. 27, p. 85-117, 2006.

Levinson S. C., *Space in Language and Cognition*, Cambridge University Press, 2003.

Mandl T., Agosti M., Di Nunzio G., Yeh A., Mani I., Doran C., Schulz J., “LogCLEF 2009: the CLEF 2009 Multilingual Logfile Analysis Track Overview”, *Proceedings of the 10th cross-language evaluation forum conference on Multilingual information access evaluation: text retrieval experiments*, CLEF’09, Springer-Verlag, p. 508-517, 2009.

Mandl T., Carvalho P., Nunzio G. M. D., Gey F. C., Larson R. R., Santos D., Womser-Hacker C., “GeoCLEF 2008: the CLEF 2008 Cross-Language Geographic Information Retrieval Track Overview”, *CLEF*, p. 808-821, 2008.

Mani I., Doran C., Harris D., Hitzeman J., Quimby R., Richer J., Wellner B., Mardis S. A., Clancy S., “SpatialML: Annotation Scheme, Resources, and Evaluation”, *Language Resources and Evaluation*, vol. 44, no. 3, p. 263-280, 2010.

Mani I., Pustejovsky J., *Interpreting Motion – Grounded Representations for Spatial Language*, vol. 5 of *Explorations in Language and Space*, Oxford University Press, Oxford, 2012.

Mani I., Pustejovsky J., Gaizauskas R. (eds), *The Language of Time: a Reader*, Oxford University Press, Oxford, 2005.

Mani I., Schiffman B., Zhang J., “Inferring Temporal Ordering of Events in News”, *Proceedings of the Human Language Technology Conference (HLT-NAACL’03)*, 2003.

Mani I., Verhagen M., Wellner B., Lee C. M., Pustejovsky J., “Machine Learning of Temporal Relations”, *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Sydney, Australia, p. 753-760, July, 2006.

McDermott D., “A Temporal Logic for Reasoning About Processes and Actions”, *Cognitive Science*, vol. 6, p. 101-155, 1982.

Moratz R., Lücke D., Mossakowski T., “A Condensed Semantics for Qualitative Spatial Reasoning About Oriented Straight Line Segments”, *Artificial Intelligence*, vol. 175, no. 16-17, p. 2099-2127, 2011.

Muller P., “A Qualitative Theory of Motion Based on Spatio-Temporal Primitives”, *in* A. Cohn, L. Schubert, S. Shapiro (eds), *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR’98)*, Morgan Kaufmann, 1998.

Nieto M., Saurí R., Poveda M., “ModeS TimeBank: a Modern Spanish TimeBank Corpus”, *Procesamiento de Lenguaje Natural*, vol. 47, p. 259-267, 2011.

Papadias D., Theodoridis Y., Sellis T. K., Egenhofer M. J., “Topological Relations in the World of Minimum Bounding Rectangles: a Study with R-Trees”, *Proceeding of the SIGMOD Conference*, p. 92-103, 1995.

Prior A. N., *Past, Present, and Future*, Oxford University Press, Oxford, 1967.

Pustejovsky J., Hanks P., Saurí R., See A., Gaizauskas R., Setzer A., Radev D., Sundheim B., Day D., Ferro L., Lazo M., “The TimeBank Corpus”, *Proceedings of Corpus Linguistics*, p. 647-656, 2003.

Pustejovsky J., Ingria R., Saurí R., Castano J., Littman J., Gaizauskas R., Setzer A., Katz G., Mani I., "The Specification Language TimeML", *in* I. Mani, J. Pustejovsky, R. Gaizauskas (eds), *The Language of Time: a Reader*, Oxford University Press, Oxford, 2005.

Pustejovsky J., Moszkowicz J., "The Qualitative Spatial Dynamics of Motion in Language", *Spatial Cognition & Computation*, vol. 11, no. 1, p. 15-44, 2011.

Randell D., Cui Z., Cohn A., "A Spatial Logic Based on Regions and Connection", *Proceedings of KR'92*, Morgan Kaufmann, San Mateo, CA, 1992.

Reichenbach H., *Elements of Symbolic Logic*, McMillan, New York, chapter The Tenses of Verbs, p. 287-298, 1947.

Roller S., Speriosu M., Rallapalli S., Wing B., Baldridge J., "Supervised Text-Based Geolocation Using Language Models on an Adaptive Grid", *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Association for Computational Linguistics, Jeju Island, Korea, p. 1500-1510, July, 2012.

Saquete E., González J. L. V., Martínez-Barco P., Muñoz R., Llorens H., "Enhancing QA Systems With Complex Temporal Question Processing Capabilities", *J. Artif. Intell. Res. (JAIR)*, vol. 35, p. 775-811, 2009.

Saurí R., Badia T., *Catalan TimeBank 1.0 Corpus Documentation*, The Linguistic Data Consortium Corpus Catalog LDC2012T10, 2012.

Skiadopoulos S., Koubarakis M., "On the Consistency of Cardinal Direction Constraints", *Artificial Intelligence*, vol. 163, no. 1, p. 91-135, 2005.

Snoussi M., Gensel J., Davoine P., "Extending TimeML and SpatialML Languages to Handle Imperfect Spatiotemporal Information in the Context of Natural Hazards Studies", *Proceedings of the 15th AGILE International Conference on Geographic Information Science*, AGILE, 2012.

Spreyer K., Frank A., "Projection-Based Acquisition of a Temporal Labeller", *Proceedings of IJCNLP 2008*, 2008.

Talmy L., "Lexicalization Patterns: Semantic Structure in Lexical Forms", *Language Typology and Syntactic Description*, vol. 3, p. 57-149, 1985.

Talmy L., *Toward a Cognitive Semantics*, The MIT Press, 2000.

Tannier X., Muller P., "Evaluating Temporal Graphs Built From Texts Via Transitive Reduction", *Journal of Artificial Intelligence Research*, vol. 40, p. 375-413, 2011.

Turner R., Sripada S., Reiter E., "Generating Approximate Geographic Descriptions", *Empirical Methods in Natural Language Generation*, Springer-Verlag, p. 121-140, 2010.

UzZaman N., Allen J., "Temporal Evaluation", *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Portland, Oregon, USA, p. 351-356, June, 2011.

Van Beek P., "Reasoning About Qualitative Temporal Information", *Artificial Intelligence*, vol. 58, no. 1-3, p. 297-326, 1992.

Van Benthem J., *The Logic of Time*, Reidel, Dordrecht, 1983.

Van de Weghe N., *Representing and Reasoning About Moving Objects: a Qualitative Approach*, PhD thesis, Ghent University, Belgium, 2004.

Verhagen M., Saurí R., Caselli T., Pustejovsky J., "SemEval-2010 Task 13: TempEval-2", *Proceedings of the 5th International Workshop on Semantic Evaluation*, Association for Computational Linguistics, p. 57-62, 2010.

Vogel A., Jurafsky D., "Learning to Follow Navigational Directions", *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Uppsala, Sweden, p. 806-814, July, 2010.