Principles for annotating and reasoning with spatial information

Paul C. Morărescu

University of Texas at Dallas
Richardson TX 75080
paul@hlt.utdallas.edu

Abstract

In this paper we present the first phase of the ongoing SpaceBank project that attempts to create a linguistic resource for annotating and reasoning with spatial information from text. SpaceBank is the spatial counterpart of TimeBank, an electronic resource for temporal semantics and reasoning. The paper focuses on building an ontology of lexicalized spatial concepts. The textual occurrences of the concepts in this ontology will be annotated using the SpaceML language, briefly described here. SpaceBank is designed to be integrated with TimeBank, for a spatio-temporal model of the textual information.

1. Introduction

Most entities, relations, events or situtations that populate any fragment of text or speech are related to time and space. Most mental and computational world models include, or are based on, a spatio-temporal component. Therefore, capturing the spatial and temporal information is a key step in building such models from text, in content-based text processing tasks such as Information Extraction or Question Answering. The temporal aspect of the problem has received much attention from the NLP community in recent years, one notable result being the creation of TimeBank (Hobbs and Pan 2004). TimeBank is an integrated electronic resource based on the DAML-Time temporal ontology and a linguistic corpus. The occurrences of the ontology concepts as parts of larger events in the corpus are annotated using a specification language called TimeML. This resource is currently used in temporal inference tasks. In contrast, the spatial aspect of the problem has received much less attention. There are multiple projects and resources which consider this domain among other domains. They are surveyed in section 2 of this paper. Their value and limitations for the spatial content-based text processing are the object of Section 2, and Section 3.1 shows that the spatial domain still needs an integrated resource similar to TimeBank. This paper attempts to fill that niche by introducing SpaceBank, which is also designed to integrate with TimeBank for a spatio-temporal model of the world. Sections 3.2 and 3.3 describe the construction of a spatial ontology. Section 4 gives a brief specification of the SpaceML language, and Section 5 draws the conclusions.

2. Previous work

This section surveys the existing research tasks and resources that are relevant to this paper. They are grouped in three categories: hierarchies, ontologies and linguistic corpora. The first two attempt to organize the spatial concepts that populate certain areas of human cognition. The third attempts to annotate the occurrence of such concepts in texts.

2.1. Hierarchies of Spatial Concepts

A hierarchy covers a certain subdomain of language or cognition. For example, Named Entity Recognition (NER) systems such as (Bikel et al.1999) recognize locations, among other named entities. All kinds of locations are tagged simply as LOCATION. Only a few NER systems are able to distinguish subclasses of locations such as cities and countries. These systems are based on *Extended Named Entity Hierarchies*, which function as shallow ontologies. Geographical Information Systems [1] (GIS) need to recognize a large variety of geographical and geopolitical units. Most units form a hierarchy of *containment*, which is a fundamental *spatial relation*.

2.2. Spatial Ontologies

OpenCyc [2] is a large-scale ontology that includes the spatial domain, represented by complex hierarchies of *spatial relations* [3], *movements* [4], *geographical entities* [5], *paths and traversals* [6] and *transportation* [7].

Smaller ontologies covering the spatial domain, among other domains, include the Basic Formal Ontology (BFO, (Grenon and Smith 2004)), the Suggested Upper Merged Ontology (SUMO, (Nills and Pease 2001) and the Region Connection Calculus (RCC, (Cohn et al.1999)).

A comprehensive and general ontology is WordNet (Fellbaum 1998), which groups semantically synonymous words in sets called *synsets* and organizes them according to the *ISA* (*hypernymy*) and other semantic relations. Significant parts of WordNet's hierarchies are made of spatial concepts. WordNet is the base of our ontology, as shown in Section 3.3.

2.3. Spatial Linguistic Corpora

PropBank (Kingsbury et al.2001) annotates predicates and their arguments, which designate the objects of predicates and the modifiers of the entire predication. The locative modifiers represent the spatial information that apply to the predication. PropBank does not distinguish subclasses of locative modifiers, and it is thus of limited use to a study of the spatial domain of English. However, there are many shallow semantic parsers trained on PropBank, and one of

[1] http://www.geo.ed.ac.uk/home/giswww.html
[2] http://www.opencyc.org/
[3] http://www.cyc.com/cycdoc/vocab/spatial-vocab.html
[4] http://www.cyc.com/cycdoc/vocab/movement-vocab.html
[5] http://www.cyc.com/cycdoc/vocab/geography-vocab.html
[6] http://www.cyc.com/cycdoc/vocab/path-traversal-vocab.html
[7] http://www.cyc.com/cycdoc/vocab/transportation-vocab.html

their most difficult tasks is recognizing locative adjuncts, as shown in Table 1 [8]. This motivates the search for new ideas in the spatial domain, that are relevant at this level of linguistic representation.

Adjunct	Training	Top 3 Avg
AM-ADV	1727	44.21
AM-CAU	283	42.54
AM-DIR	231	35.35
AM-DIS	1077	61.30
AM-LOC	**1279**	**36.75**
AM-MNR	1337	39.58
AM-MOD	1753	95.33
AM-NEG	687	93.09
AM-PNC	446	32.10
AM-TMP	3567	56.57

Table 1: CoNLL 2004: top 3 systems' average performace in recognizing adjuncts in PropBank

The PropBank annotations of locative adjuncts are illustrated in the sentences [9] in Figure 1.

[The size and pace of disbursements]$_{Arg1}$ will$_{ArgM-MOD}$ accelerate$_{REL}$ further$_{Arg2-MNR}$ [under the Brady Plan]$_{ArgM-LOC}$, which promises larger and earlier disbursements to approved countries.

[Which flights]$_{Arg1}$ stop$_{REL}$ [in Minneapolis?]$_{ArgM-LOC}$

Here$_{ArgM-LOC}$ she$_{Arg0}$ delivers$_{REL}$, [especially during her enthusiastically awful rendition of the "Candy Man," which she sings while prancing around in a little cotton candy pink angora sweater that couldn't be more perfect.]$_{ArgM-TMP}$

Figure 1: Examples of PropBank annotations of locative adjuncts

The linguistic expressions annotated as locative adjuncts in the first and the third example above are not really locatives. In the first example, the expression *under the Brady Plan* is synonymous to the expression *according to the Brady Plan*, which clearly is not a locative, despite using the *under* preposition. This serves as a simple example of spatial vocabulary borrowed by non-spatial domains. In the third example, the word *here* in the expression *Here she delivers,* is not a locative either. Therefore, PropBank annotations of locative adjuncts may be too shallow to be included in a study of the spatial language.
The Automatic Content Extraction [10] (ACE-2004,2005) evaluations annotate the text occurrences of three types of spatial relations (*Located*, *Near* and *Part-whole*). These are relations between pairs of entities (e.g., a *GeoPolitical Entity* and an *ORGanization*), and thus differ from the relations between locative adjuncts and their predications in PropBank. Figure 2 presents an example annotation for each of the three relation types mentioned. [11]
Our analysis indicates that the ACE annotation does not suffer from the problems reported above for PropBank.

They$_{PER}$ left [the Al-Rashid hotel, where foreign visitors stay in Baghdad]$_{FAC}$

Hummer$_{PER}$ overtook Austrian David Kriner who placed second in the ski jump$_{PER}$ after about one kilometer.

[A mixture of gunpowder and ammonia$_{WEA}$]$_{WEA}$, the same materials used to make the bomb which killed around 211 people in september 1999 in Moscow and other areas of Russia

Figure 2: Examples of ACE annotations of PHYS-ical relations

That is, the information annotated as locative is indeed locative. However, the ACE corpus has its own limitations with respect to the spatial language. The most important limitation is that only three types of spatial relations are annotated, and even they are restricted to few classes of arguments.
We built a simple recognizer for all ACE relations, based on 13 features an the J48 classifier in WEKA, which was empirically observed to obtain good results on this problem, compared to other algorithms implemented in the same data mining library. Table 2 shows the results obtained by this classifier on detection and classification of spatial relations.

Relation	Training	P	R	F
PHYS_Located	1019	0.779	0.946	0.854
PHYS_Near	136	0.222	0.029	0.052
PHYS_Part-Whole	514	0.811	0.869	0.839

Table 2: Performace in recognizing ACE spatial relations

The explanation for the large drop in performance observed for the PHYS_Near relation is given by the confusion matrix presented in Table 3 [12]. Most instances of the PHYS_Near relation are misclassified either as PHYS_Located or PHYS_Part-Whole.

Relation	Training	a	b	c
PHYS_Located (a)	1019	894	6	0
PHYS_Near (b)	136	42	4	76
PHYS_Part-Whole (c)	514	9	6	487

Table 3: Confusion matrix for ACE spatial relations

FrameNet (Backer et al.2003) annotates semantic frames (Filmore 1985), which represent events or situations. Each frame contains several frame elements, which encode the participants and the properties of the event or state defined by the frame. FrameNet annotates spatial information both in the form of frames and in the form of frame elements. Table 4 lists the 112 frames, out of the 609, that we consider relevant for the spatial linguistic domain.
Figure 3 illustrates the FrameNet annotation of the *Emptying*, *Departing*, and *Attaching* frames. In these examples, each annotated instance of a spatial frame has at least two conveyors of spatial information: the verb and one spatial frame element. The two entities which serve as arguments to the spatial relation are not necessarily among the frame

[8] http://www.lsi.upc.edu/~srlconll/st04/slides/intro.pdf
[9] http://www.cs.rochester.edu/~gildea/Verbs/
[10] http://www.itl.nist.gov/iad/894.01/tests/ace/
[11] Examples are taken from ACE-2004 data. The labels indicate entity types, as follows: PER: person, FAC: facility, and WEA: weapon

[12] The difference betwen the number of training examples for each relation and the sum of the numbers in the confusion matrix corresponding to that relation comes from the non-spatial ACE relations, which were omitted in this table.

Adorning	Aggregate	Amalgamation
Architectural_part	Arriving	Attaching
Being_attached	Being_located	Biological_area
Body_decoration	Body_mark	Bounded_movement
Boundary	Bounded_entity	Bringing
Building	Building_subparts	Buildings
Cause_change_of_scalar_position	Cause_confinement	Cause_expansion
Cause_fluidic_motion	Cause_impact	Cause_motion
Cause_to_amalgamate	Cause_to_be_sharp	Cause_to_fragment
Cause_to_move	Cause_to_move_in_place	Change_position_on_a_scale
Change_posture	Closure	Congregating
Connecting_architecture	Connectors	Containers
Containing	Containment	Containment_relation
Departing	Dimension	Direction
Dispersal	Emanating	Emitting
Emptying	Escaping	Evading
Expansion	Filling	Fluidic_motion
Friction	Fullness	Gathering_up
Grinding	Hair_configuration	Halt
Hit_target	Immobilization	Import_export
In	Inchoative_attaching	Inclusion
Ingest_substance	Ingestion	Inspect
Into	Isolated_places	Light_movement
Locating	Location_of_light	Locative_relation
Mass_motion	Measure_area	Measure_linear_extent
Measure_volume	Motion	Motion_directional
Motion_scenario	Moving_in_place	Observable_bodyparts
Part_edge	Part_inner_outer	Part_observed_segments
Part_orientational	Part_piece	Part_whole
Path_shape	People_by_residence	Placing
Political_locales	Position_on_a_scale	Posture
Removing	Replacing	Reshaping
Residence	Self_motion	Sending
Separation	Shaped_part	Shapes
Sharpness	Shoot_projectiles	Sound_movement
Source_path_goal	Source_of_information	Speed
To	Trajector-Landmark	Travel
Vehicle	Wearing	

Table 4: Spatial frames in FrameNet

FE	Frames	FE	Frames
Place	264	Path_of_gaze	2
Source	52	Location_of_inspector	2
Goal	51	Reference_point	2
Path	48	Projectile	2
Area	28	Location_of_source	2
Distance	27	Location_of_tester	2
Body_part	19	Bounded_area	2
Location	19	Shape	2
Subregion	18	Goal_area	2
Sub_location	15	Vantage_point	2
Part	15	Size	2
Whole	14	Roadway	2
Container	14	Base_position	2
Content	14	Itinerary	2
Direction	13	Starting_point	2
Position	13	Connected_locations	2
Origin	10	Static_object	2
Constituent_parts	9	Following_distance	2
Trajector	7	Location_of_communicator	2
Body_location	7	Constant_location	2
Landmark	7	Holding_Location	2
Components	7	Target	2
Orientation	6	Enclosed_region	2
Container_possessor	6	Supporting_Body_part	2
Handle	6	Obstruction	2
Connector	5	Mode_of_Transportation	2
Form	5	Shape_prop	2
Dimension	5	Orientational_Location	2
Locus	5	Supporting_body_part	2
Location_of_perceiver	4	Containing_object	2
Parts	4	Destination_event	2
Fixed_location	4	Address	2
Sub_region	3	Subregion_bodypart	2
Route	3	Location_of_Event	2
Endpoints	2	Subpart	2
Course	2	Path_shape	2
Undesirable_location	2		

Table 5: Spatial frame elements in FrameNet

elements. In the first example, both the Source and Path frame elements describe the trajectory of the contents of the bucket that is emptied. Neither the contents nor its trajectory are explicitly mentioned in the sentence. In the second example, the Source frame element is itself implicit [13]. In the third example, both arguments (the Item and the Goal) are explicitly mentioned, and the relation is described by the verb.

[The players]$_{Agent}$ *emptied [the bucket]*$_{Source}$ *[over his head]*$_{Path}$.

[The plane]$_{Theme}$ *leaves [at seven]*$_{Time}$ *[]*$_{Source}$.

[The robber]$_{Agent}$ *tied [Harry]*$_{Item}$ *[to the chair]*$_{Goal}$ *[with a rope]*$_{Connector}$.

Figure 3: Examples of FrameNet annotations for spatial frames.

Table 5 lists the 73 frame elements, out of the 830, that we consider relevant for the spatial linguistic domain. Note that these frame elements are not necessarily related to the frames listed in the previous table.
Figure 4 illustrates the FrameNet annotations for the *Place*, *Path* and *Container* frame elements.

We$_{Agent}$ *immediately rushed to the ladies, washed Jessica carefully in the sink and dried her*$_{Dryee}$ *[under the hand drier]*$_{Instrument+Place}$.

Carefully$_{Manner}$, *he*$_{Self_{M}over}$ *crawled [the breadth of the building]*$_{Distance+Path}$ *and, to his relief, he saw that the house beneath him was one of a row.*

Gently$_{Manner}$ *melt [the butter and syrup]*$_{Undergoer}$ *[in a small pan]*$_{Container}$ *[]*$_{Agent}$.

Figure 4: Examples of FrameNet annotations for spatial frame elements.

[13]DNI, in FrameNet annotation

The first example shows that, in a given sentence, the same entity can play both a spatial role and a non-spatial role. The second example shows an entity that can play two different spatial roles. The third example reveals the spatial content of events like *melting*, which humans would not necessarily relate to a particular place.
We include here the work by (Setzer 2001), (Pustejovsky et al.2005) and (Hobbs and Pan 2004) in temporal semantics and reasoning. They argue that events are temporally anchored in the narrative, and this anchorage forms the foundation of our reasoning about how the world changes. To facilitate the automatic recognition and reasoning about events as a means of performing content-based (as opposed to keyword-based) text processing, they developed three resources: (1) TimeML, a specification language for annotating temporal entities and events in free text; (2) DAML-Time, an ontology for temporal concepts, designed for the temporal contents of Web pages and services; and (3) TimeBank, a corpus annotated with TimeML. We borrowed this architecture for SpaceBank.

3. Building a Lexicalized Spatial Ontology

3.1. The need for a new spatial ontology

The resources and ideas surveyed in the previous section cover a wide range of aspects of the spatial language and teach us valuable lessons. PropBank and FrameNet evidence the pervasiveness of space in event semantics and at the syntax-semantics interface. While PropBank labels all types of spatial information with the same label, FrameNet differentiates among 112 frames, as well as among 73 frame elements. This brings forward the granularity problem, since, for example, in the spatial domain it is not important whether Robert *sits* on a chair (the *Posture* and *Placing* frames) or he is *tied* to the chair (the *Attaching* frame). FrameNet also shows that not all the entities that play spatial roles are overt. The confusion matrix obtained for the ACE corpus, indicates that different spatial relations may have very similar linguistic representations, which make them difficult to discriminate based on surface features alone.

These resources also have important (and sometimes inherent) limitations in the spatial domain: (1) None of them is dedicated to spatial language, thus none can be expected to be either optimal or comprehensive in this area; (2) They are difficult to integrate, since each is underlied by a different theory and operates at a different linguistic level; (3) Typical ontologies mentioned in the previous section organize concepts to allow logical inference, but they do not offer an annotated corpus for training. Linguistic corpora have the opposite problem.

The spatial domain needs a resource that combines an ontology and an annotated corpus. This is the main purpose of SpaceBank.. It is also designed to be integrated with TimeBank, to produce a spatio-temporal world representation which combines the features of both components. To this end, SpaceBank uses the same linguistic corpus as TimeBank. The specification language, named SpaceML, is designed as an extension to TimeML. The spatial ontology, called DAML-Space, follows the steps of DAML-Time.

Among the problems that the design of a spatial ontology must solve, we mention the potential difference between the *ontological type* of an entity (e.g., ORGANIZATION) and its *spatial function* (e.g., DESTINATION, such as *Oracle* in the sentence *Bill Gates visited Oracle today*.). There is a many-to-many mapping between concepts and spatial functions.

As a real world example to illustrate the spatial concepts in a text that would need to be included in a spatial ontology, we take the first sentence of an Associated Press newswire article from Oct 20, 2005: *Hurricane Wilma's outer edge battered Cancun's white-sand beaches Thursday as officials ordered hotel guests to evacuate, tourists jockeyed for spots on the last flights out, and tens of thousands of people fled from Honduras to the Florida Keys.* There are several spatial concepts and relations in this sentence: (1) A special *part-whole* relation between Hurricane Wilma and its outer edge, invoked by the genitive construction. The relation is special because the *edge* is not a real *part* of the hurricane, but rather a *property*.; (2) A typical *part-whole* relation between the city of Cancun and its beaches, invoked by the genitive; (3) A *destination* relation between Hurricane's *edge* and the *beaches*, invoked by the semantic frame (i.e., *Cause_harm*) of the verb *batter*. Note that the beaches are the *Victim* frame element, thus there are frame elements that have a spatial function which is not their main function (and, consequently, they are not listed in Table 5); (4) An *origin* relation between the *hotel* and the guests' *evacuation*, invoked by the semantic frame (i.e., *Escaping*) of the verb *evacuate*. The hotel is the *Undesirable_location* frame element, whose main function is spatial, and thus it is listed in Table 5; (5) A *destination* relation between tourists and the spots in the last flights out, invoked by the verb *jockey* (note that *the spots* is a spatial metaphor for airplane seats and *flights out* also invokes a spatial relation of *origin* between the flights and the city; (6) An *origin* relation betwen *tens of thousands of people* and *Honduras*, and a *destination* relation between the same *people* and *Florida Keys*, both relations invoked by the verb *flee*. Since Honduras is a different location from Cancun, establishing a causal link between the Hurricane's presence in the latter location and tourists' fleeing the former location is a matter of extralinguistic spatial knowledge about the proximity of the two locations and the typical size of a Hurricane.

3.2. Primitives for a Spatial Ontology

(Gambarotto 2003) notes that identifying the ontological primitives of the spatial domain assumes the isolation of the basic objects of the geometry of our perception of the world. Note that the argument operates with the geometry of our perception, rather than the geometry of the world itself. This fact is the object of a general agreement in the cognitive literature, which uses the terms *psychological space* or *mental space* to designate the space we are talking or writing about, as opposed to the *physical space* we know from Geometry or Physics. Summarizing another general agreement in cognition, (Levinson 1996) states that the main difference between the mental space and the physical space is that the former is *relative*, whereas the latter is *absolute*. The relative space is anchored in the places occupied by physical objects around us and the relations between them. The linguistic counterparts of these cognitive arguments come from (Svorou 1994) and (Vandeloise 1991), among others. Svorou notes that we reason and talk about space in ways that reveal our *beliefs* rather than our *scientific knowledge* on the subject. Vandeloise argues for the inadequacy of both the geometric models and the logic models for the semantics of spatial prepositions in French.

To isolate the basic objects of the geometry of our perception, and thus to identify the ontological primitives, cross-linguistic studies propose three ontological topics, or types of concepts: (1) Topology, which refers to inclusion and contact; (2) Orientation; and (3) Distance. The ontological problem is to choose the primitive concepts, or objects, that populate these topics. At this point, we distinguish between primitives that are abstract concepts such as *points* or *vectors*, which make use of our knowledge in the Euclidean geometry to code both the orientation and the distance in the relative space, and *regions of space* corresponding to physical objects. (Setzer 2001) mentioned temporal ontologies based on points only, as well as temporal ontologies based on intervals only. Each has advantages and drawbacks. A particular approach is to use the intervals defined in (Allen 1983) for the temporal domain. (Gambarotto 2003) notes that using Allen's relations requires the approximation of physical objects by parallelepipedic shapes.

Hobbs [14] proposes an extended list of spatial ontological topics, which are listed in Table 6 next to their temporal counterparts.

The *topology* includes points, arcs, surfaces, volumes, connectedness and boundaries. The *dimension* and *orientation* include direction, path, conversions between different coordinate systems and the frame of reference. The Shape includes various 2D and 3D shapes, and the symmetry. The subtopics of *size* include length, area, volume, precise and uncertain measures. The *geopolitical divisions* were those occurring in FrameNet and ACE corpora, including country, state, county, city and continent. The *granularity* is a

[14] http://gunsight.metacarta.com/kornai/NAACL/WS9/Conf/hobbs.ppt

Space	Time
Topology	Topology
Dimension & Orientation	-
Shape	-
Size	Duration
Lat/long elevation	Clock & calendar
Geopolitical divisions	-
Granularity	Granularity
Aggregates, distributions	Temporal aggregates

Table 6: Topics in the spatial and temporal ontologies

design feature that allows a concept like a city to be treated either as a point, space or volume. This flexibility must be provided by the design of the entire ontology.

3.3. Our Lexicalized Spatial Ontology

Our main purpose in building a new spatial ontology is to facilitate the automatic recognition of spatial concepts and relations in newswire text, and to enable the automatic logic inference in the spatial domain. As such, we focus on lexicalized concepts that have clearly identifiable spatial characteristics [15], rather than abstract concepts that attempt to organize the entire spatial domain. For these reasons, rather than taking the rationalistic top-down way of fixing some primitives and then trying to fit all the textual spatial items into their bounds, we decided to follow a more empirical path. Namely, we start with a comprehensive semantic resource for English and then follow the bottom-up path to the primitives. As such, we decided to start the ontology construction from WordNet.

We identified the WordNet synsets corresponding to the nominal topics and concepts proposed by Hobbs. From these starting points, we explored the WordNet noun hierarchy both top-down and bottom-up. The bottom-up search always stops at the {entity} synset, which is the root of the hierarchy. The top-down search, which is initiated from every synset that we reach during the search, follows only the hyponymy and instance links that lead to spatial synsets and stops when it reaches a synset that has only spatial hyponyms and instances. Figure 5 illustrates the part of the top level ontology that subsumes the physical object hierarchy. The entire nominal ontology has 11,452 word-sense pairs, which represent 7.9% of the 145,104 nominal word-sense pairs in WordNet 2.1. We computed the polysemy frequencies to evaluate the need for Word Sense Disambiguation when recognizing in text concepts from our ontology. Table 7 presents the results. The average polysemy is 1.68, which is larger than 1.23, the average polysemy of the entire WordNet nominal hierarchy. The spatial language is ambiguous, therefore Word Sense Disambiguation is required.

P	1	2	3	4	5	6	7	8	11	28
F	3932	1750	726	244	79	40	16	10	1	1

Table 7: Polysemy in the nominal spatial ontology. P = polysemy, F = frequency

[15](Bierwisch-96) classifies concepts, their properties and their relations in four categories, according to their degree of spatiality: strictly spatial, intrinsically spatial, extrinsically spatial and aspatial. He also shows that space enters language through semantics, rather than through morphology or syntax, even in highly inflected languages.

4. The SpaceML annotation language

The SpaceML specificaction language has four major data structures which mirror the TimeML structures as illustrated in Table 8 SpaceML uses the event structure defined

TimeML	SpaceML
EVENT	EVENT
TIMEX3	SPACEX
SIGNAL	SIGNALS
LINK	LINKS

Table 8: The major data structures in TimeML and SpaceML

in TimeML. SPACEX is used to mark up explicit spatial expressions such as geopolitical entities, postal addresses, latitude and longitude and distances. The spatial expressions may be (a) Fully specified, such as *2601 North Floyd Rd, Richardson, TX 75080*; (b) Underspecified expressions, such as *the next block*; and (c) *Directions*, *Paths*, *Shapes* and *Orientations*, e.g., *towards the White House*. The attributes of the SPACEX tag mirror those of TIMEX3.

The SIGNALS tag annotates sections of text, typically function words, that indicate how spatial objects are to be related to each other. SIGNALS marks up several types of indicators: (1) Prepositions that have spatial meanings (among others): *on*, *in*, *at*, *to*, *between*, *above*, *in front of* and *outside of*. (2) Nominalizations such as *entrance, transportation* and *connection*; (3) Motion verbs such as *exit*, *advance* and *move*; and (4) *Dimensional adjectives* such as *high*, *short* and *thick*, as well as their nominalizations: *height*, *shortness* and *thickness*.

The LINKS tag encodes spatial relations, which can relate two events, two entities or an event to an entity. Entities may or may not be locative. The LINKS tag does not include subordinate and aspectual links as does the temporal LINK tag, but it accounts for the *frame of reference*, which is the origin of the coordinate system in which the spatial relation is to be interpreted. Specifying the frame of reference disambiguates expressions like *on the left sidewalk*, *in front of the tree*, and *behind the door*. We distinguish frames of reference centered on a *person* (usually the speaker), the *earth* (invoked by expressinons such as *the northwest corner*), and on the *carrier* or the*driving force* in case of a motion event. Besides the frame of reference, the LINKS tag marks up the *class* of spatial relation, e.g. *location*, *part-whole*, *near* and *containment*.

5. Conclusions and Future Work

This paper presents the first phase of an ongoing effort to build an electronic corpus annotated for spatial information. The types of concepts that are annotated, entities, events and spatial relations between them, are organized in a spatial ontology. The future work include the automation of the ontology construction and the expansion of the ontology over the verbal, adjectival and adverbial WordNet hierarchies. The starting points will be the verbs, adjectives and adverbs that are used in the glosses of the concepts from the nominal ontology introduced here. The assumption is that, in order to define nominal spatial concepts, one has to use verbal, adjectival and adverbial spatial concepts.

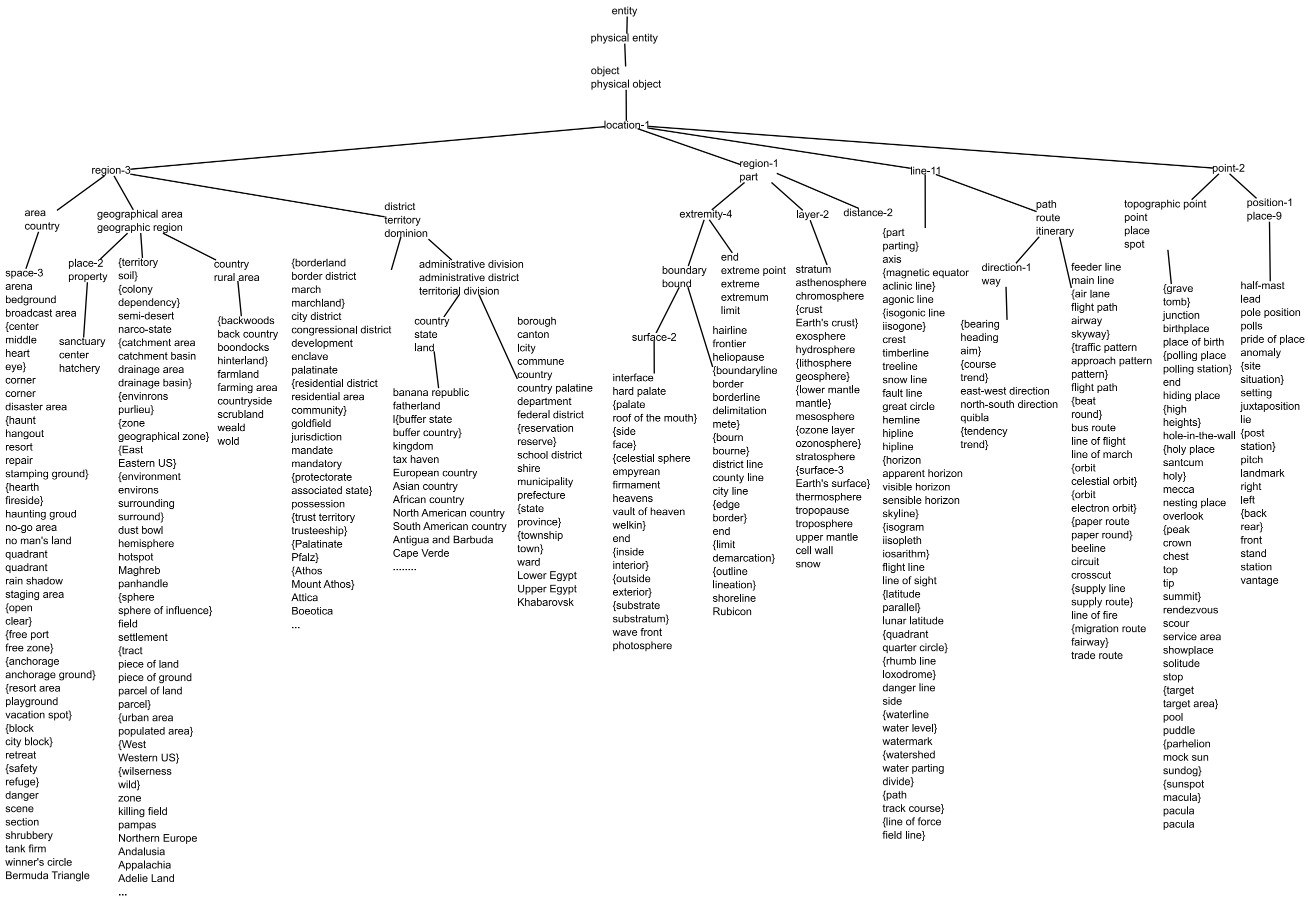

Figure 5: The top level of the nominal ontology (the physical entities). The numbers refer to WordNet senses.

6. References

James F. Allen. 1983 Maintaining Knowledge about Temporal Intervals. ACM Comm. 832-843

Collin F. Backer, Charles J. Fillmore and Beau Cronin. 2003. The Structure of the Framenet Database. In *International Journal of Lexicography*, Volume 16.3: 281-296..

Manfred Bierwisch. 1996. How Much Space Gets into Language? In (Bloom et al. 1996).

Daniel M. Bikel and Richard L. Schwartz and Ralph M. Weischedel. 1999. An Algorithm that Learns What's in a Name. In *Machine Learning*, Vol.34, No.1-3:211-231.

Paul Bloom, Mary A. Peterson, Lynn Nadel, and Merrill Garrett (Eds.). 1996. Language and Space. MIT Press, 1999.

Anthony G. Cohn and Brandon Bennett and John Gooday and Nicholas Mark Gotts. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. In *GeoInformatica*, Vol.1, No.3:275-316, 1997.

Christiane Fellbaum. 1999. WordNet - An Electronic Lexical Database. MIT Press, 1998.

Charles J. Fillmore. 1985. Frames and the semantics of understanding. In *Quaderni di Semantica*, Vol. 6.2: 222-254, 1985.

Christian Freksa, Christopher Habel, Karl F. Wender (Eds.). 1998. Spatial Cognition. Springer Verlag, 1998.

Pierre Gambarotto and Philippe Muller. 2003. Ontological Problems for the Semantics of Spatial Expressions in Natural Language. In (van der Zee 2003)

Pierre Grenon and Barry Smith. 2004. SNAP and SPAN: Towards Dynamic Spatial Ontology. In *Spatial Cognition and Computation*, 4(1), 6910, 2004.

Jerry R. Hobbs, and Feng Pan. 2004. An Ontology of Time for the Semantic Web. In *ACM Transactions on Asian Language Information Processing*, Vol. 3, No. 1, March 2004.

Paul Kingsbury, Martha Palmer, and Mitch Marcus. 2002. Adding Semantic Annotation to the Penn TreeBank. In *Proceedings of the Human Language Technology Conference*, 2002.

Stephen C. Levinson. 1996. Frames of Reference and Molyneux's Question: Crosslinguistic Evidence. In (Bloom et al. 1996).

Niles, I., and Pease, A. 2001. Towards a Standard Upper Ontology. In *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, Chris Welty and Barry Smith, (Eds.), Ogunquit, Maine, October 17-19, 2001., 2002.

James Pustejovsky, Robert Ingria, Roser Saur, Jos Castao, Jessica Littman, Robert Gaizauskas, Andrea Setzer, Graham Katz and Inderjeet Mani. 2005. The Specification Language TimeML. The Language of Time - a Reader, Oxford University Press, 2005

Andrea Setzer. 2001. Temporal Information in Newswire Articles: an Annotation Scheme and Corpus Study. PhD Thesis, University of Sheffield, 2001.

Soteria Svorou. 1994. The Grammar of Space. 1994. John Benjamins Publishing Co., 1993.

Claude Vandeloise. 1991. Spatial Prepositions. The University of Chicago Press, 1991.

Emile van der Zee and Jon Slack (Eds.). 2003. Representing Direction in Language and Space. Oxford University Press, 2003