A Corpus-based Analysis of Geometric Constraints on Projective Prepositions

Christian Hying
IMS, Universität Stuttgart
Azenbergstr. 12
70174 Stuttgart
Germany
christian.hying@ims.uni-stuttgart.de

Abstract

This paper presents a corpus-based method for automatic evaluation of geometric constraints on projective prepositions. The method is used to find an appropriate model of geometric constraints for a two-dimensional domain. Two simple models are evaluated against the uses of projective prepositions in a corpus of natural language dialogues to find the best parameters of these models. Both models cover more than 96% of the data correctly. An extra treatment of negative uses of projective prepositions (e.g. *A* is not above *B*) improves both models getting close to full coverage.

1 Introduction

This paper describes an empirical approach to finding an appropriate model of geometric constraints of projective prepositions with respect to a domain that is implicitly given by a corpus. We examine uses of the projective prepositions *above*, *below*, *to the right of*, *to the left of* and other projective prepositions whose orientation is aligned with one of the former, when they describe the location of an object relative to another object in two-dimensional space, see for example (1) and (2) relating to Figure 1:

(1) The circle is *to the right of* the rectangle.

(2) The circle is **not** *to the left of* the rectangle.

Henceforth, the term *located object* (LO) will be used to refer to the object whose location is specified and the term *reference object* (RO) to refer to the object relative to which the location is specified.

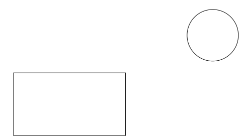

Figure 1: Example of a spatial scene.

In the examples, the *located object* is the circle in Figure 1 and the *reference object* is the rectangle. The notion *projective term* refers to the word of a projective preposition that determines the direction, e.g. the word *right* for the projective preposition *to the right of*. Let us call the use of the projective prepositions *positive use* when it is used in default context as in (1) and *negative use* when it is embedded under negation as in (2).

Geometric constraints that are associated with projective prepositions need to be such that they are met by positive uses such as (1) and violated by negative uses such as (2). Given that these sentences are appropriate uses to describe Figure 1, the spatial scene should meet the constraints that are associated with *to the right of* and violate the constraints of *to the left of*. It is obvious that this dual question of true or false invokes the issue of vagueness: We may find utterances describing a particular spatial scene and also their negations describing the same scene. For example, the following positive use of *above* may be appropriate to describe the spatial scene above – *The circle is above the rectangle* – but also the corresponding negative use in the sentence *The circle is not above the rectangle*.

We collect empirical evidence of uses of projective prepositions from the *HCRC Map Task* corpus (Anderson et al., 1991) – a corpus of human-human

Proceedings of the 4th ACL-SIGSEM Workshop on Prepositions, pages 1–8,
Prague, Czech Republic, June 2007. ©2007 Association for Computational Linguistics

dialogues. In contrast to other approaches that report empirical studies on geometric conditions of projective prepositions (Kelleher, 2003; Crawford et al., 2000; Logan and Sadler, 1996; Gapp, 1995; Abella, 1995) the resource used in this paper enables us to study their use in conversation.

This paper presents a new method for automatic evaluation of geometric constraints on projective prepositions with corpus data. We use this method to study the use of projective prepositions in human-human conversations and apply it to two models of geometric constraints with different parameters in order to evaluate the coverage for each parameter. A detailed analysis of incorrect cases leads us to a separate treatment of negative uses.

2 Related Work

This section introduces two types of spatial orientation relations that we are going to use as geometric constraints for projective prepositions in Section 4.

Orientation relations are defined with respect to a *frame of reference* that defines the actual alignment of directions (Levinson, 2003). The present study is carried out under the assumption of a fixed frame of reference such that the maps that are used as spatial data define the reference directions for *above*, *below*, *right*, and *left*. Although projective prepositions are in general sensitive to extra-geometric influences, e.g. dynamic LOs and ROs and functional relations between LO and RO (Coventry and Garrod, 2004), we do not expect that such effects play a role in the data, because the domain is static and it hardly contains any pairs of objects with a functional relationship.

In the literature, we find two paradigms for defining spatial orientation relations: the orthogonal projection paradigm and the angular deviation paradigm. For each paradigm we review a simple model and define different levels of granularity. The limitations of these simple models have been discussed at length, and more complex models have been proposed (Kelleher, 2003; Schmidtke, 2001; Crawford et al., 2000; Matsakis and Wendling, 1999; Fuhr et al., 1995; Abella and Kender, 1993; Wazinski, 1992). Nonetheless, it will turn out that

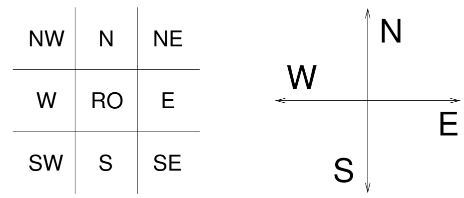

(a) Orthogonal projection model.

(b) Angular deviation model.

Figure 2: Definition of directions.

we can find for each simple model a level of granularity which covers more than 96% of the data.

Orthogonal projection. Orthogonal projection models define conditions on intervals that are the result of projecting two-dimensional or three-dimensional objects onto reference axes. (Papadias and Sellis, 1994), for example, define an orthogonal projection model with a horizontal and a vertical axis. Objects are represented by their projection onto these axes or, more illustrative, by bounding boxes. A bounding box of an object is the minimal rectangle with vertical and horizontal sides that contains the object. Lines which are defined by the sides of the bounding box of the reference object divide the space into nine regions. We refer to the regions around the bounding box of the reference object by means of the cardinal directions (*N,S,E,W,NW,NE,SW,SE*) as shown in Figure 2(a).

Let us define two relations OV and INC for expressing overlap and complete inclusion. A region A overlaps with a region B if and only if their intersection is not empty. A region A is completely included in B if and only if their intersection yields A:

(3) $OV(A,B) \Leftrightarrow A \cap B \neq \emptyset$ (*overlap*)
$INC(A,B) \Leftrightarrow A \cap B = A$ (*inclusion*)

The spatial orientation relations between LO and RO presented below are defined in terms of overlap and complete inclusion of LO with the nine regions around RO defined by the model. We exemplify the specification for the direction *north* using the auxiliary regions NHP and NXHP, where NHP $= NW \cup N \cup NE$ is the half-plane consisting of all northern regions and NXHP $=$ NHP $\cup W \cup RO \cup E$ is the (extended) half-plane which consists of all but the southern regions. For each orientation we define different levels of granularity – increasing index indicates wider interpretation. The idea is that relations on OP0 are as strict as possible and on OP7 as wide as possible. On granularity level OP0, the relation $north^0_{op}(LO, RO)$ is true if LO is completely included in the N-region. The predicate on the next granularity level is true if LO overlaps with the given N-region and is included in the northern half-plane NHP. Granularity level OP2 only requires inclusion in NHP. OP3 requires overlap with NHP and inclusion in the extended half-plane NXHP. On level OP4 the relation is true if LO is included in the extended half-plane NXHP. Relations on OP5 require overlap of LO with NXHP and LO must not overlap with S. On OP6 $north^6_{op}(LO, RO)$ is true if the LO does not overlap with S and on OP7 it is true if LO is not completely included in S. The same patterns apply to the relations $south^n_{op}$, $west^n_{op}$, and $east^n_{op}$.

OP0: $north^0_{op}(LO, RO) \Leftrightarrow INC(LO, N)$
OP1: $north^1_{op}(LO, RO) \Leftrightarrow OV(LO, N) \wedge INC(LO, \mathtt{NHP})$
OP2: $north^2_{op}(LO, RO) \Leftrightarrow INC(LO, \mathtt{NHP})$
OP3: $north^3_{op}(LO, RO) \Leftrightarrow$
$OV(LO, \mathtt{NHP}) \wedge INC(LO, \mathtt{NXHP})$
OP4: $north^4_{op}(LO, RO) \Leftrightarrow INC(LO, \mathtt{NXHP})$
OP5: $north^5_{op}(LO, RO) \Leftrightarrow$
$OV(LO, \mathtt{NXHP}) \wedge INC(LO, \mathtt{NXHP} \cup SW \cup SE)$
OP6: $north^6_{op}(LO, RO) \Leftrightarrow INC(LO, \mathtt{NXHP} \cup SW \cup SE)$
OP7: $north^7_{op}(LO, RO) \Leftrightarrow OV(LO, \mathtt{NXHP} \cup SW \cup SE)$

Note, that on granularity levels OP0 to OP3 opposite relations such as $north$ and $south$ are disjoint. Their extensions overlap on levels OP4 to OP7.

Angular deviation. Angular deviation models define conditions on one or more angles that represent how much LO deviates from a reference direction from the perspective of RO. In two-dimensional space there are four reference directions corresponding to the cardinal directions: $\vec{N}$, $\vec{S}$, $\vec{E}$, and $\vec{W}$. They are aligned with the vertical axis and the horizontal axis, respectively, as shown in Figure 2(b). Like the models presented in (Hernandez, 1994; Gapp, 1994) we use centroids to determine one single angle between RO and LO. Let the function $c(\cdot)$ return the centroid of its argument and let $\vec{o}$ be a vector from the centroid of the reference object to the centroid of the located object.

(4) $\vec{o} = \overrightarrow{c(RO)c(LO)}$

The angle between two vectors $\vec{a}$ and $\vec{b}$ is represented as $\angle(\vec{a}, \vec{b})$ and the angular deviation of $\vec{a}$ from the direction given by $\vec{b}$ is represented as $|\angle(\vec{a}, \vec{b})|$.

Orientation relations are defined via inequality conditions specifying that the deviation of the angle $\vec{o}$ from the corresponding reference direction is below or equal to a threshold. The threshold is defined as the granularity level multiplied by 10 degrees. We define 19 granularity levels ADn from n=0 to n=18 according to the pattern shown in (5). The same patterns with the reference directions $\vec{S}$, $\vec{W}$, and $\vec{E}$ apply to the relations $south^n_{ad}$, $west^n_{ad}$, and $east^n_{ad}$, respectively.

(5) ADn: $north^n_{ad}(LO, RO) \Leftrightarrow |\angle(\vec{N}, \vec{o})| \leq (n \cdot 10°)$

Note, that opposite relations such as $north$ and $south$ are disjoint on the levels from AD0 to AD8 and overlap from AD9 to AD18.

3 Data

This section describes the data that is used for the analysis of the semantics of projective prepositions. The data is an exhaustive collection of uses of projective prepositions occurring in the *HCRC Map Task* corpus (Anderson et al., 1991) where the speakers describe the location of a two-dimensional object relative to another two-dimensional object. The *HCRC Map Task* corpus is a collection of route description dialogues where one participant tries to explain a route printed on a map to another participant. It contains transcriptions of 128 dialogues which were recorded with 32 subjects. The maps are schematic maps containing line drawings of objects, so called *landmarks*. Examples of sections of the maps are shown in Section 5. The participants cannot see each other's maps so that the task can be accomplished only by means of what the participants say to one another. The two maps that are used for one dialogue are not exactly identical because not all landmarks have an identical counterpart on the other map. Therefore, the participants align their information about the maps by describing the location of landmarks.

TERM	Frequency	TERM	Frequency
above	87	under	5
left	86	up	5
below	77	west	3
right	65	north	2
underneath	52	south	2
beneath	7	east	1
bottom	7	upwards	1
top	7	over	1
down	5		

Table 1: Frequency of projective terms.

The present study selects those descriptions from the corpus that satisfy the following requirements:

Requirements:
(i) The description describes the location of one landmark relative to exactly one other landmark.
(ii) The description contains a projective preposition that is associated with one of the four cardinal directions from Figure 2(b).
(iii) The description does not contain any modifiers.

After having removed duplicates of descriptions occurring in the same dialogue, the set of data consists of 734 different uses of projective prepositions. 324 uses are filtered out by condition (iii) because they contain modifiers such as hedges (e.g. *just*), direction modifiers (e.g. *straight*), and distance modifiers (e.g. *2 cm*). The remaining set of data consists of 410 different uses of unmodified projective prepositions which further divides into 389 positive uses and 21 negative uses. Table 1 shows all projective terms ordered by frequency.

Spatial data. The corpus is supplemented by electronic copies of the maps that the participants have used. We created geometric representations of each map by redrawing the shape of each landmark and representing it as a closed polygon at the same location as the original landmark. All polygons are associated with unique identifiers. Let us define a function $polygon$ that yields the polygon definition for each landmark. Given that l is an identifier of a landmark and m an identifier of a map, the expression $polygon(l, m)$ returns the definition of the corresponding polygon.

Annotations. We identify all descriptions in the corpus that satisfy the requirements specified above. Then we mark the corresponding projective prepositions in the corpus and annotate them with the following type of information:

(6)
$$\begin{bmatrix} \text{TERM} & : \text{Projective Term} \\ \text{DIAL} & : \text{Dialogue Identifier} \\ \text{MAP} & : \text{Map Identifier} \\ \text{LO} & : \text{Landmark Identifier} \\ \text{RO} & : \text{Landmark Identifier} \\ \text{INT} & : (\texttt{pos} \mid \texttt{neg}) \end{bmatrix}$$

The feature `TERM` denotes the projective term. The feature `DIAL` holds a symbol that uniquely identifies the dialogue which the corresponding utterance occurs in. The feature `MAP` specifies the map which the corresponding utterance describes a part of. The features `LO` for located object and `RO` for reference object hold symbols that uniquely identify landmarks. Finally, the feature `INT` determines the way how to interpret the whole feature structure. It accepts one of the values `pos` and `neg`. The value `pos` indicates positive use of the projective preposition in the given utterance from the corpus: The feature structure is interpreted as the statement that the participant of dialogue `DIAL` who has map `MAP` produced utterances where the location of `LO` relative to `RO` on map `MAP` can be described correctly by the preposition in question. The value `neg` indicates a negative use of the preposition: The feature structure is interpreted as the statement that the participant of dialogue `DIAL` who has map `MAP` produced utterances where the *negation of the preposition* used is appropriate to describe the location of `LO` relative to `RO` on map `MAP`. In the corpus we find cases of explicit and implicit negation. The following two examples show cases of explicit negation.

(7) X is not below Y.

(8) A: Is X below Y?
B: No.

In the first example, the speaker makes a statement and uses a negated prepositional phrase. In the second example, the negation is triggered by a negative response to a question.

Implicit negations are triggered by rejections of alternatives. In the following example, participant A asks B about the truth of alternatives. If B chooses one alternative the others are rejected as incorrect:

(9) A: Is X above or below Y?
B: It's above.

Participant B states that the first alternative *X is above Y* is correct and thereby implicitly rejects the other alternative *X is below Y*.

4 Automatic Evaluation of Geometric Constraints on Projective Prepositions

This section describes a method of automatic evaluation of geometric constraints on projective prepositions with respect to the data described in the previous section.

For each level of granularity of the spatial orientation relations defined in Section 2 we define a model-theoretic semantics that maps projective prepositions onto truth conditions that are expressed in terms of these spatial orientation relations. In general, truth conditions determine the truth of a natural language expression with respect to a particular model of a situation. Applied to data used in this study this means that the truth conditions determine the applicability of projective prepositions with respect to a pair of landmarks that appear on the same map.

Semantics. For each projective preposition we define as many meanings as we have defined levels of granularity of spatial orientation relations in Section 2. We define a semantics on feature structure representations (6). Given the model α and the granularity level n we map a feature structure f onto the truth condition shown in (a) if $f.INT$=`pos` and onto (b) otherwise:

Let f be a feature structure of type (6), $\pi_{lo} = polygon(f.LO, f.MAP)$, and $\pi_{ro} = polygon(f.RO, f.MAP))$, then
(a) $\|f.TERM\|_{\alpha}^{n}(\pi_{lo}, \pi_{ro})$ if $f.INT$=`pos`;
(b) $\neg\|f.TERM\|_{\alpha}^{n}(\pi_{lo}, \pi_{ro})$ if $f.INT$=`neg`.

As said above, the function $polygon(\cdot,\cdot)$ yields a geometric representation of the landmark specified by a landmark identifer and a map identifier. The term $\|f.TERM\|_{\alpha}^{n}$ denotes the mapping of a projective term from Table 1 onto a spatial relation with the account α and the granularity level n. For example, the projective terms *above*, *top*, *up*, *upwards*, *over*, and *north* are all mapped onto $north_{\alpha}^{n}$-relations.[1] For example, if we evaluate the account using orthogonal projection and granularity level 0 the feature structure shown in (10) is mapped onto the formula $\neg north_{op}^{0}(\pi_1, \pi_2)$ where π_1 and π_2 are the polygons determined by `LO` and `RO`, respectively.

level	+pos	-pos	+neg	-neg	corr
OP0	79	310	21	0	100
OP1	249	140	21	0	270
OP2	346	43	19	2	365
OP3	376	13	16	5	392
OP4	385	4	11	10	396
OP5	386	3	7	14	393
OP6	387	2	2	19	389
OP7	389	0	0	21	389

Table 2: Results of the orthogonal projection models.

(10)
$$\begin{bmatrix} \text{TERM} & = \text{above} \\ \text{DIAL} & = \text{d0} \\ \text{MAP} & = \text{m2f} \\ \text{LO} & = \text{m2\_manned\_fort} \\ \text{RO} & = \text{m2\_rapids} \\ \text{INT} & = \texttt{neg} \end{bmatrix}$$

Automatic evaluation. We evaluate a semantics of projective prepositions by automatically computing truth conditions for each feature structure in the data and evaluating it with the corresponding geometric representations of RO and LO. If the truth value is *true* and the feature structure specifies positive use (*i.e.* `INT = pos`), then in this case the semantics is correct. Likewise, if the truth value is *false* and the data specifies negative use (`INT = neg`) the semantics is correct. In all other cases there is a mismatch between the semantics and the feature structure, so that the corresponding use of a projective preposition provides negative empirical evidence against the semantics.

5 Results and Discussion

The results of the evaluation are shown in Table 2 and Table 3. It comprises the evaluation of 27 semantic accounts corresponding to 8 levels of granularity of the orthogonal projection model (OP0 to

[1](O'Keefe, 1996) suggests that distinct projective prepositions can be associated with different levels of granularity, for example, *above* and *up*. For the present study the data is too sparse to compare such differences.

level	+pos	-pos	+neg	-neg	corr
AD0	0	389	21	0	21
AD1	116	273	21	0	137
AD2	179	210	21	0	200
AD3	250	139	21	0	271
AD4	291	98	21	0	312
AD5	320	69	21	0	341
AD6	347	42	20	1	367
AD7	370	19	18	3	388
AD8	382	7	17	4	399
AD9	385	4	14	7	399
AD10	386	3	12	9	398
AD11	386	3	10	11	396
AD12	386	3	7	14	393
AD13	386	3	5	16	391
AD14	387	2	5	16	392
AD15	388	1	4	17	392
AD16	388	1	3	18	391
AD17	388	1	1	20	389
AD18	389	0	0	21	389

Table 3: Results of the angular deviation models.

OP7) and 19 levels of granularity of the angular deviation model with thresholds from 0° (AD0) to 180° (AD18). The first column specifies the granularity level used. The evaluation of positive uses of projective prepositions is listed in the second and third column, the results for negative uses in the fourth and fifth column. The columns *+pos* and *+neg* report the number of correct cases in which the truth conditions are consistent with the value of the `INT` feature. The number of all correct cases is the sum of *+pos* and *+neg* and is printed in the last column with the label *corr*. The remaining columns *-pos* and *-neg* report incorrect truth conditions for positive and negative uses, respectively.

Orthogonal projection. Over all orthogonal projection models OP4 (included in extended half-plane) correctly covers a maximum number of 396 cases (96.6%).

For a more detailed analysis aiming at full coverage we take a closer look at the errors: there are 4 positive uses for which OP4 provides an incorrect semantics. The corpus reveals that three of these uses are not covered by OP4 because the speakers confused left and right. This confusion is apparent either because it is corrected by the speaker at a later point in the dialogue or because the use is obviously wrong. The remaining case is given by the following part of the corpus relating to Figure 3:

(11) *dialogue q4ec3, utterance 174f*

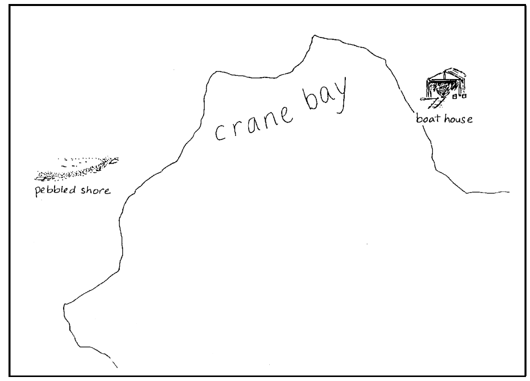

Figure 3: Pebbled shore, crane bay, and boat house.

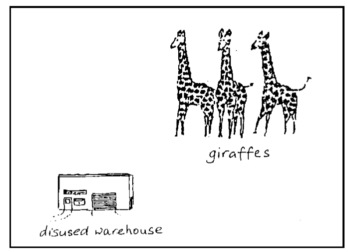

Figure 4: Disused warehouse and giraffes.

G: have you got anything **below** pebbled shore
F: washed stones and flag ship ... and bay

Note, that Figure 3 does not display the landmarks *washed stones* and *flag ship*. The participant F says that *crane bay* is below *pebbled shore*. This case is not captured by OP4 but by OP5 (overlap with extended half-plane).

All negative uses are correctly rejected by OP0 and OP1. The next level OP2 (i.e. completely included in half-plane) does not reject the following two cases:

(12) *dialogue q4nc2, utterance 264f*
G: i don't have a disused warehouse on mine
F: oh right. well it's just parallel to it ... like ... just ehm ... ... well not **underneath** the giraffes ... you know ...

(13) *dialogue q3nc7, utterance 66f*
G: is totem pole **below** the trout farm?
F: no i–, well, it's kind of opposite it

These uses are explicit negations. In (12) F says

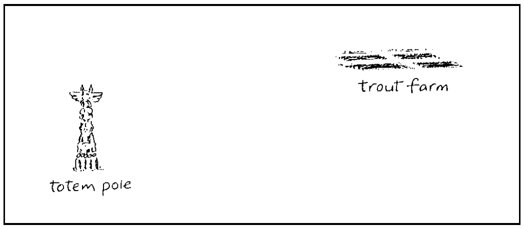

Figure 5: Totem pole and trout farm.

that the *warehouse* in Figure 4 is not underneath the *giraffes*. And in (13) F indicates that the *totem pole* is not below the *trout farm* in Figure 5. As said before, OP1 is the most general model that rejects these cases.

To summarise, a semantics that aims at covering all of the *good* data employs OP5 for positive uses and OP1 for negative uses. [2] On level OP5 and to a lesser extent on OP4, the extensions of opposite relations such as *above* and *below* overlap, because all objects that are included in the union of the regions W, RO, and E are both *above* and *below* relative to the reference object. Since on OP4 the overlap is smaller than on OP5 it is better to use OP4 instead. A combination of OP4 for positive uses and OP1 for negative uses still covers almost all of the good data (99.8%).

Angular deviation. Over all angular deviation models AD8 and AD9 correctly cover a maximum number of 399 cases (97.3%).

On level AD9 there are 4 positive uses with an incorrect semantics. Again the same three uses as above are due to confusion of left and right. The remaining use is the following utterance, which relates to the part of a map depicted in Figure 3. The narrowest model that covers this use is AD13:

(14) *dialogue q4ec3, utterance 332*
my boat house is ... down **below** crane bay

All negative uses are correctly rejected by all models from AD0 to AD5. Model AD6 does not predict rejection of the case which has already been described above in (12). AD7 additionally produces two further errors in the following two cases which describe Figure 6(a) and Figure 6(b), respectively.

[2] *Good* data means all data excluding the cases where left and right was confused.

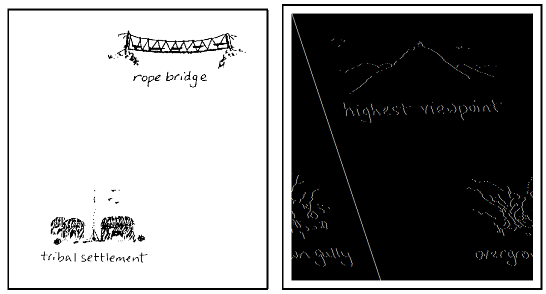

(a) Tribal settlement and rope bridge.
(b) Highest viewpoint and overgrown gully

Figure 6: Section of maps 13 and 10.

(15) *dialogue q4ec1, utterance 10f*
F: is it **underneath** the rope bridge or to the **left**?
G: it's **underneath** the rope bridge

(16) *dialogue q4ec8, utterance 41f*
G: and eh to the ... **left** or **right** of highest viewpoint
F: ... it's **beneath** it

These examples show implicit negative uses. The utterances in (15) give rise to the interpretation that the *tribal settlement* is *not* to the left *rope bridge*. And the utterances in (16) imply that the *overgrown gully* is neither to the left nor to the right of the *highest viewpoint*. These three negative uses and again the localisation of the *totem pole* in (13) have not been modelled correctly by the semantics that employs AD8.

To summarise, a semantics aiming to cover all of the *good* data uses AD13 for positive uses and AD5 for negative uses. Considering that the extensions of the opposite relations in AD13 overlap to a great extent, it is better to use a combination of AD9 for positive uses and AD5 for negative uses which still covers all of the good data except one case (99.8%).

If we compare the angular deviation model (AD9/AD5) with the orthogonal projection model (OP4/OP1), the angular deviation model is superior, because in AD9 the extensions of opposite relations such as *above* and *below* only have a very small overlap, namely when the angular deviation is exactly 90°, while in OP4 the overlap is much more significant.

6 Summary and Conclusion

This paper described a method to evaluate geometric constraints on projective prepositions with empirical data extracted from a corpus of human-human conversations. The key feature of the approach is the annotation of projective prepositions in the corpus with links to geometric representations of the objects that the arguments of the prepositions refer to. The data is used to automatically apply and evaluate different granularity levels of a semantics building upon a simple orthogonal projection model and a simple angular deviation model. Both models cover more than 96% of the data correctly. Further refinement shows that the angular deviation model covers the data almost perfectly (99.8%) if we provide an extra treatment for negative uses, so that positive uses are accepted when the angular deviation is below 90° and negative uses are accepted when the angular deviation is greater than 50°.

References

Alicia Abella and John R. Kender. 1993. Qualitatively describing objects using spatial prepositions. In *National Conference on Artificial Intelligence*, pages 536–540.

Alicia Abella. 1995. *From imagery to salience: locative expressions in context*. Ph.D. thesis, Computer Science Department, Columbia Universtity, New York, NY.

A. Anderson, M. Bader, E. Bard, E. Boyle, G. M. Doherty, S. Garrod, S. Isard, J. Kowtko, J. McAllister, J. Miller, C. Sotillo, H. S. Thompson, and R. Weinert. 1991. The HCRC Map Task Corpus. *Language and Speech*, 34(4):351–366.

Kenny Coventry and Simon Garrod. 2004. Towards a classification of extra-geometric influences on the comprehension of spatial prepositions. In *Functional Features in Language and Space*.

L. Elizabeth Crawford, Terry Regier, and Janellen Huttenlocher. 2000. Linguistic and non-linguistic spatial categorization. *Cognition*, 75(3):209–235.

T. Fuhr, G. Socher, C. Scheering, and G. Sagerer. 1995. A three-dimensional spatial model for the interpretation of image data. In *IJCAI-95 Workshop on Representation and Processing of Spatial Expressions*, Montreal, Canada.

Klaus-Peter Gapp. 1994. Basic meanings of spatial relations: computation and evaluation in 3d space. In *AAAI'94: Proceedings of the twelfth national conference on Artificial intelligence (vol. 2)*, pages 1393–1398, Menlo Park, CA, USA. American Association for Artificial Intelligence.

K.-P. Gapp. 1995. An empirically validated model for computing spatial relations. In I. Wachsmuth, C.-R. Rollinger, and W. Brauer, editors, *KI-95: Advances in Artificial Intelligence. 19th German Annual Conference on Artificial Intelligence*, pages 245–256. Springer, Berlin, Heidelberg.

Daniel Hernandez. 1994. *Qualitative Representation of Spatial Knowledge*. Springer-Verlag New York, Inc.

John Kelleher. 2003. *A Perceptually Based Computational Framework for the Interpretation of Spatial Language in 3D Simulated Environments*. Ph.D. thesis, Dublin City University, Dublin.

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

Gordon D. Logan and Daniel D. Sadler. 1996. A computational analysis of the apprehension of spatial relations. In Paul Bloom, Mary A. Peterson, Lynn Nadel, and Merril G. Garrett, editors, *Language and Space*. MIT Press.

Pascal Matsakis and Laurent Wendling. 1999. A new way to represent the relative position between areal objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(7):634–643.

John O'Keefe. 1996. The spatial prepositions in English, vector grammar, and the cognitive map theory. In Paul Bloom, Mary A. Peterson, Lynn Nadel, and Merril G. Garrett, editors, *Language and Space*. MIT Press.

Dimitris Papadias and Timos K. Sellis. 1994. Qualitative representation of spatial knowledge in two-dimensional space. *VLDB Journal: Very Large Data Bases*, 3(4):479–516.

Hedda R. Schmidtke. 2001. The house is north of the river: Relative localization of extended objects. In D.R. Montello, editor, *COSIT 2001, LNCS 2205*, pages 415–430.

Peter Wazinski. 1992. Generating spatial descriptions for cross-modal references. In *Proceedings of the third conference on Applied natural language processing*, pages 56–63. Association for Computational Linguistics.