

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

SpaceCase: A Model of Spatial Preposition Use

Permalink

<https://escholarship.org/uc/item/22q9532k>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN

1069-7977

Authors

Forbus, Ken
Lockwood, Kate
Usher, Jeffrey

Publication Date

2005

Peer reviewed

SpaceCase: A Model of Spatial Preposition Use

Kate Lockwood (kate@cs.northwestern.edu), Ken Forbus (forbus@northwestern.edu)
and Jeffrey Usher (usher@northwestern.edu)

Department of Computer Science, 1890 Maple Ave.
Evanston, IL 60201 USA

Abstract

We present SpaceCase, a computational model of spatial preposition use that combines geometric and functional influences. SpaceCase treats spatial preposition use as governed by evidential rules, each representing influences of particular factors. Our model is unique in relying on both automatically constructed visual representations from sketched input, and drawing our functional representations from an independently derived large knowledge base, both of which reduce tailorability. SpaceCase can account for the results of Feist and Gentner (2003), whose experiments about *in/on* judgments in native English speakers showed the influence of four factors: (1) geometry of the ground, (2) animacy of the ground (3) animacy of the figure and (4) function of the ground. SpaceCase also captures Feist and Gentner's (2001) result that memory for spatial relationships can be influenced by spatial language during encoding.

Keywords: Spatial language, Bayesian models

Introduction

Our aptitude for communicating and reasoning about space is key to our abilities to navigate, give directions, and to reason analogically about other subjects (Gentner, Imai, & Boroditsky, 2002). One way that we describe spatial scenes is through the use of prepositions like *in* and *on*. Traditionally, it was thought that the spatial preposition used to describe a scene depended solely on the geometric arrangement and properties of the objects in the scene. As described below, however, recent research indicates that non-geometric properties also play important roles. This raises a new difficulty for modeling the use of spatial prepositions: Representations of these other factors need to be created, ideally created independently from the spatial preposition model itself, to reduce tailorability. The same is true of course for the spatial representations used in stimuli given to models. Fortunately, progress in Artificial Intelligence has provided off-the-shelf knowledge bases and sketching systems with reasonable stand-ins for visual processing abilities. SpaceCase exploits both, by contrast with all previous models that we are aware of.

We begin by reviewing some of the evidence about spatial prepositions, focusing on the Feist and Gentner (2001, 2003) experiments. Next we describe SpaceCase, showing how it uses an independently-motivated sketch understanding system (sKEA, (Forbus & Usher, 2002)) and draws its representations from a large knowledge base (over 39,000 concepts, constrained by 1.2 million facts). We next show how SpaceCase can account for the Feist and Gentner

(2003) labeling experiment, including a sensitivity analysis that indicates the model is working for the right reasons. We show how SpaceCase models Feist and Gentner's (2001) retrieval results next. Finally, we discuss related and future work.

Psychological Evidence

The issue of how language and space interact has had a long history in cognitive science research. Early theories of spatial preposition use claimed that people assigned spatial prepositions based on the geometry of a visual scene. However, more recent work has shown that the use of spatial prepositions is influenced by a variety of functional factors in addition to the geometry of the situation. Factors such as context (Coventry, 1999; Herskovitz, 1986), functional relationships between the objects (Carlson-Radvansky et al., 1999; Coventry et al., 1994; Vandeloise, 1994), and control relationships (Feist & Gentner, 2003; Garrod et al., 1999) also influence how we use prepositions in everyday language (see Coventry & Garrod, 2004 for a review).

We focus here on modeling the results of Feist and Gentner (2003), for concreteness. They examined the role of four factors in *in/on* determinations in visual scenes involving two objects, a *figure* (located object) and a *ground* (reference object): (1) the geometry of the ground, (2) the animacy of the ground, (3) the animacy of the figure, and (4) the functional role of the ground. They found that all of these factors were involved in determining whether subjects would describe the figure as *on* or *in* the ground. Specifically, high curvature is more likely to lead to *in*, and low curvature more likely to be associated with *on*. If the ground were animate (a hand, for instance), *in* was more likely to be used, whereas if the figure is animate, *on* was more likely to be used. Moreover, subjects were more likely to use *in* than *on* if they were told that the ground was a container (say, a bowl) than when they were told it was something else (e.g., a plate), even with the same curvature. Can language affect how spatial relations are processed? Feist and Gentner (2001) showed that giving subjects a sentence involving a spatial preposition while viewing a scene affects how that scene is stored in memory. That is, suppose a subject is shown Figure 1(right) below while being told "The puppet is *on* the table.", as part of a larger set of stimuli. When later asked if they had seen Figure 1(left), which was not shown to them earlier, subjects who had heard *on* during encoding were more likely to incorrectly report that they had seen it. This suggests that information from multiple modalities (visual and linguistic)

can be combined into a single representation, such that one influences the other.



Figure 1: Sample stimuli from Feist and Gentner (2001)

The SpaceCase Model

SpaceCase has been developed to model phenomena such as those described above. We assume that the assignment of spatial prepositions rests on the knowledge and skills that people bring to bear in other spatial tasks. Spatial prepositions encode a combination of geometric and functional properties, making them both detectable in visual scenes and able to provide information about what possibilities are relevant in that scene, when detected. For instance, the distinction between *on* and *in* in English includes an aspect of location control, with more control when *in* is used than when *on* is used. We assume that there are multiple, situation-specific criteria that determine when it is appropriate to use one term over another. We view these criteria as evidential, in that they tend to suggest, rather than uniquely determine, answers. Thus we describe our model in terms of evidence rules, which, given a situation, provide levels of belief about how prepositions should be assigned.




Figure 2: Sketched input for stimuli of Figure 1

Before we present the specific rules that comprise SpaceCase’s assignment criteria, we first must describe how we compute the input properties that are needed, according to what is known about the stimulus. To reduce tailorability, we input stimuli into the model as hand-drawn sketches. We use the *sketching Knowledge Entry Associate* (sKEA), an open-domain sketch understanding system (Forbus & Usher, 2002) for this purpose. A sketch in sKEA consists of a set of glyphs. Each glyph has *ink* and *content*. The ink consists of the actual strokes drawn. The content is the entity that the glyph represents. For example, Figure 2 illustrates how the stimulus of Figure 1 is presented to sKEA.

sKEA provides a rich set of visual processing capabilities. We focus here on those used by the current version of SpaceCase. sKEA uses the ink to compute geometric properties of a glyph, such as the curvature of the object

depicted. It also uses the ink to compute relationships between glyphs, including qualitative topological relationships (i.e., Cohn’s RCC8 (1996) vocabulary) such as touching.

The content of a glyph is assigned to be an instance of one or more collections drawn from Cycorp’s Cyc knowledge base¹. This is specified by the stimulus drawer when the glyph is specified. For example, when drawing a bowl, the user first lays down the ink and then labels the glyph as an instance of **Bowl-Generic** using tools in sKEA. This identification allows us to infer additional information about the figure and the ground in the sketch. For example, the bowl mentioned above is inferred to be a **Container**, via an inheritance relationship. This provides a model of the functional information that is needed in assigning spatial prepositions. We are not claiming that the KB contents are a psychologically accurate model of human knowledge in detail; all we are relying on are very high-level, coarse distinctions. What is crucial, however, is that the vast majority of the KB contents were independently developed by other researchers. Similarly, most of sKEA’s visual processing abilities predate this project – only curvature was added to support this project, and the same curvature computations will be used in our subsequent simulations.

SpaceCase collects the following information from a given sketch: (1) the geometry of the ground (2) the animacy of the figure (3) the animacy of the ground and (4) the function of the ground. The geometry of the ground is computed from the properties of the ink in the sketch and the other three factors are collected via inference about the kinds of entities involved. Animacy of the ground and animacy of the figure are binary values. Geometry of the ground is represented as either high curvature, medium curvature, or low curvature. The function of the ground is qualitatively assigned to one of three categories: strongly functions as a container, weakly functions as a container, or functions as a surface. These distinctions are taken directly from the human-subjects data. After all of the information is gathered from a scene, it is fed as evidence into a Bayesian updating algorithm which assigns a probability that either of the prepositions accurately describes the scene. We use Everett’s (1999) evidential rule engine, which in turn uses Pearl’s (1986) hierarchical updating algorithm. Evidence contributes to the support for a preposition based on the likelihood of that evidence for that preposition. Likelihood is defined as:

$$\lambda_n = P(e|H) / P(e|\neg H)$$

After all of the evidence is considered and the model converges, if the likelihood of any preposition exceeds a threshold, that preposition is proposed as the correct descriptor for the scene. At present, the only options the model has are: *in-ContGeneric*, *on-Physical*, and *other-preposition*. The first two are formal predicates which are used in the knowledge base for covering a very large set of specialized cases, defined by a hierarchy of specialized

¹ sKEA’s KB uses content drawn from the Cyc KB, plus our own material related to analogy and qualitative reasoning.

predicates. *in-ContGeneric* has thirteen specializations, including different levels of location control (e.g., open versus closed container). *on-Physical* has two specializations, corresponding to floating on a liquid and particles strewn over a surface.

Currently SpaceCase has a total of ten evidential rules. Three of the rules are used to describe the support relationship between the ground and the figure. This is an example of using sKEA to provide perceptual information to the system, since the triggers for these rules depend on the visual relationships between the sketched figure and ground relationships. The three support relationship rule variables are:

- * Figure-completely-supported-by-ground
- * Figure-partially-supported-by-ground
- * Figure-not-supported-by-ground

They trigger based on how many of the figures bottom edge points intersect with the grounds edge points. The first rule (complete support) increases the likelihood that the figure is either *in* or *on* the ground. The second and third rules (no support or only partial support) increases the probability that another preposition would be more appropriate to describe the scene.

The other seven rules in our system collect the evidence necessary to make the *in/on* judgments as in the Feist and Gentner (2003) experiments. Therefore, they relate to the factors studied (animacy, ground function and ground geometry). The variables that represent these likelihoods are:

- * ground-high-curvature
- * ground-medium-curvature
- * figure-animate
- * ground-animate
- * ground-function-container-strong
- * ground-function-cointainer-weak
- * ground-function-slab

Based on the results from the human subjects trials, the ground high and medium curvature rules increase the likelihood that the figure is *in* the ground. The ground animate rule and the ground function container rules also increase the likelihood of *in*. The figure animate and ground function slab rules increase the likelihood that the figure is *on* the ground. The curvature rules are triggered by the curvature that sKEA computes from the digital ink for the glyph that represents the ground. The other rules are all triggered by inferences made from the knowledge base, using the concept instance information asserted when the sketch is created in sKEA.

When each rule is triggered, it creates an evidence element that contains the name of the preposition to update (*in*, *on* or *other*) as well as the amount by which to update its likelihood. The evidence values associated with each rule are parameters of the model. The values chosen were based on the pattern of results found in the human subjects experiments. For example, the function of the ground, for people, has a much stronger influence on the number of *in*

responses than the curvature of the ground. Therefore the ground function strong rule increases the likelihood of an *in* response by a greater value than the ground high curvature rule. The values of the likelihood update variables in the current incarnation of SpaceCase are described in Table 1.

As we will see below, SpaceCase is not terribly sensitive to the specific values of these parameters. As long as their ordinal relationships fit the pattern of results found earlier, SpaceCase's answers will accurately model the data.

Variable Name	Preposition	Likelihood
figure-complete-support	in/on	5
figure-partial-support	other	10
figure-no-support	other	30
ground-high-curvature	in	3
ground-medium-curvature	in	2
figure-animate	on	3
ground-animate	in	5
ground-function-container-strong	in	10
ground-function-container-weak	in	5
ground-function-slab	on	10

Table 1: Evidential parameters currently used in SpaceCase

Labeling Experiments

In Feist & Gentner (2003), subjects were shown simple pictures of figures located on grounds. The different factors (geometry, animacy, and function) were varied in the different stimuli. Subjects were given sentences the *<figure>* is *in/on* the *<ground>* (where *<figure>* and *<ground>* were replaced with appropriate labels) and asked to indicate which preposition best fit the situation displayed.

We sketched the different stimuli using sKEA and identified the items as instances of appropriate concepts from the Cyc knowledge base. The sketches were labeled to represent the same entities as in the human trials and represented the same combinations of factors as the human subjects saw, as described above.

There were a total of 36 stimuli. There are two choices for figure {firefly, coin}, 6 choices for ground {bowl, dish, plate, slab, rock, hand}, and three different curvatures for each ground item {low, medium, high}. For all trials in the original experiment, it was assumed that the relationship ground-supports-figure held, and that this would be obvious to human subjects from the stimuli. In our sketched stimuli however, we determined the support relationship as outlined in the rules in the previous section. Figure 3 below shows an example of a stimulus from the original experiment on the left and our sketched stimulus for the same trial on the right. This set of stimuli represents the condition where the figure is a firefly, the ground is a dish, and the ground exhibits medium curvature.



Figure 3: Example stimuli from the original human subjects experiments (left) and the sketched stimuli used as input to our model (right)

One difference between the original experiment and our model of it was that the original stimuli were 3-D renderings and our sketches are 2-D, due to limitations of sKEA. We have not seen any issues arising from this difference.

Results

SpaceCase was consistent with human subjects on all 36 experimental stimuli for the values of the parameters given above. Importantly, SpaceCase is not overly sensitive to the specific values chosen: As long as parameters reflect the relative strengths of the factors as found by Feist & Gentner (2003), the correct results are generated. We determined this via a series of sensitivity analyses, looking at how the results changed when parameters were varied. Because this is a large space, we have focused on two dimensional subspaces of these parameters at a time, with the other parameters keeping the values from Table 1. Below is an example plot. The lighter gray squares indicate parameter settings where SpaceCase’s answers are consistent with human subjects, and the darker squares indicate inappropriate results:




Figure 4: Sensitivity Analysis: figure-animate versus ground-container-strong

Examining why failures occur can lead to interesting insights. For example, the firefly-hand stimulus proves to be particularly interesting, since both are animate. Subjects were more likely to say “the firefly is *in* the hand”, and for many parameter values, SpaceCase does as well. However, when the figure-animate parameter is set sufficiently higher than the ground-animate parameter, we get instead “the firefly is *on* the hand”. Feist & Gentner (2003) found a

much larger positive effect on *on* usage for ground animacy than the negative effect on figure animacy. Thus when SpaceCase’s parameters violate constraints found by psychological experimentation, it fails, suggesting that it is failing for the right reasons.

Retrieval Experiment

Feist & Gentner (2001) describes a series of experiments where human subjects were shown pictures that were ambiguous with respect to whether or not the figure was *on* the ground, with or without sentences that described the scene involving spatial prepositions. For example, in Figure 5 (initial picture), it is ambiguous as to whether the block is actually on the building, but the experimental group might also be asked to rate the applicability of the sentence “The block is on the building.”, while being shown the picture. (Several variations were used to rule out alternate explanations, e.g., simply concentrating on the picture, and whether language without prepositions would have an effect.) In the retrieval phase, subjects would then be shown both the original picture and two variations, a plus variant which unambiguously satisfies the spatial preposition, and a minus variant, which is even worse with regard to exhibiting that spatial preposition than the original picture. Subjects tended to believe that they had seen the plus variant when they also were exposed to the appropriate spatial preposition during encoding, thus illustrating that language can affect the encoding and memory of spatial relations in visual stimuli.




Figure 5: Example of Feist & Gentner (2001) stimuli,

To model these results, we recreated all of the stimuli (original pictures as well as plus and minus variants) using sKEA. Because SpaceCase currently only handles *on/in* distinctions where the figure is supported by the ground, we eliminated stimulus sets that involved other pairs of prepositions. For retrieval we used Forbus *et al*’s (1994) MAC/FAC similarity-based reminding model. MAC/FAC’s memory consisted of the sketched versions of what the subjects saw.

For the first experiment, the input sketches were run against our model to determine the applicability of the preposition to the initial situation. This is similar to one experiment run against human subjects to show the applicability of the preposition to each variant of the initial stimuli. These results are summarized below

initial sketch	0.363
plus variant	0.859
minus variant	0.2428

Table 2: Applicability of the spatial preposition on for each of the input stimuli as judged by SpaceCase, all results are averaged across all sets of stimuli.

These results are consistent with the human-subject trials where the plus variant was given the highest applicability rating, the initial sketch an in-between rating, and the minus variant the lowest applicability rating. These results also pointed out some weaknesses in our current version of SpaceCase which will be addressed in the next iteration. For example, one of the stimulus sets involved a block on top of a building. For this particular stimulus the rule for the ground acting as a container fired since a building can be inferred to be a container in the KB. Clearly, a building can be a container, but in this particular case (the block on the roof of the building) the support relationship is not one of containment. SpaceCase needs to use visual properties as well as conceptual properties, to ensure that containment is actually occurring in a given picture.

To provide a baseline of comparison for the next group of experiments, each case library was probed with the initial sketch to see which sketch would be retrieved. This is similar to the control condition in the retrieval experiment where human subjects were not given the spatial language sentence along with the picture. In all cases, the initial sketch was retrieved, as expected.

For the next experiment, the sketches were run again with the model, but we added the formal equivalent of spatial preposition indicated by the sentence into the representation of each sketch. Doing this led to retrieving the plus variant of the sketch rather than the initial sketch. This is consistent with the human-subjects results that subjects with the spatial language sentence were much more likely to false alarm to the plus variant.

Related Work

Many of the early spatial preposition models are based solely on geometric properties of scenes (Logan & Sadler, 1996; Regier, 1996; Gapp, 1995). Others, such as WordsEye (Coyne & Sprout, 2001) uses hand-coded databases of objects and representations that are crafted particularly for it. By contrast, our model uses a pre-existing, independently developed representation system and a sketch understanding system developed for other purposes, and whose development continues to be constrained by multiple tasks.

More recently, computational models have begun to address the role of functional information. Reiger, Carlson, and Corrigan (2005) are currently extending the Attentional Vector Sum (AVS) model (Regier & Carlson, 2001) of spatial prepositions to account for extra-geometric information. The initial AVS model created a vector-sum representation of the direction of the located object (figure)

relative to the reference object (ground) weighted by the amount of attention paid to the point on the reference object. This model used only geometric properties to assign prepositions and was tested on afunctional geometric shapes. The new version of AVS was modified to focus attention on the functionally important parts of objects (and thus weight them more heavily). The results for this model are very promising, e.g. its' ability to recreate the results of Carlson-Radvansky et al. (1999). We think it is an interesting and complementary approach as ground-functionality is isolated in the functionally important parts of the object.

Coventry, Cangelosi, Joyce, and Richards (2002; Joyce, Richards, Cangelosi, & Coventry, 2003) are currently developing a model for spatial language comprehension and production that is based on their functional geometric framework. Objects are encoded with "what + where" information and is fed into a predictive, time-delay connectionist model.

Discussion

We have presented SpaceCase, a computational model of human use of spatial prepositions. SpaceCase assigns *in* and *on* relationships in a manner consistent with Feist & Gentner's (2003) results, and exhibits the influence of spatial language during encoding as found in Feist & Gentner (2001). These results are made possible by sKEA, our open-domain sketching system which provides a high-level model of visual processes and the input to our model. Both geometric and conceptual data from the sketched input are used as evidence in a probabilistic updating system to choose between prepositions.

Our immediate goal is to expand our model to include more prepositions and a wider range of more complex inputs. In addition to being able to determine *in/on* for more complex input, we would like the model to be able to understand other prepositions such as *over*, *under*, *above*, *near*, *next to*, and *around*. We would also like to be able to correctly use spatial prepositions in more complex scenes. For example, a commonly used example is that if a stack of books is on a table, then it is also correct to say that the top book in the stack is "on" the table. However, if a jar with a lid is on the table, very few people would say that the lid is on the table.

Expansion will necessarily involve the development of more rules for SpaceCase. Complex scenes may also require a more complete understanding of concepts like containment. Right now, SpaceCase only knows whether an item typically functions as a container. There is evidence that it is also important to determine how well a ground object is fulfilling its role (exerting location control over the figure). In particular, projective prepositions may require some naïve physics knowledge. One other difficulty we foresee is that the simple Bayesian updating algorithm might not suffice for more complex scenes. If this turns out to be the case, we plan to experiment with multiple levels of competing agents.

SpaceCase might also provide a testbed for examining other areas of spatial language. Competing theories of spatial language can be tested by manipulating the rules of the system. For example, there is a debate on how exactly non-geometric features like functionality interact with geometric factors in scenes to determine the preposition used. For now we have focused on native English speakers, however, one area we are interested in is cross-linguistic spatial preposition use. One experiment would be to see if the same basic inputs to the model could be used to also model preposition use in another language by merely changing the available prepositions and the rules for how evidence is assigned.

Acknowledgements

This research was supported by grants from the Artificial Intelligence Program of the Office of Naval Research and the ROLE program of the National Science Foundation.

References

- Carlson-Radvansky, L.A., Covey, E.S., & Lattanzi, K.M. (1999). "What" effects on "where": Functional influences on spatial relations. *Psychological Science*, 10, 516-521.
- Cohn, A.G. 1996. Calculi for Qualitative Spatial Reasoning. In Calmet, J., Campbell, J.A., and Pfalzgraph, J. (Eds.) *Artificial Intelligence and Symbolic Mathematical Computation*. LNCS 1138. Springer Verlag, 124-143.
- Coventry K.R., Carmichael, R., & Garrod, S.C. (1994). Spatial prepositions, object-specific function, and task requirements. *Journal of Semantics*, 11, 289-309.
- Coventry, K.R. (1999) Function, geometry, and spatial prepositions: Three experiments. *Spatial Cognition and Computation*, 1, 145-154.
- Coventry, K.R., Cangelosi, A., Joyce, D.W., & Richards, L.V. (2002). Putting geometry and function together: Towards a psychologically-plausible computation model for spatial language comprehension. *The Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (p.33). Mahwah, NJ: Erlbaum Associates Inc.
- Coventry, K. R., & Garrod, S. A. (2004). *Saying, Seeing, and Acting: The Psychological Semantics of Spatial Prepositions*. New York: Psychology Press.
- Coyne, B. and Sprout, R. 2001 WordsEye: An automatic text-to-scene conversion system. *SIGGRAPH 2001*.
- Everett, J.O. (1999). Topological Inference of Teleology: Deriving Function from Structure via Evidential Reasoning. *Artificial Intelligence*, 113(1-2).
- Feist, M.I. & Gentner, D. (2001). An influence of spatial language on recognition memory for spatial scenes. In J.D. Moore & K. Stenning (Eds.) *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*. (pp 279-284)
- Feist, M. I., & Gentner, D. (2003). Factors Involved in the User of In and On. *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*.
- Forbus, K. and Usher, J. (2002). Sketching for knowledge capture: A progress report. IUI'02, January 13-16, 2002. San Francisco, California.
- Forbus, K., Gentner, D., and Law, K. (1994). MAC/FAC: A Model of Similarity-Based Retrieval. *Cognitive Science* 19, 141-205.
- Gapp, K.P. (1995). Angle, distance, shape and their relationship to projective relations. In J.D. Moore & J.F. Lehman (Eds.). *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. (pp.112-117). Mahwah, NJ: Lawrence Erlbaum Associates Inc.
- Garrod, S., Ferrier, G., & Campbell, S. (1999). In and on: Investigating the functional geometry of spatial prepositions. *Cognition*, 72, 167-189.
- Gentner, D., Imai, M. and Boroditsky, L. (2002). As Time Goes By: Evidence for two systems in processing space>time metaphors. *Language and Cognitive Processes*, 17 537-565.
- Herskovits, A. (1986). *Language and spatial cognition: An interdisciplinary study of the prepositions in English*. Cambridge, England: Cambridge University Press.
- Joyce, D. W., Richards, L. V., Cangelosi, A., & Coventry, K. R. (2003). On the Foundations of Perceptual Symbol Systems: Specifying Embodied Representations via Connectionism. In F.Dretje, D. Dorner, & H. Schaub (Eds), *The logic of cognitive systems: Proceedings of the Fifth International Conference on Cognitive Modeling* (pp. 147-152). Bamberg: Universitäts-Verlag Bamberg.
- Logan, G.D. & Sadler, D. D. (1996). A computational analysis of the apprehension of spatial relations. In P.Bloom, M.A. Peterson, L.Nadel, & M.F. Garrett (Eds). *Language and Space* (pp.493-530). Cambridge, MA: MIT Press.
- Pearl, J. (1986). On Evidential Reasoning in a Hierarchy of Hypotheses. *Artificial Intelligence*, 28, 9-15.
- Regier, T. (1996). The human semantic potential: Spatial language and constrained connectionism. Cambridge, MA: MIT Press.
- Regier, T. & Carlson, L.A. (2001). Grounding Spatial Language in Perception: An Empirical and Computational Investigation. *Journal of Experimental Psychology*, 130:2, 273-298.
- Regier, T., Carlson L., & Corrigan, B. (2005). Attention in spatial language: Bridging geometry and function. In L. A. Carlson & E. van der Zee (Eds.), *Functional features in language and space: Insights from perception, categorization, and development*. Oxford: Oxford University Press.
- Vandeloise, C. (1994). Methodology and Analysis of the Preposition In. *Cognitive Linguistics*, 5, 157-184.