



2006-01-01

Spatial Prepositions in Context: The Semantics of near in the Presence of Distractor Objects.

Fintan Costello

University College Dublin, fintan.costello@ucd.ie

John D. Kelleher

Dublin Institute of Technology, john.d.kelleher@dit.ie

Follow this and additional works at: <http://arrow.dit.ie/scschcomcon>

 Part of the [Computational Engineering Commons](#)

Recommended Citation

Costello, F., Kelleher, J.: Spatial Prepositions in Context: The Semantics of near in the Presence of Distractor Objects. In Proceedings of the 3rd ACL-Sigsem Workshop on The Linguistic Dimensions of Prepositions and their Use in Computational Linguistics Formalisms and Applications, Association for Computational Linguistics (Arsenijevic, B., Baldwin, T. and Trawinski, B.)

This Conference Paper is brought to you for free and open access by the School of Computing at ARROW@DIT. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@DIT.

For more information, please contact yvonne.desmond@dit.ie, arrow.admin@dit.ie, brian.widdis@dit.ie.



School of Computing

Articles

Dublin Institute of Technology

Year 2006

Spatial Prepositions in Context: The
Semantics of near in the Presence of
Distractor Objects.

Fintan Costello*

John D. Kelleher†

*University College Dublin, fintan.costello@ucd.ie

†Dublin Institute of Technology, john.d.kelleher@dit.ie

This paper is posted at ARROW@DIT.

<http://arrow.dit.ie/scschcomart/1>

— Use Licence —

Attribution-NonCommercial-ShareAlike 1.0

You are free:

- to copy, distribute, display, and perform the work
- to make derivative works

Under the following conditions:

- Attribution.
You must give the original author credit.
- Non-Commercial.
You may not use this work for commercial purposes.
- Share Alike.
If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.

For any reuse or distribution, you must make clear to others the license terms of this work. Any of these conditions can be waived if you get permission from the author.

Your fair use and other rights are in no way affected by the above.

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License. To view a copy of this license, visit:

- URL (human-readable summary):
<http://creativecommons.org/licenses/by-nc-sa/1.0/>
 - URL (legal code):
<http://creativecommons.org/worldwide/uk/translated-license>
-

Spatial Prepositions in Context: The Semantics of *near* in the Presence of Distractor Objects

Fintan Costello

Department of Computer Science
University College Dublin
fintan.costello@ucd.ie

John Kelleher

School of Computing
Dublin Institute of Technology
john.kelleher@comp.dit.ie

Abstract

The paper examines how people’s judgments of proximity between two objects are influenced by the presence of a third object. In an experiment participants were presented with images containing three shapes in different relative positions, and asked to rate the acceptability of a locative expression such as ‘the circle is near the triangle’ as descriptions of those images. The results showed an interaction between the relative positions of objects and the linguistic roles that those objects play in the locative expression: proximity was a decreasing function of the distance between the object in the head position in the expression and that in the relative clause position, and an increasing function the distance between the head and the third, distractor object. This finding leads us to a new account for the semantics of spatial prepositions such as *near*.

1 Introduction

In this paper, we present an empirical study of the cognitive representations underpinning the uses of proximal descriptions in locative spatial expressions. A spatial locative expression consists of a locative prepositional phrase together with whatever the phrase modifies (noun, clause, etc.). In their simplest form, a locative expression consists of a prepositional phrase modifying a noun phrase, for example *the man near the desk*. People often use spatial locatives to denote objects in a visual scene. Understanding such references involves co-ordination between a perceptual event and a linguistic utterance. Consequently, the study of spatial locatives affords the opportunity to examine

some aspects of the grounding of language in non-language.

The conception of space underlying spatial locatives is fundamentally relativistic: the location of one object is specified relative to another whose location is usually assumed by the speaker to be known by the hearer. Moreover, unpinning this relativistic notion of space is the concept of proximity. Consequently, the notion of proximity is an important concept at the core of human spatial cognition. Proximal spatial relationships are often described using topological prepositions, e.g. *at*, *on*, *near*, etc.

Terminology In this paper we use the term **target** (T) to refer to the head of a locative expression (the object which is being located by that expression) and the term **landmark** (L) to refer to the relative clause in that expression (the object relative to which the head’s location is described), see Example (1).

Example 1 . [The man]_T near [the table]_L.

We will use the term **distractor** to describe any object in the visual context that is neither the landmark nor the target.

Contributions The paper reports on a psycholinguistic experiment that examines proximity. Previous psycholinguistic work on proximal relations, (Logan and Sadler, 1996), has not examined the effects other objects in the scene (i.e., distractors) may have on the spatial relationship between a landmark and a target. The experiment described in this paper compares peoples’ judgments of proximity between target and landmark objects when they are presented alone and when there are presented along with other distractor objects. Based on the results of this experiment we

propose a new model for the semantics of spatial prepositions such as *near*.

Overview In §2 we review previous work. In §3 we describe the experiment. In §4 we present the results of the experiment and our analysis. The paper finishes with conclusions, §5.

2 Related Work

In this section we review previous psycholinguistic experiments that examined proximal spatial relations. We then present example spatial contexts, that the previous experiments did not examine, which motivate the hypothesis tested in this paper: the location of other objects in a scene can interfere with the acceptability of a proximal description being used to describe the spatial relationship between a landmark and a target.

1.74	1.90	2.84	3.16	2.34	1.81	2.13
2.61	3.84	4.66	4.97	4.90	3.56	3.26
4.06	5.56	7.55	7.97	7.29	4.80	3.91
4.47	5.91	8.52	O	7.90	6.13	4.46
3.47	4.81	6.94	7.56	7.31	5.59	3.63
3.25	4.03	4.50	4.78	4.41	3.47	3.10
1.84	2.23	2.03	3.06	2.53	2.13	2.00

Figure 1: 7-by-7 cell grid with mean goodness ratings for the relation *near* as a function of the position occupied by X.

Spatial reasoning is a complex activity that involves at least two levels of representation and reasoning: a *geometric level* where metric, topological, and projective properties are handled, (Herskovits, 1986); and a *functional level* where the normal function of an entity affects the spatial relationships attributed to it in a context, see (Vandeloise, 1991; Coventry, 1998; Garrod et al., 1999).

There has been a lot of experimental work done on spatial reasoning and language: (Carlson-Radvansky and Irwin, 1993; Carlson-Radvansky and Irwin, 1994; Hayward and Tarr, 1995; Gapp, 1995; Logan and Sadler, 1996; Carlson-

Radvansky and Logan, 1997; Coventry, 1998; Garrod et al., 1999; Regier and Carlson, 2001; Kelleher and Costello, 2005). Of these only (Logan and Sadler, 1996) examined topological prepositions in a context where functional factors were excluded.

The term **spatial template** denotes the representation of the regions of acceptability associated with a preposition. It is centred on the landmark and identifies for each point in its space the acceptability of the spatial relationship between the landmark and the target appearing at that point being described by the preposition (Logan and Sadler, 1996).

The concept of a spatial template emerged from psycholinguistic experiments reported in (Logan and Sadler, 1996). These experiments examined various spatial prepositions. In these experiments, a human subject was shown sentences, each with a picture of a spatial configuration. Every sentence was of the form “The X is [relation] the O”. The accompanying picture contained an O in the center of an invisible 7-by-7 cell grid, and an X in one of the 48 surrounding positions. The subject then had to rate how well the sentence described the picture, on a scale from 1(bad) to 9(good).

Figure 1 gives the mean goodness rating for the relation “near to” as a function of the position occupied by the X, as reported in (Logan and Sadler, 1996). If we plot the mean goodness rating for “near” against the distance between target X and landmark O, we get the graph in Figure 2.



Figure 2: Mean goodness rating vs. distance between X and O.

Both the figure and the graph make it clear that the ratings diminish as we increase the distance between X and O. At the same time, we can observe that even at the extremes of the grid the ratings were still above 1 (the minimum rating). Indeed, in the four corners of the grid, the points most distant from the landmark, the mean ratings nearly average twice the minimum rating.

However in certain contexts other factors, apart from the distance between the landmark and the target, affect the applicability of a proximal relation as a description of the target’s position relative to the landmark. For example, consider the two scenes (side-view) given in Figure 2. In the scene on the left-hand side, we can use the description “the blue box is near the black box” to describe object (a). However, consider now the scene on the right-hand side. In this context, the description “the blue box is near the black box” seems inappropriate as an expression describing (a). The placing of object (c) beside (b) would appear to interfere with the appropriateness of using a proximal relation to locate (a) relative to (b), even though the absolute distance between (a) and (b) has not changed.



Figure 3: Proximity and distance

In summary, there is empirical evidence that indicates that as the distance between the landmark and the target increases the applicability of a proximal description decreases. Furthermore, there is anecdotal evidence that the location of other distractor objects in context may interfere with applicability of a proximal description between a target and landmark object. The experiment presented in this paper is designed to empirically test the affect of distractor objects on proximity judgements.

3 Experiment

This work examines the impact of distractor objects on subjects’ judgment of proximity between the target and the landmark objects. To do this, we examine the changes in participants judgements of the appropriateness of the topological preposition *near* being used to describe a spatial configuration of the target and landmark objects when a distractor object was present and when it was removed.

Topological prepositions (e.g., *at*, *on*, *in*, *near*) are often used to describe proximal spatial relationships. However, the semantics of a given topological preposition also reflects functional (Garrod et al., 1999), directional (Logan and Sadler, 1996) and topological factors.¹ Consequently, it was im-

¹See (Cohn et al., 1997) for a description different topological relationships.

portant to control for these factors during the design of the experiment.

Functional factors were controlled for by using simple shapes in the stimuli. The preposition *near* was used to control the impact of directional factors. Previous psycholinguistic work indicated that *near* was not affected by any directional preferences. Finally, the influence of topological factors was controlled for by ensuring that the landmark and target maintained a consistent topological relationship (the objects never touched, overlapped or were contained in other objects).

3.1 Material and Subjects

All images used in this experiment contained a central landmark and a target. In most of the images there was also another object, which we will refer to as the **distractor**. All of these objects were coloured shapes, a circle, triangle or square. However, none of the images contained two objects that were the same shape or the same colour.

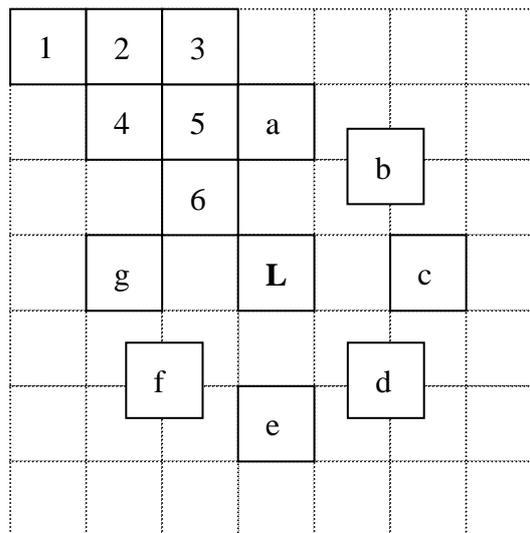


Figure 4: Relative locations of landmark (L) target positions (1..6) and distractor positions (a..g) in images used in the experiment.

The landmark was always placed in the middle of a seven by seven grid (row four, column four). There were 48 images in total, divided into 8 groups of 6 images each. Each image in a group contained the target object placed in one of 6 different cells on the grid, numbered from 1 to 6 (see Figure 4). As Figure 4 shows, we number these target positions according to their nearness to the landmark.

Each group, then, contains images with targets at positions 1, 2, 3, 4, 5 and 6. Groups are organised according to the presence and position of a distractor object. Figure 4 shows 7 different positions used for the distractor object, labelled *a, b, c, d, e, f, g* and *h*. In each of these positions the distractor is equidistant from the landmark. In group *a* the distractor is directly above the landmark, in group *b* the distractor is rotated 45 degrees clockwise from the vertical, in group *c* it is directly to the right of the landmark, in *d* is rotated 135 degrees clockwise from the vertical, and so on. Notice that some of these distractor positions (*b, d*, and *f*) are not aligned with the grid. This realignment is necessary to ensure that the distractor object is always the same distance from the landmark. Each of these groups of images used in the experiment corresponds to one of these 7 distractor positions, with a distractor object occurring at that position for every image in that group. In addition, there is an eight group (which we label as group *x*), in which no distractor object occurs.

Previous studies of how people judge proximity have typically examined judgments where the target is above, below, to the left or right of the landmark. The results of these studies showed that these distinctions are relatively unimportant, and the gradient of proximity observed tends to be symmetrical around the landmark. For this reason, in our study we ignore these factors and present landmark, target and distractor randomly rotated (so that some participants in our experiment will see the image with target at position 1 and distractor at position *a* in a rotated form where position 1 is below the landmark and position *a* is to the right of the landmark, but others will see the same relative positions at different rotations). In each image all objects present were placed exactly at the center of the cell representing their position.

During the experiment, each image was displayed with a sentence of the form *The ____ is near the ____*. The blanks were filled with a description of the target and landmark respectively. The sentence was presented under the image. 12 participants took part in this experiment.

3.2 Procedure

There were 48 trials, constructed from the following variables: 8 distractor conditions * 6 target positions. To avoid sequence effects the landmark, target and distractor colour and shape were ran-

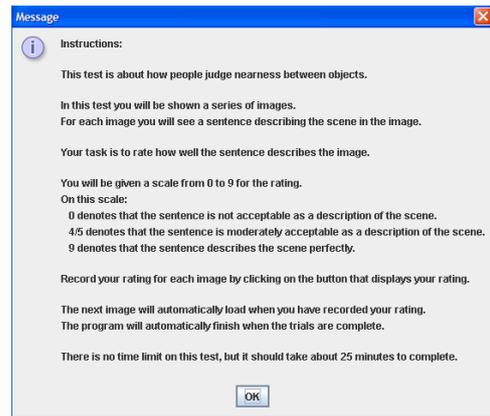


Figure 5: Experiment instructions.

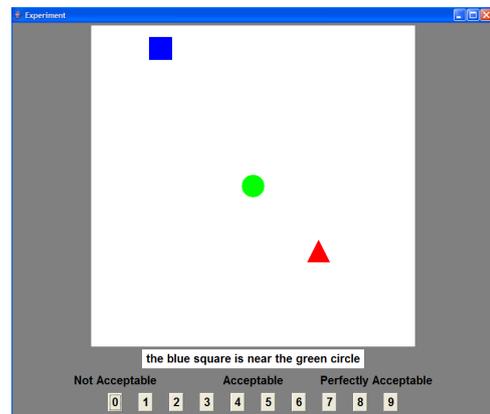


Figure 6: Sample trial from the experiment.

domly modified for each trial and the distractor condition and target location were randomly selected for each trial. Each trial was randomly reflected across the horizontal, vertical, or diagonal axes. Trials were presented in a different random order to each participant.

Participants were instructed that they would be shown sentence-picture pairs and were asked to rate the acceptability of the sentence as a description of the picture using a 10-point scale, with zero denoting not acceptable at all; four or five denoting moderately acceptable; and nine perfectly acceptable. Figure 5 presents the instructions given to each participant before the experiment. Trials were self-paced, and the experiments lasted about 25-30 minutes. Figure 6 illustrates how the trials were presented.

4 Results and Discussion

There are two questions we want to ask in our examination of people's proximity judgments in the presence of distractor objects. First, does the pres-

ence of a distractor make any noticeable difference in people’s judgements of proximity? Second, if the presence of a distractor does influence proximity judgements, how does that influence operate?

We address the first question (does the distractor object have an influence on proximity judgments) by comparing the results obtained for images in group x (in which there was no distractor) with results obtained from other groups. In particular, we compare the results from this group with those obtained from groups c , d and e : the three groups in which the distractor object is furthest from the set of target positions used (as Figure 4 shows, distractor positions c , d , and e are all on the opposite side of the landmark from the set of target positions). We focus on comparison with groups c , d , and e because results for the other groups are complicated by the fact that people’s proximity judgements are influenced by the closeness of a distractor object to the target (as we will see later).

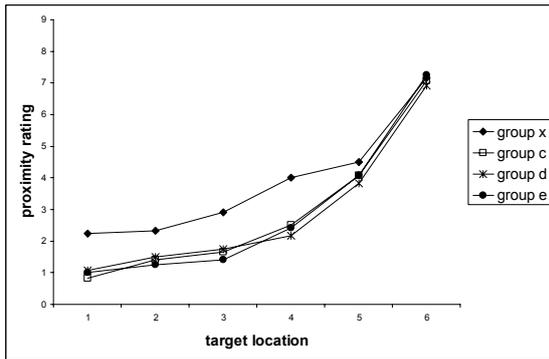


Figure 7: mean proximity rating for target locations for group x (no distractor) and groups c , d , and e (distractors present behind landmark)

Figure 7 shows the average proximity rating given by participants for the 6 targets 1 to 6 for group x (in which there was no distractor object) and for groups c , d , and e (in which distractors occurred on the opposite side of the landmark from the target). Clearly, all three sets of distractor responses are very similar to each other, and are all noticeably different from the no-distractor response. This difference was shown to be statistically significant in a by-subjects analysis comparing subjects’ responses for groups c, d and e with their responses for group x . This comparison showed that subjects produced significantly lower proximity ratings for group c than group x (Wilcoxon signed-rank test $W+ = 55.50, W- =$

$10.50, N = 11, p \leq 0.05$), lower ratings for group d than group x ($W+ = 48.50, W- = 6.50, N = 10, p \leq 0.05$) and lower ratings for group e than group x ($W+ = 51.50, W- = 3.50, N = 10, p \leq 0.01$). (We exclude one subject from this analysis because they mistakenly gave the lowest possible proximity rating of 0 to the item closest to the landmark in group x).

These results show that the presence of a distractor object reliably influences people’s proximity judgements. But how does this influence operate? We examine this by considering two factors: the relationship between people’s proximity judgement and the distance from the landmark to the target object, and relationship between people’s proximity judgement and the distance from the distractor to the target object. (Recall that in the design of our materials, the distance from landmark to distractor was kept constant so target-to-landmark and target-to-distractors are the two factors that vary in our experiment.)

We can formalise our expectations about proximity judgements as follows. Let T be the target whose proximity to the landmark we’re trying to judge, let L and D be the landmark and distractor objects respectively, and let $dist(A, B)$ be the computed distance between two objects. This relationship between proximity and distance-to-landmark can be formalised as in Equation 1:

$$prox(T, L) \cong -dist(T, L) \quad (1)$$

In other words, Equation 1 states that the smaller the distance between the target and the landmark, the higher the proximity value for that target. This equation gives a good fit to people’s proximity judgments for targets in our experiment. For group x (the set of images for which there was no distractor object, just a target and the landmark), the correlation between $-dist(T, D)$ and people’s average proximity scores for target T was high ($r = 0.95$). The first graph in Figure 8 illustrates this correlation, comparing the average proximity value given by participants for each target in group x with the computed proximity value for each target in that group from Equation 1.

Equation 1, however, takes no account of the presence of a distractor object. To address the influence of a distractor object on judgements of proximity, we propose an alternative account, in which the judged proximity of a target to a landmark rises as the target’s distance from the land-

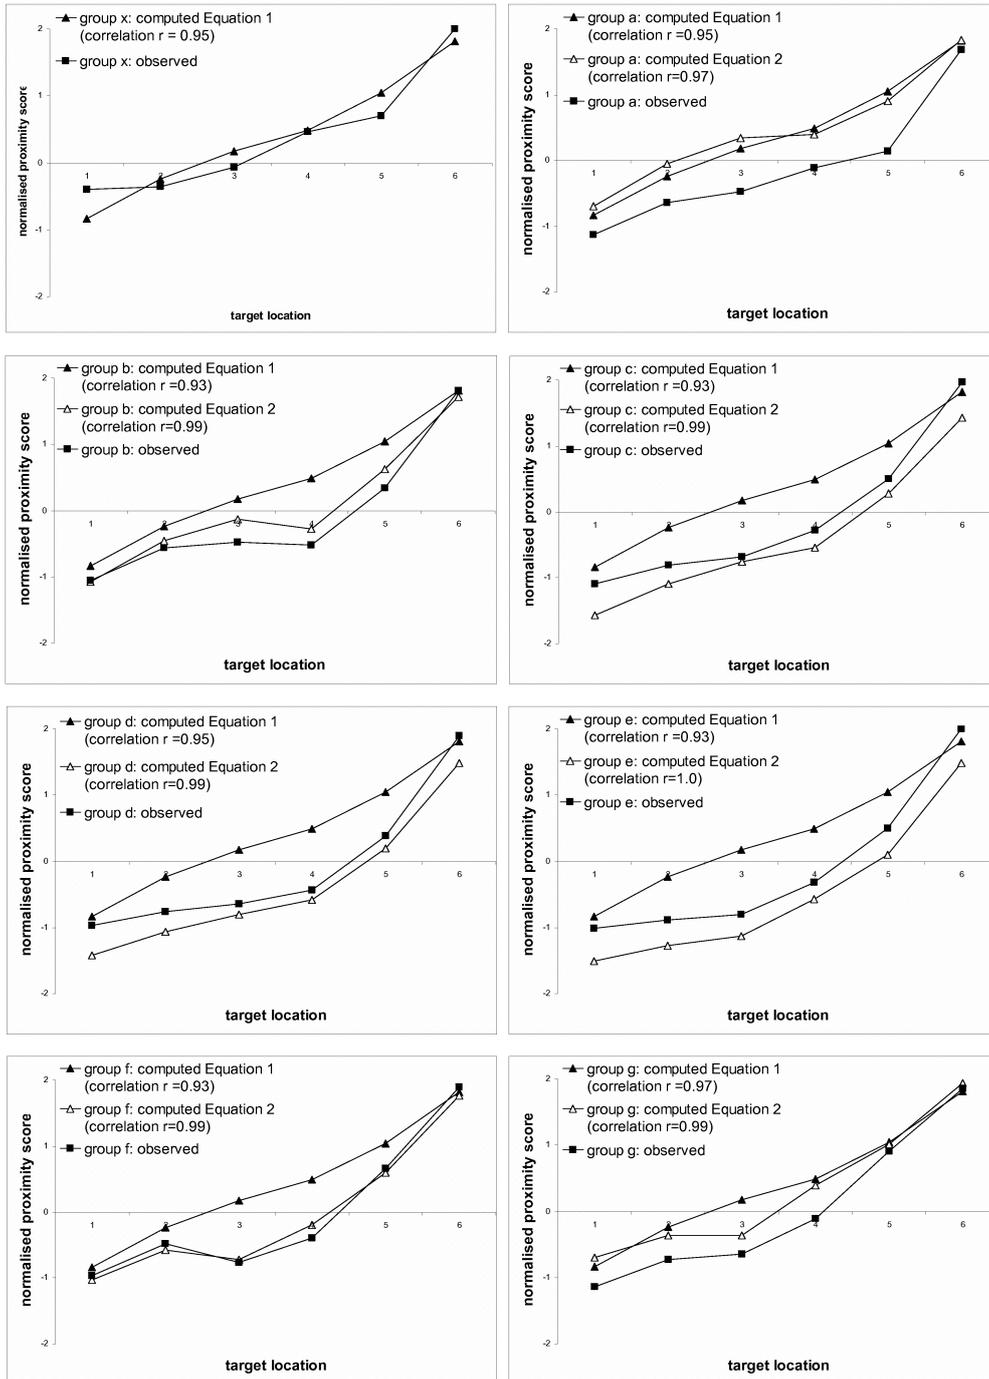


Figure 8: comparison between normalised proximity scores observed and computed for each group.

mark decreases (the closer the target is to the landmark, the higher its proximity score for the landmark will be), but *falls* as the target's distance from the distractor decreases (the closer the target is to the distractor, the lower its proximity score for the landmark will be). This relationship

between judged proximity, distance-to-landmark, and distance-to-distractor can be formalised as in Equation 2:

$$prox(T, L) \cong -dist(T, L) + dist(T, D) \quad (2)$$

Equation 2 states that if a target object is close to the landmark *and far from the distractor* it will have a high proximity score for that landmark. However, if it is close to the landmark but also close to the distractor, its proximity score will be lower.

The remaining seven graphs in Figure 8 assess this account by comparing the average proximity value given by participants for each target in the distractor groups *a* to *g* with the proximity value for each target in that group computed from Equation 1, and with the proximity value for each target computed from Equation 2. As these graphs show, for each group the proximity value computed from Equation 1 gives a fair match to people's proximity judgements for target objects (the average correlation across the seven groups is around $r = 0.93$). However, the addition of the distance-to-distractor term in the computation of proximity in Equation 2 significantly improves the correlation in each graph, giving an average correlation across the seven groups of around $r = 0.99$. We conclude that participants' proximity judgements for objects in our experiment are best represented by the model described in Equation 2, in which the proximity of a target to a landmark is a negative function of the target's distance from that landmark and a positive function of the target's distance from distractor objects.

Note that, in order to clearly display the relationship between proximity values given by participants for target objects, proximity computed in Equation 1 (using target-to-landmark distance only), and proximity computed in Equation 2 (using target-to-landmark and target-to-distractor distances) the values displayed in Figure 8 are normalised so that, across all groups and targets, the average proximity values given by participants have a mean of 0 and a standard deviation of 1, as do the proximity values computed in Equation 1 and those computed in Equation 2. This normalisation simply means that all values fall in the same region of the scale, and can be easily compared visually. This normalisation has no effect on the correlations obtained between the observed and computed proximity values.

5 Conclusions

This paper described a psycholinguistic experiment that investigated the cognitive representations underpinning spatial descriptions of proxim-

ity. The results showed that people's proximity judgments for objects in the presence of distractors can be modelled in a straightforward way using the relation described in Equation 2, in which proximity falls with the target's distance from the landmark, but rises with the target's distance from a distractor object. This means that if a target object is close to the landmark *and far from the distractor* it will have a high proximity rating for that landmark. However, if it is close to the landmark but also close to the distractor, its proximity rating will fall. This finding extends previous results on people's judgments of proximity for objects.

It's noticeable, however, that the match to people's responses obtained by Equation 2 for items in group *a* is less good than that obtained in any of the other groups. Of all the distractors, distractor *a* was closer to the target object than any other distractor. It may be that there is some other proximity or occlusion effect acting in people's judgements of proximity for items in group *a*. Future work will be necessary to clarify this point.

References

- L.A. Carlson-Radvansky and D. Irwin. 1993. Frames of reference in vision and language: Where is above? *Cognition*, 46:223–224.
- L.A. Carlson-Radvansky and D. Irwin. 1994. Reference frame activation during spatial term assignment. *Journal of Memory and Language*, 33:646–671.
- L.A. Carlson-Radvansky and G.D. Logan. 1997. The influence of reference frame selection on spatial template construction. *Journal of Memory and Language*, 37:411–437.
- A G Cohn, B Bennett, J M Gooday, and N Gotts. 1997. RCC: a calculus for region based qualitative spatial reasoning. *Geoinformatica*, 1:275–316.
- K.R. Coventry. 1998. Spatial prepositions, functional relations, and lexical specification. In P. Olivier and K.P. Gapp, editors, *Representation and Processing of Spatial Expressions*, pages 247–262. Lawrence Erlbaum Associates.
- K.P. Gapp. 1995. An empirically validated model for computing spatial relations. In *KI - Kunstliche Intelligenz*, pages 245–256.
- S. Garrod, G. Ferrier, and S. Campbell. 1999. In and on: investigating the functional geometry of spatial prepositions. *Cognition*, 72:167–189.
- W.G. Hayward and M.J. Tarr. 1995. Spatial language and spatial representation. *Cognition*, 55:39–84.
- A Herskovits. 1986. *Language and spatial cognition: An interdisciplinary study of prepositions in English*. Studies in Natural Language Processing. Cambridge University Press.
- J. Kelleher and F. Costello. 2005. Cognitive representations of projective prepositions. In *Proceedings of the Second ACL-Sigsem Workshop of The Linguistic Dimensions of*

Prepositions and their Use in Computational Linguistic Formalisms and Applications.

- G.D. Logan and D.D. Sadler. 1996. A computational analysis of the apprehension of spatial relations. In M. Bloom, P. and Peterson, L. Nadell, and M. Garrett, editors, *Language and Space*, pages 493–529. MIT Press.
- T Regier and L. Carlson. 2001. Grounding spatial language in perception: An empirical and computational investigation. *Journal of Experimental Psychology: General*, 130(2):273–298.
- C. Vandeloise. 1991. *Spatial Prepositions: A Case Study From French*. The University of Chicago Press.