Domain Specific Understanding of Spatial Expressions

Debra T. Burhans, Rajiv Chopra and Rohini K. Srihari *
Department of Computer Science & CEDAR
State University of New York at Buffalo
520 Lee Entrance Suite 202
Buffalo, New York U.S.A. 14228-2567
(716) 645–6164
{burhans,rchopra,rohini}@cedar.buffalo.edu

Abstract

This paper explores the notion of spatial expression understanding relative to a domain. A model for spatial expression understanding is proposed, and ways in which domain specific information can be incorporated at different stages of the model are discussed. We describe the domain of newspaper captions and show that spatial information in this domain is primarily expressed in the open class and by projective spatial prepositions. A discussion of types of open class spatial information and the role of *a priori* knowledge about a domain in interpreting this information is presented. Many of the above observations have been incorporated into *PICTION*, a system that extracts spatial and visual information from newspaper captions for collateral-based image understanding. The representation and processing of spatial information in *PICTION* is presented as an application of the proposed model. The paper concludes with a discussion of our approach, and the applicability of the model to other domains.

1 Introduction

A spatial expression conveys information about objects in space, their locations, or the frame of reference from which objects are viewed. Understanding a spatial expression includes representing the information it conveys so that it can be used for spatial reasoning. Domain specific understanding of spatial expressions means using information known *a priori* about a domain to identify and interpret spatial expressions. *Domain* here refers to a specific application or research task. Examples include newspaper caption understanding, robot navigation, three-dimensional modeling of spatial prepositions, and model-based scene analysis.

Examples of *domain specific* information are a fixed or default frame of reference, syntactic restrictions on the input text, and a predefined set of objects and spatial relations of interest. In contrast to domain specific knowledge, *scene specific* information is conveyed by a particular text. Scene specific information is present in *as seen from above* and *to Mary's left*. The meaning of *to Mary's left* depends on the frame of reference relative to which *left* is interpreted. *As seen from above* gives information about the frame of reference to be used in interpreting contextually related spatial expressions. The problem of determining what *is* contextually related is one of discourse analysis, and is not addressed in this paper.

Domain and scene specific information together comprise the *background* with respect to which a spatial expression is understood. When a person visualizes a spatial scene (in this case one described linguistically), he or she does so with respect to a unique frame of reference, whether it is imposed by the description or selected dynamically by the person. This frame of reference, minimally, comprises the *background* from a cognitive point of view. The distinction we make between domain and scene specific knowledge is based not on cognitive principles, but on considerations that arose in developing a computer application that involves spatial expression understanding.

This paper presents a computational model for spatial expression understanding and explores the role of domain specific knowledge in this task. The newspaper caption domain is described, and types of spatial information in this domain are identified. The prevalence of spatial information expressed in the open class in this domain is demonstrated. The role of the open class in conveying spatial information, and ways in which domain and context specific assumptions affect the interpretation of open class spatial expressions, is discussed. The representation and processing of spatial expressions in *PICTION* [Srihari, 1995] is presented as an application of the model to the newspaper caption domain. The paper concludes with a discussion of our approach and its applicability to other domains.

2 A Model for Spatial Expression Understanding

Given natural language text, what are the tasks involved in understanding the spatial information it conveys? A paradigm for understanding spatial expressions comprises four functionally distinct tasks (Fig. 1).

*This work was supported in part by ARPA Contract 93-F148900-000.

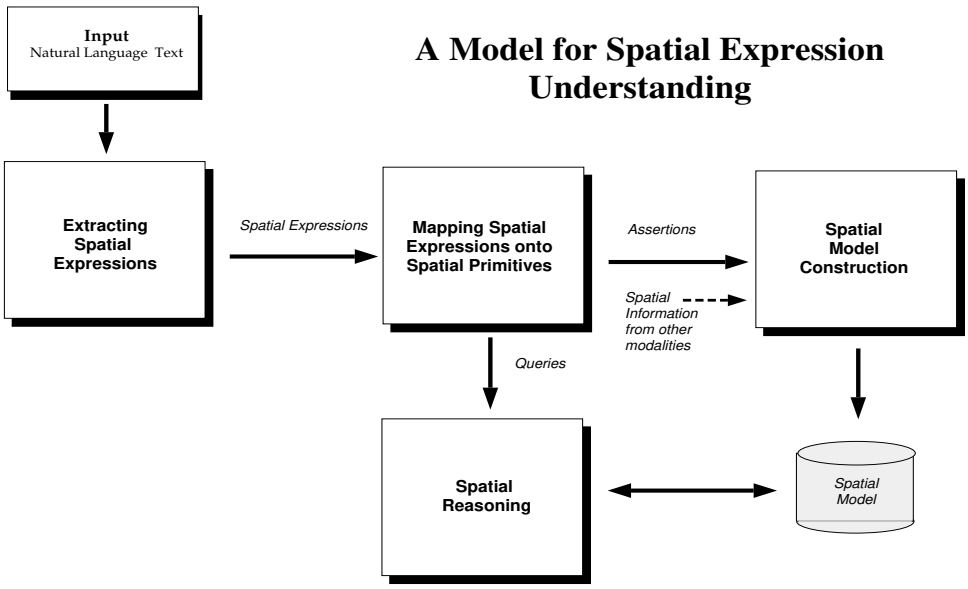

Figure 1: A Model for Spatial Expression Understanding

Extracting spatial expressions from text involves identifying and representing spatial information contained therein. Domain specific knowledge can be used to predetermine a set of objects and spatial relations of interest. In *PICTION*, for example, the objects of interest are human faces and some rectangular objects (*e.g.* television sets, picture frames). Hence, information about other objects and spatial relations between them is not extracted from the input text. This approach presupposes the ability to determine the object class of common and proper nouns in the input text, a problem that will be discussed in the next section of the paper. A system that processes and represents *all* spatial information will not employ any domain specific information at this stage.

Spatial expressions are mapped onto spatial primitives that comprise the objects and spatial relations in the underlying spatial knowledge representation. A spatial expression can occur in the form of an assertion or a question. An assertion is mapped to a non–empty set of spatial primitives that are structured as a set of assertions. Similarly, a question is mapped to a set of queries. When spatial expressions extracted from text are already in the form of assertions or queries composed from spatial primitives, this mapping is unnecessary.

If domain specific knowledge has been used to constrain the set of spatial primitives or the set of spatial expressions that are input to this mapping function, then the mapping is necessarily domain specific.

In *PICTION*, mapping the spatial expression *from left, the president, vice-president and treasurer of the Glee Club* onto a set of assertions involves the following:

- Associate each person mentioned with a *human-face* object, for example, associate *H1* with the president, *H2* with the vice-president and *H3* with the treasurer.
- Construct assertions about the presence of these objects in the scene: {(*in-picture H1*), (*in-picture H2*), (*in-picture H3*)}
- Represent *from left* as a set of *left-of* relations: {(*left-of, H1, H2*), (*left-of, H2, H3*)}

The above mapping employs domain specific knowledge that has been used to fix a frame of reference, define the semantics of *from left*, and determine the set of spatial primitives.

A **spatial model** is a set of assertions relating to a particular scene. In some cases the assertions themselves constitute the underlying representation, while in others there are definitions of spatial primitives that *ground* the assertions in a perceptual domain. The task of spatial model construction includes verifying assertions and updating the spatial model to reflect new or modified information.

The application for which spatial expression understanding is required influences the choice of the knowledge representation used for the spatial model. For example, in machine translation, this representation might be a language–independent set of spatial concepts. In *PICTION*, the fact that spatial information is grounded in the two–dimensional image domain has been used in developing definitions for the spatial primitives. These definitions are presented later in the paper.

Queries composed of spatial primitives are sent to a **spatial reasoner** that makes use of the spatial model and the definitions of spatial primitives. While not all systems that process spatial expressions employ a spatial reasoner, from a cognitive point of view spatial expression understanding seems to include being able to reason about the spatial scene described. We have not implemented spatial reasoning in the current version of *PICTION*.

The above discussion demonstrates ways in which domain specific knowledge can be applied in performing, or obviating the need to perform, the tasks defined in

Spatial Information in Newspaper Caption	
Type of Information	*Percentage of captions*
spatial prepositions (topological)	2
projective prepositions	53
left	35
right	9
center	4
misc. (*e.g.*, in front)	4
order of mention (implicit spatial relation)	13
spatial verb	50
spatial adjective	3
no spatial information	14

Table 1: The above categories are not mutually exclusive. This data is based on 262 captions taken from *The Buffalo News* and *The New York Times*.

this model. Incorporating domain specific information into an application of the model presupposes that the domain has been characterized in terms of the types of spatial information it embodies, and the ways in which that information will be used.

3 The Newspaper Caption Domain

The purpose of a newspaper caption is to help a reader identify objects in an image. This is accomplished by associating linguistic references with visual references. Visual data in this domain is two-dimensional. Every image has an orientation: top, bottom, left and right have fixed meanings. Projective spatial terms in captions are interpreted relative to this orientation and the frame of reference of the newspaper reader. *Left* means *the left hand side of the picture as viewed by a reader.* Captions refer to fixed scenes, hence, scene understanding does not entail understanding of temporality or motion.

Based on the sample of captions to which Table 1 refers, several observations about spatial language in captions can be made.

- Topological spatial prepositions are nearly absent.
- Projective prepositions are used in about half of the captions.
- *Order of mention* is a spatial relation, defined according to the following convention (stated here with respect to the object class *people*):
  if people mentioned in a caption (and predicted to be in the associated image) are not well known, and can not be distinguished based on gender or age, then the order of mention corresponds to a left to right ordering of the associated persons in the image.
- Verbs are used to convey spatial information almost as frequently as prepositions.
- Adjectives are rarely employed in spatial description.
- When spatial terms are absent from a caption, it is because the corresponding photograph depicts a single, salient object or a scene in which no specific object can be identified.

The prevalent use of open class spatial terms in this domain led us to investigate types of spatial information found in the open class.

4 Spatial Information in the Open Class

In the newspaper caption domain, spatial information is frequently expressed in the open class. *Spatial* verbs may be viewed as complex bundles of information that entail spatial relations (flank), instruments (cut), topological object properties (border), or other objects (floss). Talmy [1985] has described this implicit inclusion of information in verbs as conflation. Dorr and Voss [1993] define *lexically implicit* spatial relations as those implicitly entailed by lexical terms, corresponding to Talmy's conflated *PATH* and *LOCATION*.

Open class words have three important functions in conveying spatial information:

- They convey non-topological spatial information that can not be expressed in the closed class.
- They refer to objects in space in a variety of ways, leading to the problem of reference.
- They convey closed class, topological information, often in conjunction with other spatial information.

Topological spatial properties are defined as those which are spatially invariant, *i.e.*, they do not depend on a frame of reference or object properties such as size and shape.

Information about object orientation, absolute or relative location, and other potentially **non-topological spatial properties** is important for spatial expression understanding. Lang *et al.* employ this type of information in the representation of objects and spatial relations in *OSKAR* [1991]. In the caption domain, the phrase *Joseph Biden, back row at left* denotes a specific area of an image, and an object contained in that area (the face of Joseph Biden).

Object location may also be specified by a verb. In *John and Peggy Stern hold their adopted children*, the location of the children is given relative to the location of the parents. Consider the quite different spatial configuration described by *John and Peggy Stern wave to their adopted children*, where only the verb has changed: in this second example, there is no contact relation between the children and parents. Siskind [1994] has described the importance of the notions of support, contact and attachment in understanding spatial terms.

The verb *hold* has a complex spatial semantics, a small fragment of which is:

- if the object being held is a physical object, there is a non–empty set of spatial relations between it and the subject performing the act of holding, including contact
- if the subject is a person, and the object is not a body part of the subject, the default location for the held object is *in* the hand(s) of the subject (understanding this requires a spatial semantics for *in hand*)

Information about the meaning of *hold* when the subject is a person is important in *PICTION*. Such considerations have led us to to develop semantic representations for a small set of spatial verbs found frequently in captions. These representations are domain independent, but the choices of which verbs to represent and how much/which parts of the spatial semantic information to include have been determined by our domain.

Identifying the reference of an object name or description is important. Zernik and Vivier [1988] state that language relating to spatial scenes has two functions: object location, and object recognition. In their definition, object recognition refers to the problem of reference. Work in understanding spatial prepositions has demonstrated the importance of knowing the properties of the objects involved in the spatial relation denoted by the preposition (*e.g.* the boat in the water *vs.* the fish in the water). Object recognition is important if domain specific information has been used to constrain the set of objects of interest, as has been done *PICTION*. Properties of objects are often related to the class to which an object belongs. In the case of common nouns, determining object class may be accomplished using a lexical ontology. In *PICTION*, we have used Wordnet [Beckwith *et al.*, 1991] for this purpose. We have marked the superconcept for the class of humans, as well as several concepts representing rectangular objects or objects that appear as arguments to spatial verbs, as objects of interest. Determining the class of (unknown) proper nouns is an area of research whose results we hope to incorporate in *PICTION*.

Over half of the captions referred to in Table 1 included information (titles, descriptions in relative clauses) that allowed an object to be referred to in at least two different ways. The caption *Mary Retton, right, poses with Toukie Smith, the sister of the late fashion designer Willi Smith, ...* is a typical example. *Toukie Smith* and *Willi Smith's sister* refer to the same individual in this caption. The caption indicates that Toukie is to the left of Mary Retton. If we neglect the problem of reference, we will be incapable of answering the query, *Is Willi Smith's sister next to Mary Retton?* despite the fact that this information has been given to the system.

In *PICTION*, there are two different problems of reference. One is determining which linguistic references refer to the same object, and the other is the problem of linking each unique linguistic referent with its visual counterpart.

Topological information may be encoded in verbs, adjectives and nouns. Examples include *border, encircle, frame, surrounded (by)*, and *picture (of)*. Two percent of the captions in our sample refer to framed pictures of people. Consider the caption *Sharon Bottoms and April Wade hold picture of Tyler Doustou.* The fragment *picture of Tyler ...* expresses a topological relationship of containment between the border of the picture and the face of Tyler.

Determining the meaning of an open class spatial expression is a complex task that subsumes understanding closed class spatial expressions. Identifying and representing this information is an important part of lexical semantics, as other researchers have recognized [Dorr and Voss, 1993; Dorr, 1993]. Domain specific knowledge can be used to constrain this task for particular applications, which is the approach taken in *PICTION*.

5 *PICTION*: An Application of the Model

Figure 2. illustrates the application of the model in *PICTION*. This section describes the process of spatial expression understanding in *PICTION*, following the control flow shown in Figure 2. First, the domain specific lexicon is described.

5.1 Lexicon

A small set of lexicon entries including humans, rectangular objects and spatial verb arguments (see above section) has been manually augmented with spatial information. Systematic extraction of spatially relevant information from machine readable sources is an area of current research [Dolan, 1994]. Until the process of spatially annotating words can be automated, spatial information can only be represented for a small number of words.

Nouns: Spatial information for nouns is represented using an object schema associated with the lexical entry for the noun. The schema is designed to include default values for spatial and visual properties that may be overridden by scene specific information. This is similar to the object schema used by Lang *et al.* [1991].

For example, the object schema for "hat" is:

```
geometry: polygon
typical-location: head
```

Verbs: Our approach to verb representation is to decompose the meaning of a complex spatial verb into a combination of spatial prepositions and functions.

A partial representation of the verb "wear" is:

```
VERB FRAME:
person wear object
  [optional prepositional phrase]

MEANING:
on(object,
    body-part(person,
     get-typical-location
      (object)))
```

Wear means that the object being worn is located on the body of the wearer. Specifically, it is located on the body part returned by the get-typical-location function, which retrieves the value from the typical-location slot in the object's schema. If the optional prepositional phrase is present, the information it conveys will override this default value.

Related work in verb representation is found in Siskind [1994]. He observed that principles of support, contact, and attachment identified as important by Herskovits [1985; 1986] in understanding spatial prepositions are equally important in understanding many spatial motion verbs. While Siskind has abstracted these properties away from prepositions, we have employed spatial

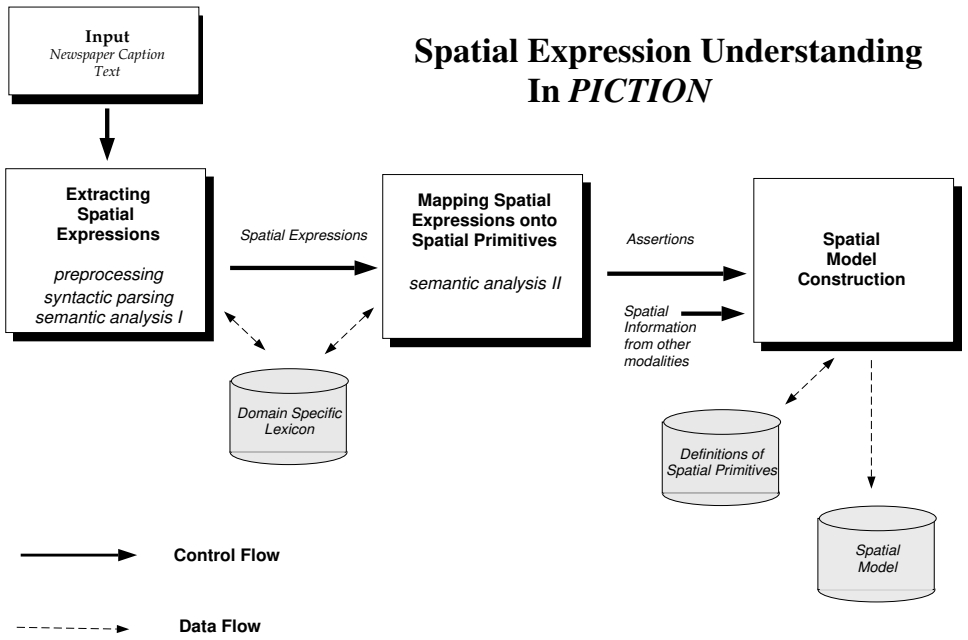

Figure 2: Spatial Expression Understanding in *PICTION*

prepositions directly in our definitions. This is sufficient in *PICTION*, because temporal and motion information is not required.

Chakravarthy [1994] discusses the lack of perceptual information in traditional verb semantics. His lexical representation of verbs has two principal components: the first is similar to an object schema, and the second is an action definition. The lack of temporal information in our domain obviates the need for an action definition, but his object schema is similar to our approach. Our domain allows us to employ a small number of primitives for defining spatial verbs. In general, this approach succumbs to the same problems he notes in Schank's work. In our domain, however, it has thus far been sufficient. While 135 of the captions described above contain spatial verbs, most of these were drawn from a set of about 10 verbs that includes *walk, flank, join, hold, wear* and *sit*.

Adjectives: We are not currently modeling spatial information for adjectives in *PICTION*, and our current data indicate it may not be important in this domain.

Prepositions: Spatial prepositions in *PICTION* are geometrically defined predicates used to verify relationships between rectangular boxes (the visual representation of a face is a *bounding box* around the actual face) in two dimensions. Spatial relations from the text are mapped onto the following set of spatial primitives: *left-of, right-of, above, below, strict-left-of, strict-right-of, strict-above, strict-below, between, surrounded-by and near*. Definitions of the spatial primitives are discussed below.

5.2 Extracting Spatial Expressions from Captions

Our approach to extracting spatial expressions from captions involves three stages: preprocessing, syntactic parsing, and semantic analysis I [Srihari and Burhans, 1994].

Preprocessing

The goal of preprocessing is to mark relative clauses containing spatial terms, and to identify and group proper nouns. Spatial terms are those which Dorr and Voss [1993] term *lexically explicit* spatial relations. An example exhibiting a relative clause containing spatial information and a proper noun complex is *Rep. John Lewis, D-Ga., left foreground.* Rules that perform keyword searches for spatial terms, and make use of punctuation to identify relative clauses, will tag *left foreground* as a phrase containing spatial information. Proper noun identification and classification is important in the newspaper caption domain, although we are not currently pursuing this as a research area. *PICTION* uses simple heuristics and multiple-word lexicon entries to recognize proper nouns in the input text. *Rep. John Lewis, D-Ga.* is a proper noun complex that refers to a single individual.

Syntactic Parsing

Syntactic parsing is performed on the preprocessor output. A lexical function grammar that covers basic features of English is used in conjunction with the Tomita parser [1987] to perform syntactic parsing. Certain semantic features (*e.g.*, animate/inanimate) that can be obtained through lexicon lookup are incorporated into the output structure of the parser.

Semantic Analysis I

Initial semantic analysis of a syntactically parsed caption encompasses the following tasks:

- identifying spatial information not marked by the preprocessor

- mapping complex spatial expressions (implicit spatial information) onto simple spatial expressions
- determining what objects mentioned in the caption are present in the accompanying image, and for each such object, identifying and associating all references to that object

The first task is accomplished using spatial and object class information in the lexicon. Information contained in the lexicon entries for complex spatial terms is used to rewrite these terms in a lexically explicit manner. Predicting which objects are present in the associated image is based on sentence structure and heuristics [Srihari, 1995]. A small grammar that includes domain specific information about object classes, object properties, and the syntactic parse structure is employed in this task. The problem of reference in *PICTION* is relatively simple, because we only process single sentence captions.

The output of semantic analysis I is a list of objects (and their aliases), and spatial relations: these constitute the spatial expressions that must be mapped onto spatial primitives.

5.3 Mapping Spatial Expressions onto Spatial Primitives

Semantic analysis II refers to the process of mapping simple spatial expressions and objects onto the set of spatial primitives. These spatial primitives are then expressed as a set of constraints: assertions in *PICTION* are represented as sets of constraints on a spatial scene [Chopra and Srihari, 1995].

First, objects are associated with picture objects and are mapped to *in-picture* constraints, which assert the existence of an object in the accompanying photograph.

Next, spatial relations are mapped onto spatial primitives. For example, *beside* and *next to* are used to indicate proximity in captions, but are not spatial primitives. They are both mapped onto the spatial primitive *near*. *Left of* is an example of a spatial relation that maps onto itself.

Finally, the picture objects and the primitive spatial relations in which they participate are mapped onto a set of *spatial* constraints.

The output from the spatial expression mapping phase is a set of assertions representing the *spatial* and *in-picture* constraints on a spatial scene.

5.4 Spatial Model Construction

PICTION performs coreferencing of the caption information with its accompanying photographic data in order to construct a spatial model of a scene. Simplifying assumptions about the shape representations of objects, as well as the domain specific fixed frame of reference and the pre-selected set of spatial primitives, facilitate the integration of linguistic and visual information.

Shape representation: *Human-face* objects representing people asserted to be present in the photograph are represented as simple upright bounding boxes in two–dimensional space. Rectangular objects are similarly represented. The reference point of the objects is the centroid of the box.

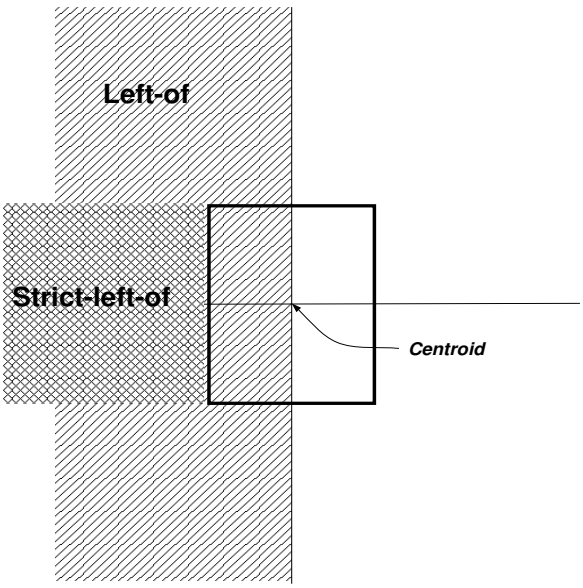

Figure 3: Interpretation of binary spatial primitives: The textured area marks out the acceptable positions for the reference point of the other object

Frame of reference: The image coordinate space fixes the single frame of reference. All the spatial primitives are interpreted relative to the viewpoint of the newspaper reader.

Spatial primitives: The set of spatial primitives that we have found to be adequate for modeling spatial relations in this domain are: *left-of, right-of, above, below, strict-left-of, strict-right-of, strict-above, strict-below, between, surrounded-by and near.*

The goal of *PICTION* is to coreference visual entities with objects derived from captions. This allows spatial primitives to be implemented as predicates that verify spatial relations rather than functions that generate optimal spatial configurations. Based on the assumptions specified above, spatial primitives are quantified in the following manner.

The first eight spatial primitives are simple binary relations that are verified using the centroid of the bounding box as the reference point (Fig. 3). For example,

$$left\text{–}of(Obj_a, Obj_b) = TRUE$$

if the x_value of the reference point of Obj_a is less than the x_value of the reference point of Obj_b. However,

$$strict\text{–}left\text{–}of(Obj_a, Obj_b) = TRUE$$

if in addition to the above condition, the y_value of the reference point of Obj_a is between the minimum and maximum y_values of the extent of Obj_b. In Fig. 3, this corresponds to the centroid lying in the horizontal band to the left of the box. Overlap of objects does not affect the interpretations of spatial primitives in this domain.

The definitions of *between* and *surrounded-by* are based on euclidian distance. For example:

$$surrounded\text{–}by(Obj_a, Obj_{b_1} \ldots Obj_{b_n}) = TRUE$$

if

$$\forall j \sum_{i=1}^{n} dist(Obj_a, Obj_{b_i})$$
$$< \quad dist(Obj_a, Obj_{b_j}) \quad +$$
$$\sum_{i=1, i \neq j}^{n} dist(Obj_{b_j}, Obj_{b_i})$$

In other words, if the sum of the distances between Obj_a and the other objects is less than the the sum of the distances between any Obj_{b_j} and the other objects, then Obj_a is the subject of the *surrounded-by* predicate. The primitive *between* can be viewed as a ternary specialization of *surrounded-by*.

Near is often interpreted qualitatively [Abella and Kender, 1993]. In our domain, an object is *near* another object if there is no other object *nearer* the latter.[1] Verifying *near* in our domain entails computing relationships between all the objects. If we defined the set $M = \{AllObjects - \{Obj_a, Obj_b\}\}$, then $near(Obj_a, Obj_b, M) = TRUE$ **if** the distance between Obj_a and Obj_b is less than the distance between Obj_b and any object in the set M.

6 Discussion

We have proposed a model that we believe will be useful in building domain specific systems for understanding spatial expressions. We have described *PICTION* as an application of this model.

We have shown that understanding spatial expressions involves identifying spatial information expressed in both the open and closed classes, and that the problem of identifying and representing open class spatial information may be simplified when background assumptions are taken into account. We have described a way of representing spatial nouns and verbs that reflects our observations about spatial language in the newspaper caption domain.

Areas of future work for *PICTION* include expanding the coverage of spatial terms, looking at a much larger data set, and formalizing some of the heuristics currently used in the system.

The model can be applied to the domain of machine translation as follows: the extraction of spatial information from the input text corresponds to the first task we identified. Mapping spatial expressions onto spatial primitives involves constructing a language independent representation of the spatial information. The spatial model corresponds directly to this interlingua, whose spatial terms comprise the spatial primitives. Translating from an interlingua representation to a target language mirrors the problem of coreferencing. In such a task, spatial reaoning is not generally necessary.

Work done by Olivier [1994] in representing spatial expressions has focused on developing complex, three–dimensional definitions for spatial prepositions. In this case, construction of the spatial model, definitions of spatial primitives, and spatial reasoning are the focus of the system, while the complexity of going from input text to spatial primitives is minimized.

Much of the research in spatial expression understanding makes use of background assumptions, and focuses on one or two of the tasks highlighted in the model we have described. We have attempted to show that this is not necessarily limiting, but instead can lead to success in a clearly defined domain.

[1] In the case of multiple object types, an object A is near another object B if there is no other object C, of the same class as A, nearer to B.

References

[Abella and Kender, 1993] Alicia Abella and John R. Kender. Qualitatively Describing Objects Using Spatial Prepositions. In *Proceedings of AAAI-93*, pages 536–540, 1993.

[Beckwith *et al.*, 1991] R. Beckwith, C. Fellbaum, D. Gross, and G. A. Miller. WordNet: A Lexical Database Organized on Psycholinguistic Principles. In *Lexicons: Using On-line Resources to Build a Lexicon*. Lawrence Erlbaum, 1991.

[Chakravarthy, 1994] Anil S. Chakravarthy. Towards a Perceptually-Based Semantics of Action Verbs. In *AAAI-94 Workshop on the Integration of Natural Language and Vision Processing*, pages 161–164, 1994.

[Chopra and Srihari, 1995] Rajiv Chopra and Rohini Srihari. Control structures for incorporating picture-specific context in image interpretation. In *IJCAI-95*, Montreal,Canada, August 1995. Morgan Kaufmann.

[Dolan, 1994] William B. Dolan. Exploiting Lexical Information for Visual Processing. In *AAAI-94 Workshop on the Integration of Natural Language and Vision Processing*, pages 185–188, 1994.

[Dorr and Voss, 1993] Bonnie J. Dorr and Clare R. Voss. Machine Translation of Spatial Expressions: Defining the Relation between an Interlingua and a Knowledge Representation System. In *Proceedings of AAAI-93*, pages 374–379, 1993.

[Dorr, 1993] Bonnie J. Dorr. *Machine Translation: A View from the Lexicon*. MIT Press, Cambridge, Mass., 1993.

[Herskovits, 1985] Annette Herskovits. Semantics and Pragmatics of Locative Expressions. *Cognitive Science*, 9:341–378, 1985.

[Herskovits, 1986] Annette Herskovits. *Language and Spatial Cognition*. Cambridge University Press, 1986.

[Lang *et al.*, 1991] Ewald Lang, Kai-Uwe Carstensen, and Geoffrey Simmons. *Modelling Spatial Knowledge on a Linguistic Basis*. Lecture Notes in Computer Science. Springer-Verlag, 1991.

[Olivier, 1994] Patrick Olivier. Prepositional Semantics in the WINTOP System. In *AAAI-94 Workshop on the Integration of Natural Language and Vision Processing*, pages 139–144, 1994.

[Siskind, 1994] Jeffrey Mark Siskind. Axiomatic Support for Event Perception. In *AAAI-94 Workshop on the Integration of Natural Language and Vision Processing*, pages 153–160, 1994.

[Srihari and Burhans, 1994] Rohini K. Srihari and Debra T. Burhans. Visual Semantics: Extracting Visual Information from Text Accompanying Pictures. In *Proceedings of AAAI-94*, pages 793–798. AAAI Press, 1994.

[Srihari, 1995] Rohini K. Srihari. Use of Collateral Text in Understanding Photos. In Press. *Artificial Intelligence Review*, 8(5-6), 1995. Special Issue on Integrating Language and Vision.

[Talmy, 1985] Leonard Talmy. Lexicalization patterns. In Timothy Shopen, editor, *Grammatical Categories and the Lexicon*, pages 57–149. Cambridge University Press, 1985.

[Tomita, 1987] M. Tomita. An Efficient Augmented-Context-Free Parsing Algorithm. *Computational Linguistics*, 13(1-2):31–46, 1987.

[Zernik and Vivier, 1988] Uri Zernik and Barbara J. Vivier. How Near Is Too Far? Talking about Visual Images. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, pages 202–208. Lawrence Erlbaum Associates, 1988.